

# Oracle® Cloud

## Exadata Database Service on Exascale Infrastructure



F92371-17  
March 2025



Oracle Cloud Exadata Database Service on Exascale Infrastructure,

F92371-17

Copyright © 2023, 2025, Oracle and/or its affiliates.

Primary Authors: Nirmal Kumar, Doug Williams

Contributors: Ravi Chennoju, Bryce Cracco, Yue Deng, Gayathri Dubagunta, Anil Kothuri, Natrajan Krishnamoorthi, Shuo Li, Ranjit Murali, Sanjay Narvekar, Babak Sanaee, Peter Sciarra, Lokesh Shrivastava, Nihal Shrivastava, Jason Tang, Bhuvanewari Thiagarajan, Sujeet Vasudevan, Nayana Vishwa, Ravi Wijayaratne

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

<b>1</b>	<b>Oracle Exadata Database Service on Exascale Infrastructure Overview</b>	
	About Oracle Exadata Database Service on Exascale Infrastructure	1-1
	Accessing the Exadata Database Service on Exascale Infrastructure Using the OCI Console	1-4
	Licensing Considerations for Oracle Exadata Database Service on Exascale Infrastructure	1-4
	Supported Database Edition and Versions for Oracle Exadata Database Service on Exascale Infrastructure	1-5
	Subscription Types	1-5
	Service Limits for Exadata Database Service on Exascale Infrastructure	1-6
	Metering Frequency and Per-Second Billing	1-6
	Exadata Cloud Management Interfaces	1-6
	Introduction to Exadata Cloud Management Interfaces	1-7
	OCI Control Plane Interfaces for Oracle Exadata Database Service on Exascale Infrastructure	1-7
	Local VM Command-Line Interfaces	1-9
<b>2</b>	<b>What's New in Oracle Exadata Database Service on Exascale Infrastructure</b>	
	New Regions and Realms for ExaDB-XS	2-1
	Scale ECPUs to Zero	2-3
	Deploy Single Node VM Clusters	2-3
<b>3</b>	<b>Preparing for Oracle Exadata Database Service on Exascale Infrastructure</b>	
	Oracle Cloud Infrastructure (OCI) Requirements for Oracle Exadata Database Service on Exascale Infrastructure	3-1
	Required IAM Policy for Oracle Exadata Database Service on Exascale Infrastructure	3-2
	Network Setup for Oracle Exadata Database Service on Exascale Infrastructure Instances	3-3
	VCN and Subnets	3-3
	Option 1: Public Client Subnet with Internet Gateway	3-4
	Option 2: Private Subnets	3-6
	Requirements for IP Address Space	3-7
	Configuring a Static Route for Accessing the Object Store	3-8

Setting Up DNS for an Oracle Exadata Database Service on Exascale Infrastructure Instance	3-8
DNS: Short Names for the VCN, Subnets, and Oracle Exadata Database Service on Exascale Infrastructure instance	3-8
Configure Private DNS	3-9
Node Access to Object Storage: Static Route	3-10
Object Storage IP allocations	3-10
To configure a static route for Object Storage access	3-11
Service Gateway for the VCN	3-11
Option 1: Service Gateway Access to OCI Services	3-12
Option 2: Service Gateway Access to Both Object Storage and YUM Repos	3-13
Security Rules for the Oracle Exadata Database Service on Exascale Infrastructure	3-14
Rules Required for Both the Client Network and Backup Network	3-18
Rules Required Specifically for the Client Network	3-19
Rule Required Specifically for the Backup Network	3-21
Rules Required for Events Service	3-22
Rules Required for Monitoring Service	3-22
Ways to Implement the Security Rules	3-22
If you use network security groups	3-23
If you use security lists	3-23
Network Requirements for Oracle Database Autonomous Recovery Service	3-24
Create a Service Gateway to Object Storage	3-24

## 4 Getting Started with Oracle Exadata Database Service on Exascale Infrastructure Deployment

---

Tagging Oracle Exadata Database Service on Exascale Infrastructure Resources	4-1
Restarting a VM for Planned Maintenance	4-5
Connecting to an Oracle Exadata Database Service on Exascale Infrastructure VM	4-6
Prerequisites for Accessing Oracle Exadata Database Service on Exascale Infrastructure	4-6
SCAN Listener Port Setting	4-7
Connecting to a Virtual Machine with SSH	4-7
Connecting from a Unix-Style System	4-8
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY	4-8
Accessing a Database After You Connect to the Virtual Machine	4-9
Using Oracle Net Services to Connect to a Database	4-11
Prerequisites for Connecting to a Database with Oracle Net Services	4-11
Connecting to a Database with SQL Developer	4-12
Connecting to a Database Using SCAN	4-13
Connecting to a Database Using a Node Listener	4-15
Connect to the Oracle Exadata Database Service on Exascale Infrastructure Service	4-16
Connecting to a Database with SQL Developer	4-16

Connecting to a Database with Oracle Net Services	4-17
Capacity Limits for Exadata Database Service on Exascale Infrastructure	4-21
Best Practices for Oracle Exadata Database Service on Exascale Infrastructure VMs	4-22
Moving to Oracle Cloud Using Zero Downtime Migration	4-23

## 5 How-to Guides

---

Manage Database Security with Oracle Data Safe	5-2
About Oracle Data Safe	5-2
Get Started	5-3
Using Oracle Data Safe	5-3
Connecting to an Oracle Exadata Database Service on Exascale Infrastructure VM	5-5
Connection Prerequisites	5-6
About Connecting to a VM with SSH	5-6
Connecting from a Unix-Style System	5-6
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY	5-7
To access a database after you connect to the VM	5-8
Connect to the Oracle Exadata Database Service on Exascale Infrastructure Service	5-8
Connecting to a Database with SQL Developer	5-9
Connecting to a Database with Oracle Net Services	5-9
Manage Oracle Exadata Database Service on Exascale Infrastructure	5-13
Using the Console to Provision Oracle Exadata Database Service on Exascale Infrastructure	5-14
Lifecycle Management Operations	5-14
Network Management Operations	5-15
Management Tasks for the Oracle Cloud Infrastructure Platform	5-16
Oracle Database License Management Tasks	5-18
Using the API to Create Infrastructure Components	5-19
Using the API to Manage Oracle Exadata Database Service on Exascale Infrastructure Instance	5-19
Manage VM Clusters	5-20
Using the Console to Manage VM Clusters on Oracle Exadata Database Service on Exascale Infrastructure	5-21
To create a cloud VM cluster	5-22
Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection	5-27
Using the Console to Update the License Type on a VM Cluster	5-28
To scale VM Clusters	5-28
To add SSH keys to a VM cluster	5-30
Using the Console to Add SSH Keys After Creating a VM Cluster	5-30
Using the Console to Stop, Start, or Reboot a VM Cluster Virtual Machine	5-31
Using the Console to Check the Status of a VM Cluster Virtual Machine	5-31
Using the Console to Move a VM Cluster to Another Compartment	5-32
To change the VM cluster display name	5-32

Using the Console to Terminate a VM Cluster	5-33
To view details about private DNS configuration	5-33
Adding or Removing a VM From a VM Cluster	5-34
Add a VM to a VM Cluster	5-34
Terminate a VM from a VM Cluster	5-35
Overview of Automatic Diagnostic Collection	5-35
Incident Logs and Trace Files	5-36
Health Metrics	5-40
Using the API to Manage Oracle Exadata Database Service on Exascale Infrastructure Instance	5-43
Manage Exascale Database Vaults on Exadata Database Service on Exascale Infrastructure	5-44
Manage Software Images	5-45
Using Software Images in Oracle Cloud Infrastructure	5-45
Creation and Storage of Software Images	5-45
Using the OPatch Isinventory Command to Verify the Patches Applied to an Oracle Home	5-46
Using a Software Image with an Exadata Cloud Infrastructure Instance	5-47
Using the Console for Software Images	5-47
To create a database software image	5-47
To create a database software image from a Database Home	5-48
To update database software using custom database software image	5-48
To delete a software image	5-49
Using the API to manage database software images	5-49
Create Oracle Database Homes on an Oracle Exadata Database Service on Exascale Infrastructure System	5-50
About Creating Oracle Database Homes on an Oracle Exadata Database Service on Exascale Infrastructure System	5-50
To create a new Database Home in an existing Oracle Exadata Database Service on Exascale Infrastructure instance	5-51
To create a database software image from a Database Home	5-52
Using the API to Create Oracle Database Home on Oracle Exadata Database Service on Exascale Infrastructure	5-52
Managing Oracle Database Homes on an Oracle Exadata Database Service on Exascale Infrastructure Instance	5-53
Manage Database Home Using the Console	5-53
To view information about a Database Home	5-53
To delete a database home	5-54
To manage tags for your Database Home	5-54
Using the Console to Move a Database to Another Database Home	5-55
Using the API to Manage Oracle Database Home on Oracle Exadata Database Service on Exascale Infrastructure	5-55
Manage Databases on Oracle Exadata Database Service on Exascale Infrastructure	5-56
Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Exascale Infrastructure	5-56

Oracle Database Releases Supported by Oracle Exadata Database Service on Exascale Infrastructure	5-57
Provisioning and Managing Exadata Databases	5-57
Database Memory Initialization Parameters	5-58
Customer-Managed Keys in Oracle Exadata Database Service on Exascale Infrastructure	5-58
Using the Console to Manage Databases on Oracle Exadata Database Service on Exascale Infrastructure	5-60
Using the API to manage Databases	5-76
Create and Manage Exadata Pluggable Databases	5-76
Limitations for Pluggable Database Management	5-78
Creating an Exadata Pluggable Database	5-78
Managing an Exadata Pluggable Database	5-82
Cloning an Exadata Pluggable Database	5-85
Restoring an Exadata Pluggable Database	5-92
Using the Console to Perform an In-Place Restore of a Pluggable Database (PDB)	5-92
Using the Console to Perform an Out-of-Place Restore of a Pluggable Database (PDB)	5-93
Changing the Database Passwords	5-93
To Change the SYS Password for an Oracle Exadata Database Service on Exascale Infrastructure Database	5-94
To Change Database Passwords in a Data Guard Environment	5-94
To Change the TDE Wallet Password for an Oracle Exadata Database Service on Exascale Infrastructure Database	5-94
Manage Database Backup and Recovery on Oracle Exadata Database Service on Exascale Infrastructure	5-94
Oracle Recommended Options to Perform Backup and Recovery Operations	5-95
Managing Exadata Database Backups	5-97
Managed Backup Types and Usage Information	5-97
Default Backup Channel Allocation	5-99
Prerequisites for Backups on Oracle Exadata Database Service on Exascale Infrastructure	5-99
Using the Console to Manage Backups	5-101
To configure automatic backups for a database	5-101
To create an on-demand backup of a database	5-104
To view backup status	5-104
To cancel a backup	5-104
To delete full backups from Object Storage	5-105
To delete standalone backups from Object Storage	5-106
To designate Autonomous Recovery Service as a Backup Destination for an Existing Database	5-106
Recovering an Exadata Database from Backup Destination	5-107
Using the Console to restore a database	5-107
Managing Exadata Database Backups by Using dbaascli	5-108

Default Backup Configuration	5-109
To create a backup configuration file	5-109
To create an on-demand backup	5-112
To remove the backup configuration	5-113
To delete a local backup	5-113
To delete a backup in Object Storage	5-114
Using the API to Manage Backup and Recovery	5-115
Using the API to manage backups	5-115
Alternative Backup Methods	5-115
Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management	5-116
Recovering a Database Using Oracle Recovery Manager (RMAN)	5-119
Patch and Update an Oracle Exadata Database Service on Exascale Infrastructure System	5-119
User-Managed Maintenance Updates	5-119
Patching and Updating an Oracle Exadata Database Service on Exascale Infrastructure System	5-120
Patching and Updating VM Cluster's GI and Database Homes	5-120
Updating an Exadata Cloud VM Cluster Operating System	5-137
Upgrading Exadata Databases	5-139
Manual Software Updates	5-144
Create Software Update	5-144
Download an Interim Software Update	5-145
Delete an Interim Software Update	5-146
Move an Interim Software Update Resource to Another Compartment	5-146
Using the API to Manage Interim Software Updates	5-146
Use Oracle Data Guard with Oracle Exadata Database Service on Exascale Infrastructure	5-147
About Using Oracle Data Guard with Oracle Exadata Database Service on Exascale Infrastructure	5-147
Prerequisites for Using Oracle Data Guard with Oracle Exadata Database Service on Exascale Infrastructure	5-148
Network Requirements for Data Guard	5-148
Password Requirements	5-150
Known Issues for Exadata Cloud Infrastructure and Data Guard	5-150
Adding a Node to a VM Cluster	5-151
Removing a Node from a VM Cluster	5-151
Working with Oracle Data Guard	5-151
Switchover	5-152
Failover	5-152
Reinstate	5-152
Using the Console to Manage Oracle Data Guard Associations	5-152
To enable Data Guard on Exadata Database Service on Exascale Infrastructure	5-153
To view Data Guard associations of databases in a Cloud VM Cluster	5-156
To enable automatic backups on a standby database	5-156



To perform a database switchover	5-157
To edit the Oracle Data Guard association	5-158
To perform a database failover	5-159
To reinstate a database	5-159
To terminate a Data Guard association on an Oracle Exadata Database Service on Exascale Infrastructure instance	5-159
Using the API to manage Data Guard associations	5-160
Configure Oracle Database Features for Oracle Exadata Database Service on Exascale Infrastructure	5-160
Using Oracle Multitenant on an Oracle Exadata Database Service on Exascale Infrastructure Instance	5-161
To determine if you need to create and activate an encryption key for the PDB	5-161
To create and activate the master encryption key in a PDB	5-162
To export and import a master encryption key	5-164
Managing Tablespace Encryption	5-164
Migrate to Oracle Exadata Database Service on Exascale Infrastructure	5-166
Connect Identity and Access Management (IAM) Users to Oracle Exadata Database Service on Exascale Infrastructure	5-166
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database	5-167
About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database	5-167
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication	5-168
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication	5-169
Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database	5-171
Prerequisites for IAM Authentication on Oracle Database	5-171
Disable External Authentication Scheme	5-172
Configure TLS to Use IAM Tokens	5-172
Enable, Disable, and Re-enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database	5-173
Enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database	5-174
Disable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database	5-174
Using Oracle Database Tools with Identity and Access Management (IAM) Authentication	5-175
Manage Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies, Users, Roles, and Database Passwords	5-175
Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies for IAM Users	5-176
Authorize Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users on Oracle Database	5-177
To Exclusively Map a Local IAM User to an Oracle Database Global User	5-178

Add Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Roles on Oracle Database	5-178
Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password for IAM Users	5-179
Configuring Client Connection	5-179
Configure a Client Connection for SQL*Plus that Uses an IAM Database Password Verifier	5-180
Configure Client Connection for SQL*Plus that Uses an IAM Token	5-180
Client Connections That Use a Token Requested by an IAM User Name and Database Password	5-182
Use Instance Principal to Access Database with IAM Authentication	5-183
Configure Proxy Authentication	5-183
Use Database Link with IAM Authenticated Users	5-184

## 6 Reference Guides for Oracle Exadata Database Service on Exascale Infrastructure

---

Using the dbaascli Utility on Oracle Exadata Database Service on Exascale Infrastructure	6-1
About Using the dbaascli Utility on Oracle Exadata Database Service on Exascale Infrastructure	6-2
Creating Databases Using dbaascli	6-3
Listing Available Software Images and Versions for Database and Grid Infrastructure	6-3
Creating Oracle Database Home	6-4
Creating Oracle Database In the Specified Oracle Database Home	6-5
Changing the Database Passwords	6-8
To Change the SYS Password for an Oracle Exadata Database Service on Exascale Infrastructure Database	6-8
To Change Database Passwords in a Data Guard Environment	6-9
To Change the TDE Wallet Password for an Oracle Exadata Database Service on Exascale Infrastructure Database	6-9
Managing Oracle Exadata Database Service on Exascale Infrastructure Software Images Using the Dbaascli Utility	6-9
Listing Available Software Images and Versions for Database and Grid Infrastructure	6-10
To download a software image	6-11
Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check Using dbaascli	6-12
Collecting Tooling Log Data Examples	6-12
Performing a Health Check Examples	6-16
Updating Cloud Tooling Using dbaascli	6-17
Creating a Duplicate Database	6-18
Using dbaascli to Duplicate a Cloud Database	6-18
Considerations When Using OCI Vault for the Key Management	6-19
dbaascli Command Reference	6-20
dbaascli admin updateStack	6-25
dbaascli cswlib deleteLocal	6-27

dbaascli cswlib download	6-28
dbaascli cswlib listLocal	6-30
dbaascli cswlib showImages	6-31
dbaascli database addInstance	6-33
dbaascli database backup	6-34
dbaascli database bounce	6-38
dbaascli database changepassword	6-40
dbaascli database convertToPDB	6-42
dbaascli database create	6-43
dbaascli database delete	6-46
dbaascli database deleteInstance	6-48
dbaascli database duplicate	6-50
dbaascli database getDetails	6-53
dbaascli database getPDBs	6-54
dbaascli database modifyParameters	6-55
dbaascli database move	6-56
dbaascli database recover	6-64
dbaascli database runDatapatch	6-66
dbaascli database createTemplate	6-68
dbaascli database start	6-68
dbaascli database status	6-69
dbaascli database stop	6-70
dbaascli database upgrade	6-71
dbaascli dataguard prepareStandbyBlob	6-73
dbaascli dataguard updateDGConfigAttributes	6-75
dbaascli dbhome create	6-76
dbaascli dbHome delete	6-78
dbaascli dbhome getDatabases	6-80
dbaascli dbHome getDetails	6-82
dbaascli dbHome patch	6-84
dbaascli dbimage purge	6-92
dbaascli diag collect	6-92
dbaascli diag healthCheck	6-93
dbaascli grid configureTCPS	6-95
dbaascli grid patch	6-99
dbaascli grid removeTCPSCert	6-106
dbaascli grid rotateTCPSCert	6-109
dbaascli grid upgrade	6-111
dbaascli job getStatus	6-113
dbaascli patch db apply	6-114
dbaascli patch db prereq	6-114
dbaascli pdb backup	6-115

dbaascli pdb bounce	6-117
dbaascli pdb close	6-119
dbaascli pdb getConnectString	6-121
dbaascli pdb create	6-122
dbaascli pdb createSnapshot	6-125
dbaascli pdb configureSnapshot	6-126
dbaascli pdb delete	6-127
dbaascli pdb deleteSnapshot	6-129
dbaascli pdb getDetails	6-130
dbaascli pdb getSnapshot	6-131
dbaascli pdb list	6-132
dbaascli pdb listSnapshots	6-133
dbaascli pdb localClone	6-134
dbaascli pdb open	6-136
dbaascli pdb recover	6-138
dbaascli pdb refresh	6-140
dbaascli pdb relocate	6-142
dbaascli pdb remoteClone	6-145
dbaascli system getDBHomes	6-150
dbaascli tde changePassword	6-151
dbaascli tde addSecondaryHsmKey	6-153
dbaascli tde enableWalletRoot	6-155
dbaascli tde encryptTablespacesInPDB	6-158
dbaascli tde fileToHsm	6-160
dbaascli tde getHsmKeys	6-162
dbaascli tde getMkidForKeyVersionOCID	6-164
dbaascli tde getPrimaryHsmKey	6-166
dbaascli tde hsmToFile	6-168
dbaascli tde listKeys	6-170
dbaascli tde removeSecondaryHsmKey	6-172
dbaascli tde rotateMasterKey	6-175
dbaascli tde setKeyVersion	6-178
dbaascli tde setPrimaryHsmKey	6-181
dbaascli tde status	6-184
Database Service Events	6-186
Overview of Database Service Events	6-186
Monitor Metrics for VM Cluster Resources	6-189
View Metrics for VM Cluster	6-189
View Metrics for a Database	6-190
View Metrics for VM Clusters in a Compartment	6-192
View Metrics for Databases in a Compartment	6-193
Manage Oracle Trace File Analyzer	6-194

Manage Database Service Agent	6-195
Metrics for Oracle Exadata Database Service on Exascale Infrastructure in the Monitoring Service	6-195
Oracle Exadata Database Service on Exascale Infrastructure Events	6-199
About Event Types on Oracle Exadata Database Service on Exascale Infrastructure	6-200
Prerequisites for Event Service	6-200
Oracle Exadata Database Service on Exascale Infrastructure Event Types	6-201
Oracle Exadata Database Service on Exascale Infrastructure Maintenance Event Types	6-204
Exadata Cloud Infrastructure Critical and Information Event Types	6-210
Exadata Cloud Infrastructure VM Cluster Event Types	6-214
VM Node Subsetting Event Types	6-226
Data Guard Association Event Types	6-230
Oracle Database Home Event Types	6-230
Database Event Types	6-231
Pluggable Database Event Types	6-233
Database Service Events	6-240
Overview of Database Service Events	6-241
Receive Notifications about Database Service Events	6-243
Database Service Event Types	6-244
Temporarily Restrict Automatic Diagnostic Collections for Specific Events	6-252
Application VIP Event Types	6-255
Interim Software Updates Event Types	6-267
Serial Console Connection Event Types	6-273
Viewing Audit Log Events	6-278
Monitor Metrics to Diagnose and Troubleshoot Problems with Pluggable Databases	6-279
About Database Management	6-280
Using the Console to Enable Database Management for a Container Database (CDB)	6-280
Enable Database Management	6-281
Using the Console to Enable Database Management for a Pluggable Database (PDB)	6-283
Enable Database Management	6-284
Using the Console to Edit Database Management for a Pluggable Database (PDB)	6-286
Edit Database Management	6-287
Using the Console to Disable Database Management for a Pluggable Database (PDB)	6-289
Using the Console to View Performance Hub for a Container Database (CDB)	6-289
Using the Console to View Performance Hub for a Pluggable Database (PDB)	6-290
Using the API to Enable, Disable, or Update Database Management Service	6-291
Oracle Cloud Database Metrics	6-291
Using the Console View Metrics for a Container Database (CDB)	6-291
Using the Console to View Metrics for a Pluggable Database (PDB)	6-292
Policy Details for Oracle Exadata Database Service on Exascale Infrastructure	6-293
About Resource-Types	6-294
Resource-Types for Exadata Cloud Service Instances	6-294

Supported Variables	6-294
Details for Verb + Resource-Type Combinations	6-294
Database-Family Resource Types	6-295
Permissions and API operation details for DB Backups	6-295
Permissions and API operation details for Databases (CDBs)	6-296
Permissions and API operation details for Data Guard Association	6-297
Permissions and API operation details for DB Nodes	6-297
Permissions and API operation details for DB Homes	6-297
Permissions and API operation details for Database Software Image	6-299
exadb-vm-clusters	6-299
exascale-db-storage-vaults	6-300
Permissions and API operation details for Key Stores	6-301
Permissions Required for Each API Operation	6-302
Permissions and API operation details for Pluggable Databases (PDBs)	6-306
Oracle Cloud Infrastructure Operations Insights	6-308
Managing Exadata Resources with Oracle Enterprise Manager Cloud Control	6-309
Overview of Oracle Enterprise Manager Cloud Control	6-309
Features of Enterprise Manager Cloud Control	6-310
Analyzing Exadata Database Service Database Performance	6-310
Troubleshooting Oracle Exadata Database Service on Exascale Infrastructure Systems	6-311
Known Issues for Exadata Database Service on Exascale Infrastructure	6-311
Troubleshooting Oracle Data Guard	6-311
Troubleshooting Data Guard using logfiles	6-312
Troubleshooting the Data Guard Setup Process	6-314
Obtaining Further Assistance	6-317
Collecting Cloud Tooling Logs	6-317
Collecting Oracle Diagnostics	6-317

# 1

## Oracle Exadata Database Service on Exascale Infrastructure Overview

This topic is an overview of the Oracle Exadata Database Service on Exascale Infrastructure formerly Exadata Cloud Service.

- [About Oracle Exadata Database Service on Exascale Infrastructure](#)
- [Accessing the Exadata Database Service on Exascale Infrastructure Using the OCI Console](#)  
Learn how to access the Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS) service.
- [Licensing Considerations for Oracle Exadata Database Service on Exascale Infrastructure](#)  
Subscription to Oracle Exadata Database Service on Exascale 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 Exascale Infrastructure.
- [Supported Database Edition and Versions for Oracle Exadata Database Service on Exascale Infrastructure](#)  
Oracle Exadata Database Service on Exascale Infrastructure databases require Enterprise Edition - Extreme Performance subscriptions or you can bring your own Oracle Enterprise Edition software licenses.
- [Subscription Types](#)  
Learn about the available subscription types for Oracle Exadata Database Service on Exascale Infrastructure
- [Service Limits for Exadata Database Service on Exascale Infrastructure](#)  
Limits apply to virtual machine (VM) instance counts, total ECPU count, total local storage, and total High Capacity storage.
- [Metering Frequency and Per-Second Billing](#)  
See the Per-Second billing, minimums, and limitations on billing.
- [Exadata Cloud Management Interfaces](#)  
Oracle Exadata Database Service on Exascale Infrastructure provides a variety of management interfaces to fit your use case and automation needs.

## About Oracle Exadata Database Service on Exascale Infrastructure

Exadata Database Service on Exascale Infrastructure (ExaDB-XS) is Oracle's newest deployment option for Exadata Database Service. ExaDB-XS provides a cloud service experience similar to Exadata Database Service on Dedicated Infrastructure. Customers can start with a small virtual machine (VM) cluster, and easily scale as needs grow. Oracle manages all of the physical infrastructure in a shared multitenancy infrastructure service model.

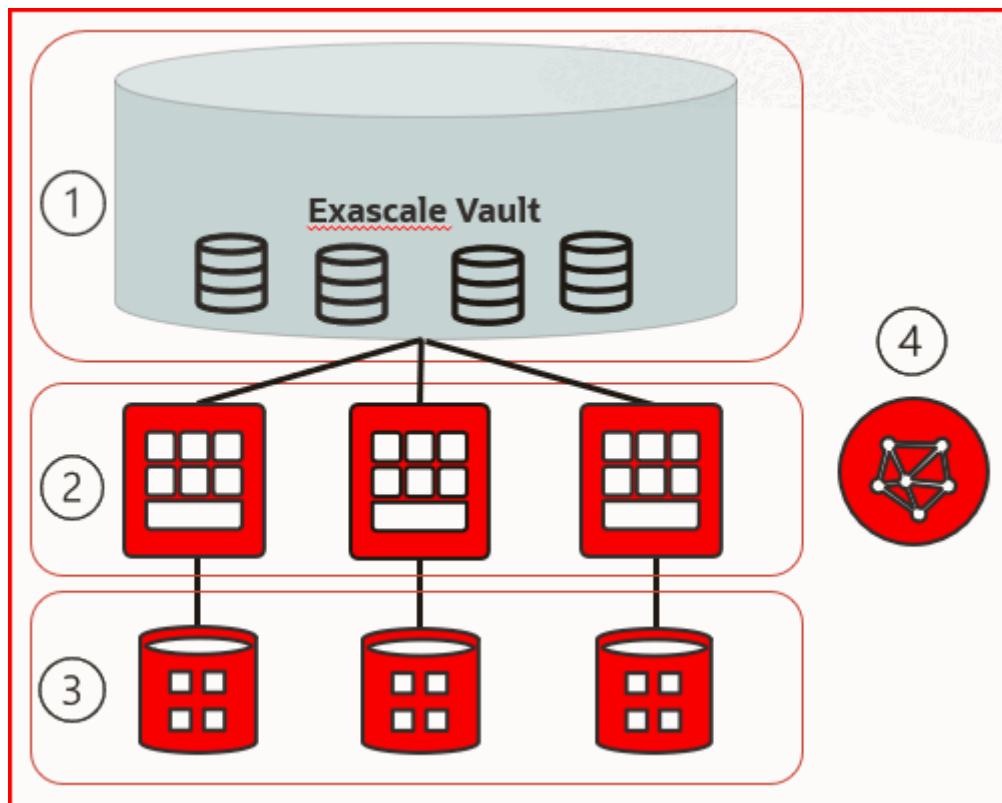
 **Note:**

The What's New chapter contains the current list of regions that are available for Exadata Database Service on Exascale Infrastructure.

Exascale is the underlying technology that serves as the foundation for this service. Exadata Database Service on Exascale Infrastructure is the next generation architecture of Oracle Exadata. It increases storage efficiency, simplifies database provisioning, and combines the extreme performance of Exadata smart software with the cost and elasticity benefits of modern clouds. Storage for database files resides in an Oracle Exadata Exascale Storage Vault. The Storage Vault provides high performance and scalable Exadata smart storage. Storage can be scaled online as needed, with a single command, and that storage becomes available for immediate use. Unlike Dedicated Infrastructure Exadata Database Service on Exascale Infrastructure does not require you to manage adding storage servers to the system, or manage storage allocations.

The following schematic overviews the overall high-level architecture of your VM Cluster and associated resources:

**Figure 1-1 ExaDB-XS Architecture**



The architecture consists of the following elements:

1. A single Exascale Vault, which provides storage for the databases
2. A set of VMs run on Oracle-managed multitenant physical database servers




3. VM filesystems, which are centrally hosted by Oracle
4. A virtual client network (VCN), which provides client and backup network connectivity

The basic unit of consumption in ExaDB-XS is a VM cluster. To facilitate VM portability, Exascale hosts storage for VM file systems on shared storage that is fully managed by Oracle. Oracle can migrate VMs across a pool of physical servers, because the VM filesystems that host the database binaries do not reside on local physical servers. VMs are migrated automatically as required for maintenance, or in the event of a system failure. VMs can also be scaled vertically by changing the number of Elastic Compute Processing Unit (ECPU) units, and changing VM memory allocation. An ECPU is an abstracted measure of compute resources. ECPUs are based on the number of cores elastically allocated from a pool of compute servers. You need at least 8 ECPUs per VM to provision a VM Cluster. VMs can be scaled in increments of 4 ECPUs. For more information about ECPUs, see: [Compute Models in Autonomous Database](#).

In addition to Enabled ECPUs, which are active in the VM, you can also add Additional Reserved ECPUs for your VM. These Additional Reserved ECPUs are physically allocated on a physical server for future scaling of the Enabled ECPUs on your VM, so that you can scale up to meet future workload demands without requiring a restart or relocation of your VM. This option helps to control costs for variable workloads, because database licensing is based on enabled ECPUs. Also, when you reserve additional ECPUs, memory is added to the VM Cluster based on the Total ECPU count. For this reason, reserving additional ECPUs also provides a way for you to provision additional memory without the licensing expense associated with the additional cores.

The following schematic illustrates conceptual details associated with scaling CPU and memory resources:

**Figure 1-2 Core Reservation and Scaling**



The illustration shows the following active and reserved cores in a VM:

1. Eight ECPUs, which are in use and active.
2. Four ECPUs in reserve, which are guaranteed to be available, and standing by, although not in use.

3. A total number of 12 ECPUs in the VM, which is the sum of Enabled cores and Reserved cores.

The total number of cores in a core reservation consists of the sum of Enabled cores and Reserved cores. To scale up your resources without restarting your systems, you can enable the reserved cores. If you want to further scale up your resources, you can add more ECPUs in units of four to your core reservation, and scale up your Enabled and Reserved cores, using a rolling restart as ECPUs are added.

Exascale also provides the benefits of redirect-on-write storage technology. With ExaDB-XS, you can provision thin clones of pluggable databases (PDBs) quickly, with space efficiency, because unchanged blocks are shared between parent and clone PDBs without being duplicated. This feature can be especially useful for development and test environments. You can create numerous thin clones of a PDB economically. For example, you can potentially give each of your developers their own PDB clone on which to work. Because Exadata Exascale has all of the performance advantages of Exadata, development environments provisioned with thin clones are representative of Exadata production environments, and not merely copies of the data.

#### Related Topics

- [Compute Models in Autonomous Database](#)

## Accessing the Exadata Database Service on Exascale Infrastructure Using the OCI Console

Learn how to access the Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS) service.

When the ExaDB-XS service is enabled in your OCI Tenancy you can sign in and select your tenancy region. Then, in the services menu, navigate to **Oracle Database**, and then to **Exadata Database Service on Exascale Infrastructure**. After you navigate to the main page for the service, notice that there are two main objects for this service: VM Clusters and Exascale Storage Vaults.

VM Clusters provide the compute environment where your Oracle Database instances will run. The databases themselves, which are accessed by those Oracle Database instances, are stored in the Exascale Storage Vault. Each VM Cluster has an Exascale Storage Vault assigned to it. You will create and associate the Exascale Storage Vault when creating the VM Cluster as a single, inline experience. However, if any lifecycle operations are then necessary for the Exascale Storage Vault (for example, scaling the total database storage to obtain more free space for expansion), then you complete those lifecycle operations from the Exascale Storage Vaults menu. For most other actions, including provisioning or management of databases, the correct starting point is the VM Clusters page.

## Licensing Considerations for Oracle Exadata Database Service on Exascale Infrastructure

Subscription to Oracle Exadata Database Service on Exascale 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 Exascale Infrastructure.

If you choose to include Oracle Database software licenses in your Oracle Exadata Database Service on Exascale 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). Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure

Oracle Exadata Database Service on Exascale 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).

At the time of release, Oracle Exadata Database Service on Exascale Infrastructure supports Oracle Database 23ai

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

Learn about the available subscription types for Oracle Exadata Database Service on Exascale Infrastructure

The available purchase models are as follows:

### 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.

### Related Topics

- [Universal Credit Pricing FAQ](#)

## Service Limits for Exadata Database Service on Exascale Infrastructure

Limits apply to virtual machine (VM) instance counts, total ECPU count, total local storage, and total High Capacity storage.

The limits set for Exadata Database Service on Exascale Infrastructure (ExaDB-XS) can be revised over time. The following table describes current service limits for ExaDB-XS resources:

**Table 1-1 Service Limits for Exadata Database Service on Exascale Infrastructure**

Limits Name	Description	Limits	Value
exadbxs-vm-instance-base-count	Exadata Database Service on Exascale Infrastructure - Instance Count	Number of VM Instances	4
exadbxs-total-cpu-base-count	Exadata Database Service on Exascale Infrastructure - Total ECPU Count	TotalCpuCores	64
exadbxs-local-storage-base-gb	Exadata Database Service on Exascale Infrastructure - Local Storage (GB)	Local Storage (in GB)	1500
exadbxs-hc-storage-base-gb	Exadata Database Service on Exascale Infrastructure - High Capacity Storage (GB)	High capacity storage (in GB)	2000

## Metering Frequency and Per-Second Billing

See the Per-Second billing, minimums, and limitations on billing.

For each Oracle Exadata Database Service on Exascale Infrastructure virtual machine you provision, you are billed for the infrastructure for a minimum of 48 hours, and then by the second after that. Each ECPU you add to the system is billed by the second, with a minimum usage period of 1 minute.

## Exadata Cloud Management Interfaces

Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure](#)  
The OCI control plane accepts input from the OCI APIs, the OCI Console, and custom interfaces built with kits, tools 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 the following My Oracle Support note: *Exadata Cloud API/CLI Alignment Matrix (Doc ID 2768569.1)*.

### Related Topics

- [Oracle Database console overview](#)
- [Using the Console](#)
- [Exadata Database Service API/CLI Alignment Matrix \(Doc ID 2768569.1\)](#)

## OCI Control Plane Interfaces for Oracle Exadata Database Service on Exascale Infrastructure

The OCI control plane accepts input from the OCI APIs, the OCI Console, and custom interfaces built with kits, tools 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. Oracle Exadata Database Service on Exascale 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 Enable, Disable, or Update Database Management Service*
  - *Using the API to Manage Backup and Recovery*
  - *Using the API to manage Data Guard associations*
  - *Using the API to manage database software images*
  - *Using the API to manage Databases*
  - *Using the API to Manage Oracle Exadata Database Service on Exascale Infrastructure Instance*
  - *Using the API to Manage Oracle Database Home on Oracle Exadata Database Service on Exascale Infrastructure*
  - *Using the API to manage pluggable databases*
  - *Using the API to Patch an Oracle Exadata Database Service on Exascale Infrastructure Instance*
  - *Using the API to upgrade Databases*

### 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](#)

## 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.

**dbaascli:** Use the `dbaascli` utility to perform various database lifecycle and administration operations on the Oracle Exadata Database Service on Exascale Infrastructure such as

- changing the password of a database user
- starting a database
- managing pluggable databases (PDBs)

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.

# 2

## What's New in Oracle Exadata Database Service on Exascale Infrastructure

Oracle is constantly adding new capabilities to Oracle Exadata Database Service on Exascale Infrastructure.

- [New Regions and Realms for ExaDB-XS](#)  
Oracle continues to add new regions for Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).
- [Scale ECPUs to Zero](#)  
You can now scale ECPUs enabled per VM to zero on ExaDB-XS.
- [Deploy Single Node VM Clusters](#)  
You can deploy and run databases in a single-node cluster without requiring an Oracle Real Application Clusters (Oracle RAC) license.

### New Regions and Realms for ExaDB-XS

Oracle continues to add new regions for Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).

 **Note:**

Where the region supports multiple availability domains (AD), only a single availability domain is supported with Oracle Exadata Database Service on Exascale Infrastructure.

**Release Date: January 22, 2025**

ExaDB-XS is now available in the following regions:

- PHX: US West (Phoenix)
- VCP: Brazil Southeast (Vinhedo)

**Release Date: January 14, 2025**

YUL: Canada Southeast (Montreal)

**Release Date: December 19, 2024**

MRS: France South (Marseille)

**Release Date: December 12, 2024**

ExaDB-XS is now available in the following regions:

- AUH: UAE Central (Abu Dhabi)
- KIX: Japan Central (Osaka)



- YNY: South Korea North (Chuncheon)

**Release Date: November 5, 2024**

ExaDB-XS is now available in the following regions:

- LHR: UK South (London)
- ZRH: Switzerland North (Zurich)
- BOG: Colombia Central (Bogota)

**Release Date: November 1, 2024**

ExaDB-XS is now available in the following regions:

- NRT: Japan East (Tokyo)
- ICN: South Korea Central (Seoul)
- YYZ: Canada Southeast (Toronto)

**Release Date: October 24, 2024**

ExaDB-XS is now available in the following regions:

- CDG: France Central (Paris)
- GRU: Brazil East (Sao Paulo)
- JED: Saudi Arabia West (Jeddah)
- MEL: Australia Southeast (Melbourne)
- SIN: Singapore (Singapore)
- SYD: Australia East (Sydney)

**Release Date: October 1, 2024**

ExaDB-XS is now available in the following regions:

- BOM: India West (Mumbai)
- HYD: India South (Hyderabad)

**Release Date: September 2024**

The initial release of ExaDB-XS is available in four regions:

- SJC: US West (San Jose)
- IAD: US East (Ashburn)
- FRA: Germany Central (Frankfurt)
- JNB: South Africa Central (Johannesburg)

**Related Topics**

- [Regions and Availability Domains](#)

## Scale ECPUs to Zero

You can now scale ECPUs enabled per VM to zero on ExaDB-XS.

**Release Date: January 14**

We are pleased to announce general availability of a new feature that allows users to scale the Enabled ECPUs of an ExaDB-XS VM Cluster to zero. This ability enables you to temporarily shut down the VM cluster and avoid billing related to usage of Enabled ECPUs while in this shutdown state.

## Deploy Single Node VM Clusters

You can deploy and run databases in a single-node cluster without requiring an Oracle Real Application Clusters (Oracle RAC) license.

**Release Date: January 22**

We are pleased to announce the general availability (GA) of deploying a VM Cluster running on a single VM for Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS). Single VM support allows customers with smaller and lower availability needs the ability to avoid the cost and complexity of having a multi-node RAC cluster, instead having a single VM only, but still all the Oracle Real Application Clusters (Oracle RAC) software installed and running. VM Clusters can, therefore, be either permanently configured as a single node, or can be scaled back and forth from single node to multiple node VM Cluster configurations depending on workload and availability needs during a period of time.

# 3

## Preparing for Oracle Exadata Database Service on Exascale Infrastructure

Review OCI as well as the site, network and storage requirements to prepare and deploy Oracle Exadata Database Service on Exascale Infrastructure in your data center.

- [Oracle Cloud Infrastructure \(OCI\) Requirements for Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn the basic concepts to get started using Oracle Cloud Infrastructure.
- [Network Setup for Oracle Exadata Database Service on Exascale Infrastructure Instances](#)  
This topic describes the recommended configuration for the VCN and several related requirements for the Oracle Exadata Database Service on Exascale Infrastructure instance.

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

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

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

Before you can provision Oracle Exadata Database Service on Exascale Infrastructure infrastructure, your Oracle Cloud Infrastructure tenancy must be enabled to use Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure](#)  
Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Exascale Infrastructure systems.

#### Related Topics

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

## Required IAM Policy for Oracle Exadata Database Service on Exascale Infrastructure

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

For administrators: The policy in "Let database admins manage DB systems" lets the specified group do everything with databases, and related database resources.

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 Oracle Exadata Database Service on Exascale Infrastructure".

#### Related Topics

- [Let database admins manage DB systems](#)
- [Getting Started with Policies](#)
- [Common Policies](#)
- [Policy Details for the Database Services](#)

- [Policy Details for Oracle Exadata Database Service on Exascale Infrastructure](#)  
This topic covers details for writing policies to control access to Oracle Exadata Database Service on Exascale Infrastructure resources.

## Network Setup for Oracle Exadata Database Service on Exascale Infrastructure Instances

This topic describes the recommended configuration for the VCN and several related requirements for the Oracle Exadata Database Service on Exascale Infrastructure instance.

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

- [VCN and Subnets](#)  
To launch an Oracle Exadata Database Service on Exascale Infrastructure VM cluster, you must have a Virtual Cloud Network and at least two subnets.
- [Node Access to Object Storage: Static Route](#)
- [Service Gateway for the VCN](#)  
Your VCN needs access to both Object Storage for backups and Oracle YUM repos for OS updates.
- [Security Rules for the Oracle Exadata Database Service on Exascale Infrastructure](#)  
This section lists the security rules to use with Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure VM cluster, you must have a Virtual Cloud Network and at least two subnets.

To launch an Oracle Exadata Database Service on Exascale Infrastructure VM cluster, 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 Oracle Exadata Database Service on Exascale Infrastructure VM cluster
- 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*

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](#) for each subnet. You will also create [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](#)  
You must create a VCN with two subnets and ensure that there are enough addresses for the size of your VM cluster.
- [Configuring a Static Route for Accessing the Object Store](#)
- [Setting Up DNS for an Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure instance](#)
- [Configure Private DNS](#)  
Review the prerequisites needed to use Private DNS.

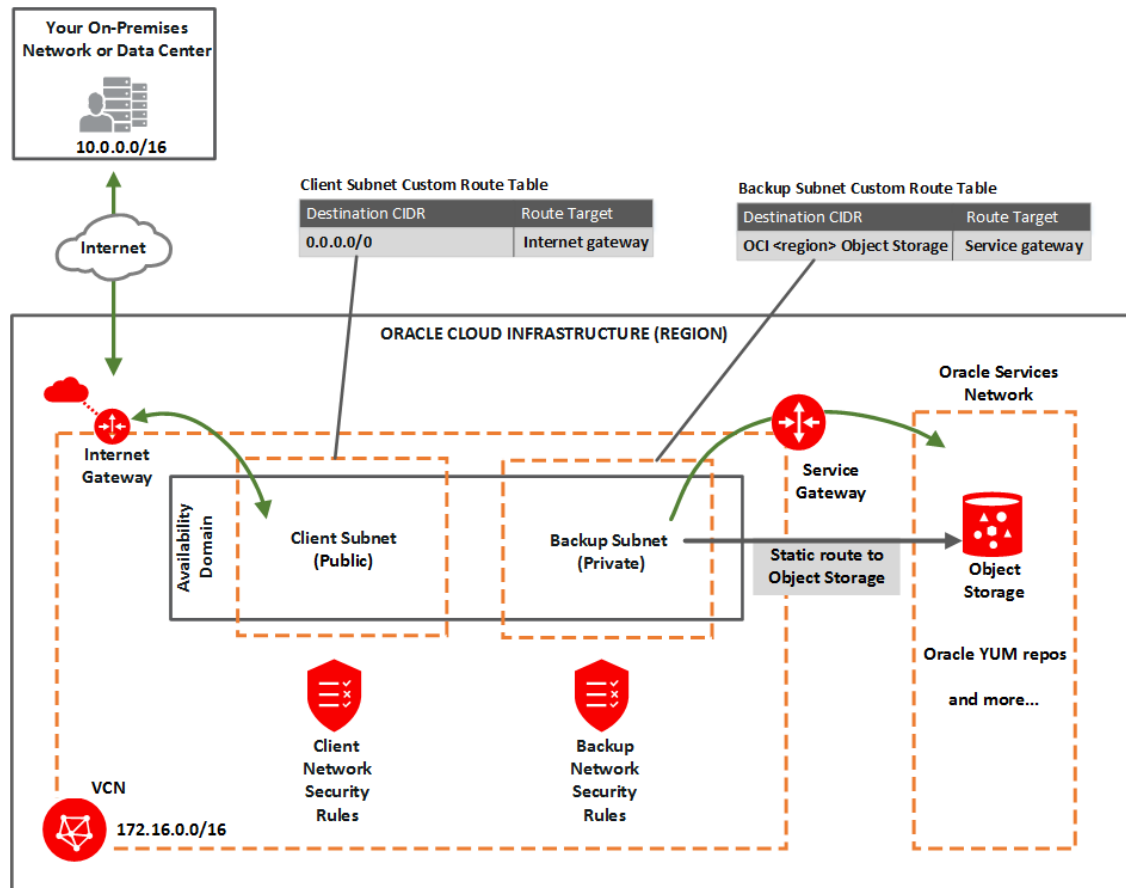
#### **Related Topics**

- [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).
- **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).
- [Security rules](#) to enable the desired traffic to and from the Exadata virtual machines compute nodes.
- [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).

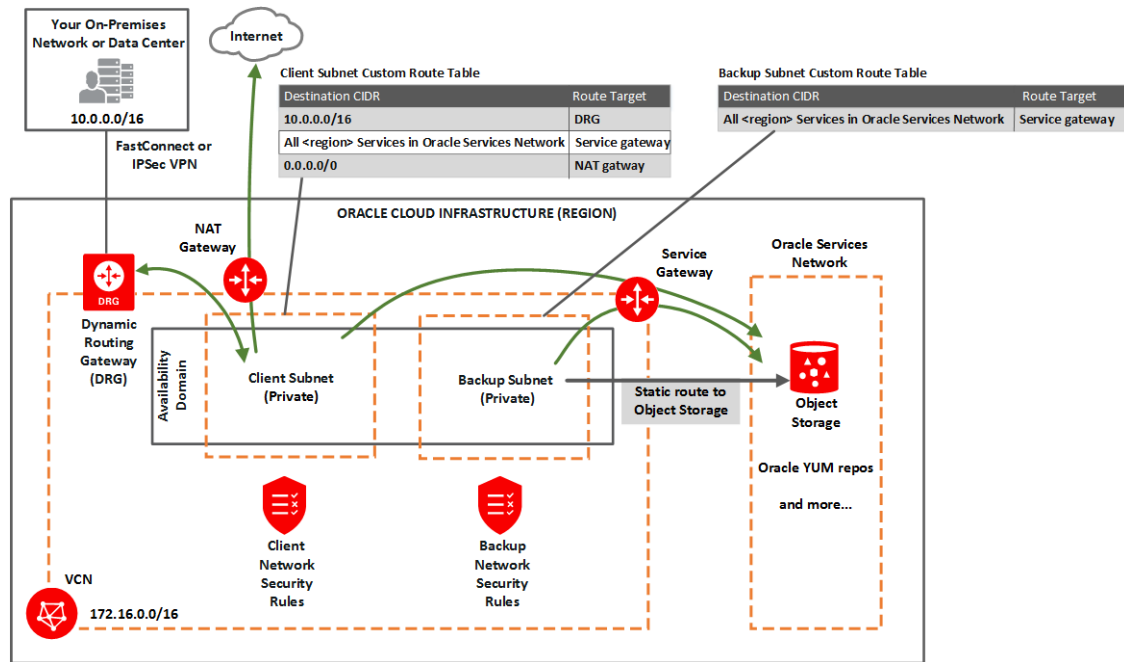
**Important:**

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 Both Object Storage and YUM Repos](#).
  - **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

You must create a VCN with two subnets and ensure that there are enough addresses for the size of your VM cluster.

### Note:

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

If you're setting up VM Clusters (and thus VCNs) in more than one region, then ensure that 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 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. The Networking service reserves three IP addresses in each subnet.

Use the following formula to calculate the minimum number of IP addresses where the variable *n* is the number of VMs in the VM cluster:

The minimum number of client addresses =  $4*n+6$

The minimum number of backup addresses =  $3*n+3$

### Note:

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. To plan for future growth, add addresses that you expect to require as you scale up your VM Cluster, not only the number of VMs you plan to provision for immediate need.

## Configuring a Static Route for Accessing the Object Store

All the traffic in an Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure instance's cloud VM cluster or DB system 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 or DB system. 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.

## 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

Review the prerequisites needed to use Private DNS.

- Private view and private zone must be created before launching DB system 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure instance.

## Service Gateway for the VCN

Your VCN needs access to both Object Storage for backups and Oracle YUM repos for OS updates.

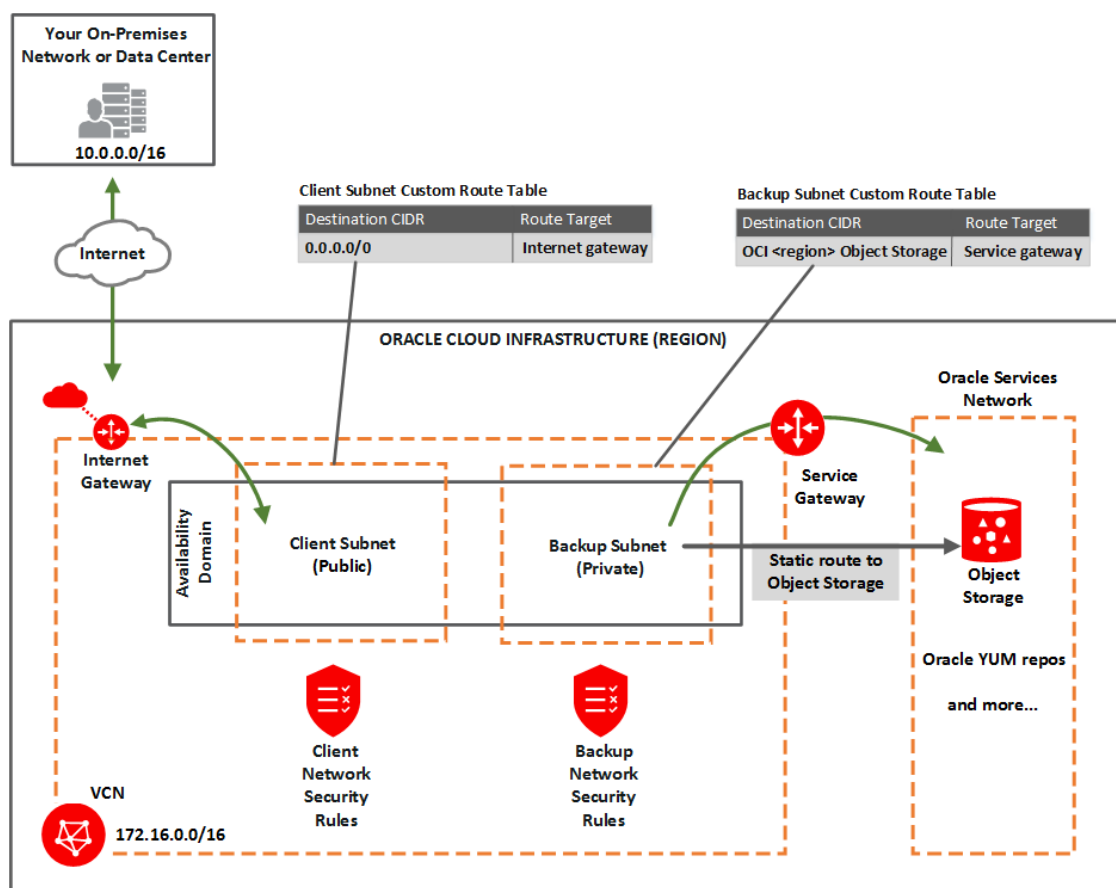
- [Option 1: Service Gateway Access to OCI Services](#)  
You configure the *backup subnet* to use the service gateway for access only to Object Storage.

- [Option 2: Service Gateway Access to Both Object Storage and YUM Repos](#)  
You configure *both the client subnet and backup subnet* to use the service gateway for access to the Oracle Services Network, which includes both Object Storage and the Oracle YUM repos.

## Option 1: Service Gateway Access to OCI Services

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" [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 DB system 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).

## Option 2: Service Gateway Access to Both Object Storage and YUM Repos

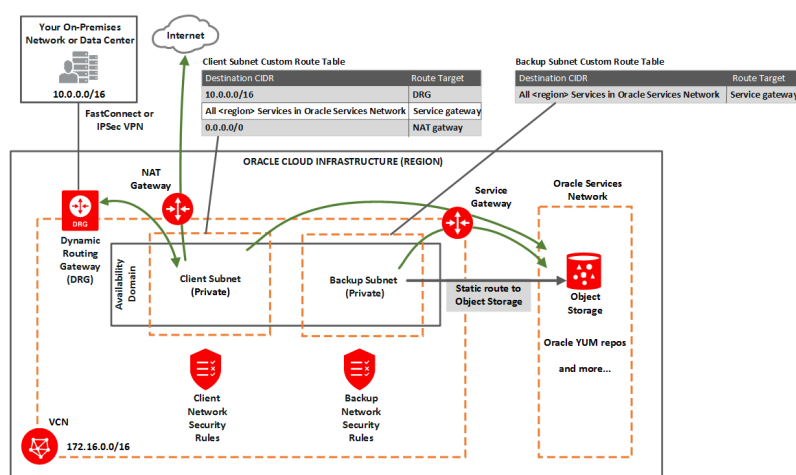
You configure *both the client subnet and backup subnet* to use the service gateway for access to the Oracle Services Network, which includes both Object Storage and the Oracle YUM repos.



### 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*. 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 Exascale Infrastructure

This section lists the security rules to use with Oracle Exadata Database Service on Exascale Infrastructure.

Security rules control the types of traffic allowed for the client network and backup network of the virtual machines. 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.

### Rules Required for Both the Client Network and Backup Network

There are several general rules that enable essential connectivity for hosts in the VCN.

If you use security lists to implement your security rules, then 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
- **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
- **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:** Your VCN's CIDR
- **IP Protocol:** ICMP
- **Type:** 3
- **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
- **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 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 cloud VM cluster or DB system 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.

- **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:** 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:** Client subnet's 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**

- **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**

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

**Required IAM Policies for Oracle Database and Oracle Grid Infrastructure Patching**

Grant Identity and Management (IAM) policies to access subnets, virtual network interface cards (vNICs) and private IP addresses (private-ips) to the users or groups that manages the database and Oracle Grid Infrastructure. For example, suppose the group `admin-group` manages compartment `ABC`. In that case you would set up the following policies:

- Allow group `admin-group` to use subnets in compartment `ABC`
- Allow group `admin-group` to use vNICs in compartment `ABC`
- Allow group `admin-group` to use private-ips in compartment `ABC`

**Rule Required Specifically for the Backup Network**

The following security rule is important for the backup network because it enables the DB system 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 DB system 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).
- [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
- **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
- **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:** Your VCN's 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
- **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 cloud VM cluster or DB system 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 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:** Client subnet's 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:** Client subnet's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521
- **Description:** An optional description of the rule.

### 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
- **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
- **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 DB system 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 1: Allows all egress traffic* in (and in the ). 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), then 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. As the database administrator, when you create an Exadata instance on Exadata Database Service on Exascale Infrastructure, you must choose several networking components (for example, which VCN and subnets to use):
  - When you choose the client subnet, you can also choose which NSG or NSGs to use. Choose the client network's NSG.
  - When you choose the backup subnet, you can also choose which NSG or NSGs to use. Choose the backup network's NSG.

You can instead create a separate NSG for the general rules. Then when database administrators choose which NSGs to use for the client network, 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 *Configuring 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.

### Related Topics

- [Configuring 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.

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](#)

# 4

## Getting Started with Oracle Exadata Database Service on Exascale Infrastructure Deployment

After completing the preparation tasks in [Preparing for Oracle Exadata Database Service on Exascale Infrastructure](#), get started with deploying your Oracle Exadata Database Service on Exascale Infrastructure system following these procedures.

- [Tagging Oracle Exadata Database Service on Exascale 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.
- [Restarting a VM for Planned Maintenance](#)  
To facilitate maintenance of Oracle Exadata Database Service on Exascale Infrastructure virtual machines (VM), Oracle notifies you of planned VM restarts.
- [Connecting to an Oracle Exadata Database Service on Exascale Infrastructure VM](#)  
Learn how to connect to an Oracle Exadata Database Service on Exascale Infrastructure virtual machine (VM) using SSH or SQL Developer.
- [Capacity Limits for Exadata Database Service on Exascale Infrastructure](#)  
To understand the scalability features and resource capacity of the ExaDB-XS service, review these tables and lists.
- [Best Practices for Oracle Exadata Database Service on Exascale Infrastructure VMs](#)  
Oracle recommends that you follow these best practice guidelines to ensure the manageability of your Oracle Exadata Database Service on Exascale Infrastructure virtual machines (VMs).
- [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.

## Tagging Oracle Exadata Database Service on Exascale 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 in accordance with your organizational scheme, which enables you to group resources, manage costs, and give insights into usage. Tags also help you to 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 predefined 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 want 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, and so on. You can achieve this by applying appropriate tags to these resources.

For example, you can 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 to 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 (in this case, 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 Exascale 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:

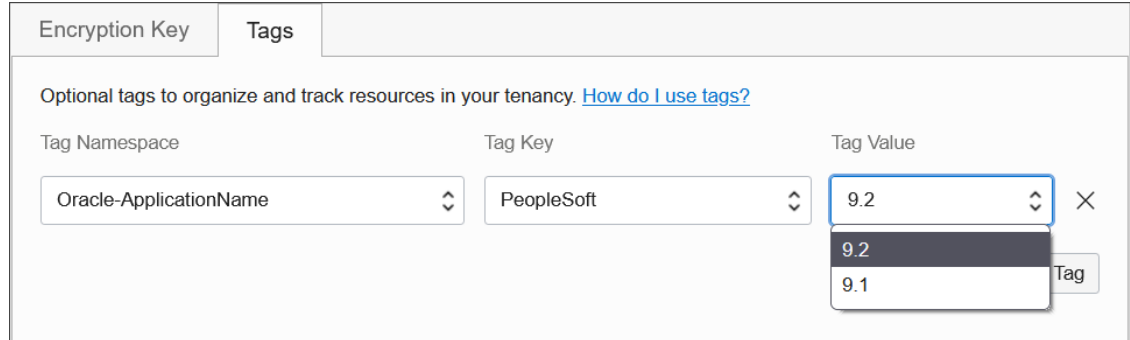
For a "try it out" tutorial that demonstrates implementing tags in Oracle Autonomous Database, refer to *Lab 14: Oracle Standard Tags in Oracle Autonomous Database Dedicated for Fleet Administrators Workshop* on Oracle LiveLabs.

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 (in this example, PeopleSoft) out of the box.

**Figure 4-1 Tagging Example**



**Oracle Standard Tags**

Your tenancy governance administrators can deploy the standard tags at the tenancy level. Your administrators can 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"> <li>• Extreme</li> <li>• High</li> <li>• Medium</li> <li>• Low</li> </ul>	<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"> <li>• Dev</li> <li>• Test</li> <li>• Prod</li> <li>• Pre-Prod</li> <li>• Staging</li> <li>• Trial</li> <li>• Sandbox</li> <li>• User Testing</li> </ul>	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"> <li>• Public</li> <li>• Internal</li> <li>• Sensitive</li> <li>• Highly Sensitive</li> <li>• Extremely Sensitive</li> </ul>	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"> <li>• Public</li> <li>• Customers</li> <li>• Partners</li> <li>• Company</li> <li>• Division</li> <li>• Department</li> <li>• Workgroup</li> </ul>	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"> <li>• 1</li> <li>• 10</li> <li>• 100</li> <li>• 1000</li> <li>• 10000</li> <li>• 100000</li> <li>• 1000000</li> <li>• 10000000</li> </ul>	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@example.com</i> or <i>app_support_grp@example.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"> <li>• HR</li> <li>• Finance</li> <li>• Marketing</li> <li>• Sales</li> <li>• Legal</li> <li>• R&amp;D</li> <li>• Customer Support</li> <li>• Internal Support</li> <li>• Manufacturing</li> </ul>	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"> <li>• 12345</li> <li>• WebMarketing</li> </ul>	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.

**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](#)

## Restarting a VM for Planned Maintenance

To facilitate maintenance of Oracle Exadata Database Service on Exascale Infrastructure virtual machines (VM), Oracle notifies you of planned VM restarts.

The Oracle Exadata Database Service on Exascale Infrastructure VMs use underlying physical hosts that periodically must undergo maintenance. When such maintenance is required, Oracle schedules a restart of your VM, and notifies you of the upcoming restart. The restart enables your VM to be migrated to a new physical host that is not in need of maintenance. Stopping and starting the node will also result in the migration to a new physical host. The only effect to your VM is the restart itself. The planned maintenance of the original physical hardware takes place after your VM has been migrated to its new host, and has no effect on your VM. If you do not restart your VM during the notification period, then Oracle will restart the VM at the end of the notification period.



 **Note:**

When Oracle schedules a restart of your VM, other VMs in that VM Cluster will not be affected by the planned maintenance. The other nodes in your cluster continue to stay available as part of your high availability (HA) strategy.

## Connecting to an Oracle Exadata Database Service on Exascale Infrastructure VM

Learn how to connect to an Oracle Exadata Database Service on Exascale Infrastructure virtual machine (VM) 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:**

Oracle Exadata Database Service on Exascale 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 for Accessing Oracle Exadata Database Service on Exascale Infrastructure](#)  
To use SSH to access a compute node in an Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS) instance, you need this information.
- [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 Oracle Exadata Database Service on Exascale Infrastructure system by using a Secure Shell (SSH) connection.
- [Using Oracle Net Services to Connect to a Database](#)  
Oracle Database Oracle Exadata Database Service on Exascale Infrastructure supports remote database access by using Oracle Net Services.
- [Connect to the Oracle Exadata Database Service on Exascale Infrastructure Service](#)  
Learn how to connect to an Oracle Exadata Database Service on Exascale Infrastructure system using SSH, and how to connect to an Oracle Exadata Database Service on Exascale Infrastructure database using Oracle Net Services (SQL\*Net).

## Prerequisites for Accessing Oracle Exadata Database Service on Exascale Infrastructure

To use SSH to access a compute node in an Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS) instance, you need this information.

 **Note:**

Before you can access ExaDB-XS, you must have configured Exadata Database service on Exascale Infrastructure.

- 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 Oracle Exadata Database Service on Exascale 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 Infrastructure Console. 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 Oracle Exadata Database Service on Exascale Infrastructure instance.

## 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. With the console, you have the option to designate a different SCAN listener port number at VM Cluster provisioning. 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 Oracle Exadata Database Service on Exascale 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-XS system from a Unix-style system using SSH, use this procedure.
- [Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY](#)  
Learn how 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-XS 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 how to access a virtual machine from a Microsoft Windows system using PuTTY.

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.

#### Before you begin

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 that you want to use to connect.
  - To perform operations that require `root`, connect as the user `opc`.
  - To access to the virtual machine for user operations (for example, to run backups), connect as the user `oracle`. (This user can also 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 **Browse** next to the **Private key file for authentication** field. 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, and 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, then 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. Access the VM using SSH as the `opc` user.
2. Log in as the Oracle user. For example: `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 Oracle Exadata Database Service on Exascale Infrastructure supports remote database access by using Oracle Net Services.

Because Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure using Oracle Net Services.

To connect to an Oracle Database on Oracle Exadata Database Service on Exascale 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 or DB system 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 Oracle Exadata Database Service on Exascale 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 DB System'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, `extscaname.example.com` is defined in the on-premises DNS server.

```
testdb =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = <extscaname.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 Exascale 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 Exascale 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 Exascale 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 Exascale 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 Oracle Exadata Database Service on Exascale 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 4-1 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:

```
exanode01.example.com:1521/S12C1
```

Or

```
exanode01.example.com:1521/PDB1.example.oraclecloudatcust.com
```

## Connect to the Oracle Exadata Database Service on Exascale Infrastructure Service

Learn how to connect to an Oracle Exadata Database Service on Exascale Infrastructure system using SSH, and how to connect to an Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 or DB system 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure system using Oracle Net Services.

- [Using Oracle Net Services to Connect to a Database](#)  
Oracle Database Oracle Exadata Database Service on Exascale 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 Exadata Database Service on Exascale Infrastructure 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure supports remote database access by using Oracle Net Services.

Because Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Exadata Database Service on Exascale Infrastructure using Oracle Net Services.

To connect to an Oracle Database on Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Exascale 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 Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Exascale 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 Exascale 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 Oracle Exadata Database Service on Exascale 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:

```
exa1node01.example.com:1521/S12C1
```

Or

```
exa1node01.example.com:1521/PDB1.example.oraclecloudatcust.com
```

## Capacity Limits for Exadata Database Service on Exascale Infrastructure

To understand the scalability features and resource capacity of the ExaDB-XS service, review these tables and lists.

### Minimum VM Cluster Size

- Single-node VM cluster
- 8 total ECPUs per VM
- 280 GB file system storage per VM
- 300 GB Exascale Vault database storage per VM cluster

### VM Scalability Options

VM clusters can be scaled quickly and easily to meet your demands:

- It is possible to enable a subset of the total ECPUs assigned to the VM.
- When you enable more of your reserved ECPUs for the VM, you can scale the ECPUs without having to restart the VM.
- You can scale memory at 2.75 GB per total ECPU.
- You can perform hot additions or removals of VMs
- You can scale VM storage (however, scaling VM storage requires a restart).
- You can perform live, online scaling of Exascale database storage

### Maximum VM Cluster Size

The following list provides maximum VM cluster configuration options

- Number of VMs in the VM Cluster: 10 VMs.
- ECPUs per VM: 200 ECPUs
- File system storage per VM: 2 TB
- Exascale Vault storage per VM cluster: 100 TB



### Database Storage Vault Minimum Capacity

The total minimum capacity billed for ExaDB-XS vaults is 300 GB. Images are stored in an Oracle Advanced Cluster File System (ACFS), and the remainder of space is available for a first database, as described in the following table.

**Table 4-2 ExaDB-XS Minimum Database Storage Vault Capacity for Systems and Database Use**

Purpose	Minimum Capacity
System use (images stored in ACFS)	50 GB
Database use (provisioning a first database)	250 GB

### VM File System Storage Minimum Capacity

The total minimum capacity billed for virtual machine (VM) storage is 280 GB. File system minimum capacities are listed in the following table.

**Table 4-3 ExaDB-XS VM File System Storage Minimum Billed Capacity**

File System	Minimum Total Capacity (GB)	Minimum Usable Capacity (GB)
/boot	0.512	0.412
/ (mirrored)	30	15
/tmp	10	10
/var (mirrored)	10	5
/var/log	18	18
/var/log/audit	3	3
/home	4	4
Swap space (/swap)	16	16
/crashfiles	20	20
/u01	82	80
/u02	84	81
Overhead	2	Not applicable
<b>All file systems (total minimum)</b>	<b>280</b>	<b>Not applicable</b>

## Best Practices for Oracle Exadata Database Service on Exascale Infrastructure VMs

Oracle recommends that you follow these best practice guidelines to ensure the manageability of your Oracle Exadata Database Service on Exascale Infrastructure virtual machines (VMs).

When followed, best practice guidelines can prevent problems that can affect the manageability and performance of your Oracle Exadata Database Service on Exascale Infrastructure VMs:

- 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 Oracle Exadata Database Service on Exascale Infrastructure VM. 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 Guest OS users or manually manipulate SSH key settings associated with your VM.
- 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](#)  
Learn how to use Oracle Data Safe with Oracle Exadata Database Service on Exascale Infrastructure
- [Connecting to an Oracle Exadata Database Service on Exascale Infrastructure VM](#)  
Learn how to connect to an Oracle Exadata Database Service on Exascale Infrastructure virtual machine (VM) using SSH or SQL Developer.
- [Manage Oracle Exadata Database Service on Exascale Infrastructure](#)  
Use the provided tools to manage the Infrastructure.
- [Manage VM Clusters](#)  
Learn how to manage your VM clusters on Oracle Exadata Database Service on Exascale Infrastructure.
- [Manage Exascale Database Vaults on Exadata Database Service on Exascale Infrastructure](#)  
You can view, scale, and delete Exascale Database Storage Vaults on Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).
- [Manage Software Images](#)
- [Create Oracle Database Homes on an Oracle Exadata Database Service on Exascale Infrastructure System](#)  
Learn to create Oracle Database Homes on Oracle Exadata Database Service on Exascale Infrastructure.
- [Managing Oracle Database Homes on an Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure](#)
- [Manage Database Backup and Recovery on Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Exascale Infrastructure.
- [Patch and Update an Oracle Exadata Database Service on Exascale Infrastructure System](#)
- [Manual Software Updates](#)  
For authorized environments, learn how to download manual software updates.
- [Use Oracle Data Guard with Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn to configure and manage Data Guard groups in your VM cluster.
- [Configure Oracle Database Features for Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn how to configure Oracle Multitenant, tablespace encryption, and other options for your Oracle Exadata Database Service on Exascale Infrastructure instance.

- [Migrate to Oracle Exadata Database Service on Exascale Infrastructure](#)  
For general guidance on methods and tools to migrate databases to Oracle Cloud Infrastructure database services, including Oracle Exadata Database Service on Exascale Infrastructure see "Migrating Databases to the Cloud".
- [Connect Identity and Access Management \(IAM\) Users to Oracle Exadata Database Service on Exascale Infrastructure](#)  
You can configure Exadata Database Service on Exascale 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.

## Manage Database Security with Oracle Data Safe

Learn how to use Oracle Data Safe with Oracle Exadata Database Service on Exascale Infrastructure

- [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 Database 23ai 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 23ai 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 Databases 23ai.

## Connecting to an Oracle Exadata Database Service on Exascale Infrastructure VM

Learn how to connect to an Oracle Exadata Database Service on Exascale Infrastructure virtual machine (VM) 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:

Oracle Exadata Database Service on Exascale 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.

- [Connection Prerequisites](#)  
Review the requirements for SSH access to a virtual machine (VM) in Oracle Exadata Database Service on Exascale Infrastructure.
- [About Connecting to a VM with SSH](#)  
You can connect to the virtual machines (VMs) in an Oracle Exadata Database Service on Exascale Infrastructure system by using a Secure Shell (SSH) connection.

- [Connect to the Oracle Exadata Database Service on Exascale Infrastructure Service](#)  
Learn how to connect to an Oracle Exadata Database Service on Exascale Infrastructure system using SSH, and how to connect to an Oracle Exadata Database Service on Exascale Infrastructure database using Oracle Net Services (SQL\*Net).

## Connection Prerequisites

Review the requirements for SSH access to a virtual machine (VM) in Oracle Exadata Database Service on Exascale Infrastructure.

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 Oracle Exadata Database Service on Exascale Infrastructure VM.

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:* On the **Exadata VM Cluster Details** page, click **Virtual Machines** in the **Resources** list.
- *DB systems:* On the **DB System Details** page, click **Nodes** 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 Oracle Exadata Database Service on Exascale Infrastructure VM.

## About Connecting to a VM with SSH

You can connect to the virtual machines (VMs) in an Oracle Exadata Database Service on Exascale 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-XS system from a Unix-style system using SSH, use this procedure.
- [Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY](#)  
Learn how to access a virtual machine from a Microsoft Windows system using PuTTY.
- [To access a database after you connect to the VM](#)  
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-XS 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 how to access a virtual machine from a Microsoft Windows system using PuTTY.

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.

### Before you begin

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 that you want to use to connect.
  - To perform operations that require `root`, connect as the user `opc`.
  - To access to the virtual machine for user operations (for example, to run backups), connect as the user `oracle`. (This user can also use the 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 **Browse** next to the **Private key file for authentication** field. 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, and 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, then 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 VM

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]  
The Oracle base has been set to /u01/app/grid
```

## Connect to the Oracle Exadata Database Service on Exascale Infrastructure Service

Learn how to connect to an Oracle Exadata Database Service on Exascale Infrastructure system using SSH, and how to connect to an Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 or DB system 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure system using Oracle Net Services.

- [Using Oracle Net Services to Connect to a Database](#)  
Oracle Database Oracle Exadata Database Service on Exascale 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 Exadata Database Service on Exascale Infrastructure 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure supports remote database access by using Oracle Net Services.

Because Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure using Oracle Net Services.

To connect to an Oracle Database on Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Exascale 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 Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Exascale 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 Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure

Use the provided tools to manage the Infrastructure.

- [Using the Console to Provision Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn how to provision an Oracle Exadata Database Service on Exascale Infrastructure system.
- [Using the API to Create Infrastructure Components](#)  
See how to use the API for common administrative tasks

- [Using the API to Manage Oracle Exadata Database Service on Exascale Infrastructure Instance](#)  
Use these API operations to manage Oracle Exadata Database Service on Exascale Infrastructure instance components.

## Using the Console to Provision Oracle Exadata Database Service on Exascale Infrastructure

Learn how to provision an Oracle Exadata Database Service on Exascale Infrastructure system.

- [Lifecycle Management Operations](#)
- [Network Management Operations](#)
- [Management Tasks for the Oracle Cloud Infrastructure Platform](#)
- [Oracle Database License Management Tasks](#)  
Learn about licensing for Oracle Exadata Database Service on Exascale Infrastructure

### Lifecycle Management Operations

- [To check the status of a cloud VM cluster](#)
- [To start, stop, or restart an Oracle Exadata Database Service on Exascale Infrastructure cloud VM cluster](#)

#### To check the status of a cloud VM cluster

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under **Oracle Exadata Database Service on Exascale 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 restart an Oracle Exadata Database Service on Exascale Infrastructure cloud VM cluster

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to start, stop, or reboot:

Under **Oracle Exadata Database Service on Exascale 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 the Actions menu (

⋮

) 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 restart a node, the floating IP address might take several minutes to be updated and display in the Console.

## 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 or DB system 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 Oracle Exadata Database Service on Exascale Infrastructure Instances](#).

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to manage:  
*Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Exascale 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.  
*DB systems*: Under Bare Metal, VM, and Exadata, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
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 Oracle Exadata Database Service on Exascale Infrastructure resources](#)
- [To move an Exadata Database Service on Exascale Infrastructure resource to another VM cluster](#)
- [To manage tags for your Oracle Exadata Database Service on Exascale Infrastructure resources](#)

## To view a work request for your Oracle Exadata Database Service on Exascale Infrastructure resources

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Find the Cloud Exadata infrastructure, cloud VM cluster, DB system 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 Database Service on Exascale Infrastructure resource to another VM cluster

#### Note:

- To move resources between compartments, VM cluster users must have sufficient access permissions on the compartment to which the VM cluster is being moved, as well as the current compartment. For more information about permissions for Database resources, see "Details for the Database Service".
- If your Oracle Exadata Database Service on Exascale Infrastructure VM cluster is in a security zone, then the destination compartment must also be in a security zone. See "Security Zone Policies" for a full list of policies that affect Database service resources.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the Cloud Exadata infrastructure, cloud VM cluster that you want to move:  
Under **Exadata Database Service on Exascale 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](#)

To manage tags for your Oracle Exadata Database Service on Exascale Infrastructure resources

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Find the Cloud Exadata infrastructure, cloud VM cluster, DB system 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](#)

## Oracle Database License Management Tasks

Learn about licensing for Oracle Exadata Database Service on Exascale Infrastructure

- [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 ECPUs on the instance. These additional licenses are metered separately.
- [To change the license type of a cloud VM cluster or DB system](#)

### 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 ECPUs on the instance. These additional licenses are metered separately.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to scale:  
Under **Exadata Database Service on Exascale 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** (for cloud VM clusters) or **Scale CPU;Cores** (for DB systems) and then specify a new number of CPU cores. The text below the field indicates the acceptable values, based on the shape used when the DB system was launched.
5. Click **Update**.

### To change the license type of a cloud VM cluster or DB system

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to manage:  
Under **Exadata Database Service on Exascale 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**.

## Using the API to Create Infrastructure Components

See how to use the API for common administrative tasks

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 Oracle Exadata Database Service on Exascale Infrastructure components.

### Exascale Database Storage Vault resource

- [CreateExascaleDbStorageVaul](#)
- [GetExascaleDbStorageVault](#)
- [ListExascaleDbStorageVaults](#)

### Exadata VM cluster resource

- [CreateExadbVmCluster](#)
- [GetExadbVmCluster](#)
- [ListExadbVmClusters](#)

### Databases

- [GetDatabase](#)
- [ListDatabases](#)

### Database Versions

- [ListDbVersions](#)

### Database Homes

- [CreateDbHome](#)
- [GetDbHome](#)
- [ListDbHomes](#)

## Using the API to Manage Oracle Exadata Database Service on Exascale Infrastructure Instance

Use these API operations to manage Oracle Exadata Database Service on Exascale Infrastructure instance 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](#).

### Exascale Database Storage Vault resource

- [ChangeExascaleDbStorageVaultCompartment](#)
- [CreateExascaleDbStorageVault](#)

- [DeleteExascaleDbStorageVault](#)
- [GetExascaleDbStorageVault](#)
- [ListExascaleDbStorageVaults](#)
- [UpdateExascaleDbStorageVault](#)

#### Exadata VM cluster

- [ChangeExadbVmClusterCompartment](#)
- [CreateExadbVmCluster](#)
- [DeleteExadbVmCluster](#)
- [GetExadbVmCluster](#)
- [ListExadbVmClusters](#)
- [RemoveVirtualMachineFromExadbVmCluster](#)
- [UpdateExadbVmCluster](#)

#### Virtual machines nodes (all Oracle Exadata Database Service on Exascale Infrastructure instances)

- [DbNodeAction](#)
- [ListDbNodes](#)
- [GetDbNode](#)

## Manage VM Clusters

Learn how to manage your VM clusters on Oracle Exadata Database Service on Exascale Infrastructure.

- [Using the Console to Manage VM Clusters on Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn how to use the console to create, edit, and manage your VM Clusters on Oracle Exadata Database Service on Exascale Infrastructure.
- [Adding or Removing a VM From a VM Cluster](#)  
You can scale VM Clusters horizontally by adding or removing VMs to or from an existing VM Cluster.
- [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 Oracle Exadata Database Service on Exascale Infrastructure Instance](#)  
Use these API operations to manage Exadata Cloud Infrastructure virtual machines (VMs) and databases on Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).

#### Related Topics

- [Application Checklist for Continuous Service for MAA Solutions](#)

## Using the Console to Manage VM Clusters on Oracle Exadata Database Service on Exascale Infrastructure

Learn how to use the console to create, edit, and manage your VM Clusters on Oracle Exadata Database Service on Exascale Infrastructure.

- [To create a cloud VM cluster](#)  
Create a VM cluster in an Oracle Exadata Database Service on Exascale Infrastructure instance.
- [Using the Console 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 Console 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 scale VM Clusters](#)  
Increase or decrease the ECPUs, memory or storage available to a VM cluster in Oracle Exadata Database Service on Exascale Infrastructure
- [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.
- [Using the Console to Add SSH Keys After Creating a VM Cluster](#)
- [Using the Console to Stop, Start, or Reboot a VM Cluster Virtual Machine](#)  
Use the console to stop, start, or reboot a virtual machine.
- [Using the Console to Check the Status of a VM Cluster Virtual Machine](#)  
Review the health status of a VM cluster virtual machine.
- [Using the Console to Move a VM Cluster to Another Compartment](#)  
To change the compartment that contains your VM cluster on Oracle Exadata Database Service on Exascale Infrastructure, use this procedure.
- [To change the VM cluster display name](#)
- [Using the Console to Terminate a VM Cluster](#)  
Before you can terminate a VM cluster, you must first terminate the databases that it contains.
- [To view details about private DNS configuration](#)

## To create a cloud VM cluster

Create a VM cluster in an Oracle Exadata Database Service on Exascale Infrastructure instance.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Click **Exadata VM Clusters**.
3. Click **Create VM Cluster**.

The **Create 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 DB system. Avoid entering confidential information.
6. **Provide the cluster name:** Select the name of the VM cluster.
7. **Select an availability domain:** Select the availability domain from the displayed options available.
8. **Configure the VM cluster:** Provide the following information:
  - **Number of VMs in the cluster:** Specify the number of the VMs that you want to configure for the cluster, between 1 and 10.
  - **ECPUs enabled per VM:** Specify the number of ECPU cores that you want to enable for the VM cluster. The minimum ECPUs is 8. The maximum number of ECPUs is 200 per VM, or limited by the number of total ECPUs you have specified for the VM. The value you select must be a multiple of 4. You can open the reserve additional ECPU section to reserve additional ECPUs.
9. (Optional) To reserve additional ECPUs, click **Show reserve additional ECPU**. Provide the following information:
  - **ECPUs additional reserved per VM (read only):** Indicates the additional reserved ECPUs. The number of additional ECPUs will be automatically calculated based on the total enabled ECPUs. Additional reserved ECPUs are not active for licensing purposes but are reserved for your VM, and ready and waiting for scaling the Enabled ECPUs. You can review additional read-only fields that show more information about the ECMUs.
  - **Total ECPU per VM:** Provide a total number of ECPUs to allocate per VM. The total must be a number between 8 and 200 in multiples of 4 ECPUs. The read-only field `Total ECPUs across VM Cluster` automatically updates to show you the total number of ECPUs allocated for all VMs in the cluster.
  - **Memory per VM (GB):** This is a read-only field. It displays amount of memory allocated to each VM. Memory is calculated based on 2.75 GB per Total ECPUs. The **Total memory across VM Cluster (GB)** field automatically updates to provide you with the total amount of memory allocated across the VM cluster, based on the memory allocation per VM that you specify.
10. The **VM file system storage** section contains the input field **File system storage capacity per VM (GB):** Specify storage capacity per VM in gigabytes (GB).

Provide how much storage you want for all VM file systems together. The VM file systems storage includes `/u02` capacity, where your Database Homes will go, along with all of the



other VM file systems (/ , /boot, /tmp, /var, /var/log, /var/log/audit, /home, swap, kdump, /u01, grid, /u02). Any extra capacity selected beyond system minimums will go into /u02. The read-only field **Total memory across VM Cluster (GB)** automatically updates to show the total memory allocated across the VM cluster.

 **Note:**

For information about reserved and enabled cores, and an overview of the ExaDB-XS architecture, see "About Exadata Database Service on Exascale Infrastructure"

- 11. Exascale Database Storage Vault:** Select either **Create new vault** or **Select existing vault**. If you select an existing vault, then select the vault in the compartment. Click **Change compartment** to select a vault in a different compartment.

When you create a new vault, the Provisioning status window opens to provide you with the status of vault creation, and the name of the vault that is being created in the format `Vault-YYYYMMDDHHMM` indicating the creation date, where `YYYY` is the year, `MM` is the month, `DD` is the day, `HH` is the hour, and `MM` is the minute.

 **Note:**

If the vault creation failed, then the Provisioning status window provides you with the work request error message indicating the point where the vault creation operation failed, and the work request ID. Make a note of this work request ID, and open a Service Request with My Oracle Support.

- 12. Configure Exascale Database Storage Vault:** Select the storage configuration to use for your database's storage. To begin, select whether you want to create a new Vault, or use an existing Vault.

For a new Vault, specify the following:

- **Storage Vault Name:** Name the new Exascale Vault. *Optional:* Use the link provided to change to another compartment where you want to place the Vault.
- **Enter the Storage Capacity for Databases:** The amount of usable disk storage capacity that will be available for storing databases that is desired. Specify the size in gigabytes (GB) between 300 to 100,000.
- (Optional) **Add smart flash as a percentage of storage capacity provisioned (%):** Select this option to purchase and specify an additional amount of flash cache over and above the amount of default flash cache that is included in the normal Storage capacity for Databases. Additional flash cache can potentially enable increased performance without adding additional storage capacity in some workloads. Additional flash cache also includes additional memory cache. Specify the additional flash cache as a percentage of the total storage provisioned. If you wish to provision additional flash cache, you must add at least 100 GB of additional flash cache. The amount of smart flash cache in GB that will be added is specified in the read-only field **Smart flash cache to be added (GB)**.

The minimum size configuration for an Exascale Database Storage Vault is 300 GB. 50 GB of the space that you allocate in your Vault is reserved for a 200 GB ACFS file system. This ACFS file system resides within your Exascale Database Storage Vault, but is reserved for system use. Thus, if you provisioned the minimum of 300 GB in your Exascale

Database Storage Vault, then 250 GB of that 300 GB capacity will be available storage for your databases.

13. **Select the Oracle Grid Infrastructure version:** This field displays the Oracle Grid Infrastructure versions available for deployment in the VM cluster.
14. **Add SSH key:** Add the public key portion of each key pair that you want to use for SSH access to the DB system:
  - **Generate SSH key pair** (Default option) Select this option 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.

 **Note:**

Download the private key so that you can connect to the database system using SSH. It will not be shown again.

- **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.
15. **Configure the network settings:** Specify the following:
    - **Virtual cloud network:** Select the virtual cloud network (VCN) for the compartment in which you want to create the VM cluster. Click **Change Compartment** to select a VCN in a different compartment.
    - **Client subnet:** Select the client subnet in the compartment. This is the subnet to which the VM cluster should attach. Click **Change Compartment** to select a subnet in a different compartment.

 **Note:**

You must select the VCN before you can select a client subnet.

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:** Select 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.

 **Note:**

You must select the VCN before you can select a backup client subnet.

- **Use network security groups to control traffic:** 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 Oracle Exadata Database Service on Exascale Infrastructure VM cluster.

Note that if you choose a subnet with a **security list**, then 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.
- **Hostname prefix** Provide your choice of hostname for the Exadata DB system. 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 DB system is 12.

 **Caution:**

The hostname must be unique within the subnet. If it is not unique, then the VM cluster will fail to provision.

- **Host domain name:** The domain name for the VM cluster. This is a read-only field. Make a note of the host domain name for your reference.

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.

 **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.

16. **Choose a license type:** The type of license that 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 you want to use your license with Oracle Cloud Infrastructure. This option removes the need for separate on-premises licenses and cloud licenses.
17. Click **Create Exadata VM Cluster**.
  18. (Optional) **Provide a contact for your VM Cluster**. Exadata Database Service on Exascale Infrastructure leverages the OCI Announcements Service. Oracle recommends that you provide your contact details here. Oracle then automatically subscribes you to announcements relevant to this service, including maintenance and outage notifications, among others. If you do not choose to provide a contact now, then you will have to subscribe to announcements manually later, leveraging the OCI Announcements Service directly. To learn more about subscribing, see [Subscribing to Announcements](#).
  19. 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 DB system 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. For more information, see *DB System Time Zone* .

 **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.
- **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.

### Related Topics

- [Network Security Groups](#)
- [Security Lists](#)
- [Oracle Cloud Infrastructure Zero Trust Packet Routing](#)
- [Getting Started with Zero Trust Packet Routing](#)
- [Resource Tags](#)
- [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.

## Using the Console 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 **Database**, click **Exadata Database Service on Exascale 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.

## Using the Console 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 Database**, click **Oracle Exadata Database Service on Exascale 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 Changes**.
  - **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 Oracle Exadata Database Service on Exascale 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 scale VM Clusters

Increase or decrease the ECPUs, memory or storage available to a VM cluster in Oracle Exadata Database Service on Exascale Infrastructure

#### Note:

Oracle doesn't stop billing when a VM or VM Cluster is stopped. To stop billing for a VM Cluster, lower the ECPU count to zero.

You can scale ECPUs enabled per VM. Keep in mind that memory scales with the total ECPU count.

Scaling up or down VM cluster resources requires thorough auditing of existing usage and capacity management by the customer 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 Exascale 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 below this minimum value.

 **Note:**

When scaling a VM Cluster, setting the number of ECPUs to zero will shut down the VM Cluster and eliminate billing related to enabled ECPU usage.

1. Navigate to the **VM Cluster Details** page

2. Click **Scale VM Cluster**.

The **Configure the VM Cluster** window opens, and displays the current configuration of your VM cluster. .

3. Scale your VM cluster as required:

- **ECPUs enabled per VM:** Specify the number of ECPU cores that you want to enable for the VM cluster. The minimum value is zero. If you do not select zero ECPUs, then the minimum enabled ECPUs for each VM is eight. The maximum number of ECPUs is 200 per VM, or limited by the number of total ECPUs you have specified for the VM. The value you select must be a multiple of 4. You can open the reserve additional ECPU section to reserve additional ECPUs.

 **Note:**

Enabled ECPU can be scaled to zero after initial provisioning to temporarily shut down VMs and stop usage billing. Infrastructure billing (for Total ECPU) will continue.

- **ECPUs additional reserved per VM (read only):** Indicates the additional reserved ECPUs. The number of additional ECPUs will be automatically calculated based on the total enabled ECPUs. Additional reserved ECPUs are not active for licensing purposes but are reserved for your VM, and ready and waiting for scaling the Enabled ECPUs.
  - **Total ECPUs per VM:** Provide a total number of ECPUs to allocate per VM. The total must be a number between 8 and 200.
  - **Memory per VM (GB):** This is a read-only field. It displays amount of memory allocated to each VM. Memory is calculated based on 11 GB per total cores. The **Total memory across VM Cluster (GB)** field automatically updates to provide you with the total amount of memory allocated across the VM cluster, based on the memory allocation per VM that you specify.
4. **VM file system storage capacity per VM (GB):** Specify storage capacity per VM in gigabytes (GB).

Provide how much storage you want for all VM filesystems together. The VM Filesystems storage includes /u02 capacity, where your Database Homes will go, along with all of the other VM filesystems (/ , /boot, /tmp, /var, /var/log, /var/log/audit, /home, swap, kdump, /u01, grid, /u02). Any extra capacity selected beyond system minimums will go into /u02.

 **Note:**

For information about reserved and enabled cores, and an overview of the ExaDB-XS architecture, see "About Exadata Database Service on Exascale Infrastructure"

## 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 Database**, then click **Exadata Database Service on Exascale 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 Changes**.

## Using the Console to Add SSH Keys After Creating a VM Cluster

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Exascale 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 Changes**.

#### Related Topics

- [Managing Key Pairs on Linux Instances](#)

## Using the Console 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 Database**, click **Exadata Database Service on Exascale 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.

## Using the Console 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 Database**, click **Exadata Database Service on Exascale 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.

## Using the Console to Move a VM Cluster to Another Compartment

To change the compartment that contains your VM cluster on Oracle Exadata Database Service on Exascale 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 Database**, click **Exadata Database Service on Exascale 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 **Move Resource**.
6. In the resulting dialog, choose the new compartment for the VM cluster, and click **Move Resource**.

## To change the VM cluster display name

### **Note:**

This topic only applies to Oracle Exadata Database Service on Exascale Infrastructure instances using the new Oracle Exadata Database Service on Exascale Infrastructure instance resource model.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under **Oracle Exadata Database Service on Exascale 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 Display Name**.

## Using the Console to Terminate a VM Cluster

Before you can terminate a VM cluster, you must first terminate the databases that it contains.

Terminating a VM cluster removes it from the Cloud Control Plane. In the process, the virtual machines and their contents are destroyed.

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Exascale 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 VM Cluster** to confirm the action.

### Note:

The database stays in a terminated state with backups listed until all backups are expired.

The Exascale Vault that had been associated with the VM Cluster survives the deletion of the VM Cluster. This is because Exascale Vaults can be shared among multiple VM Clusters. If the VM Cluster you've terminated was the only one using the VM Cluster, then you should also terminate the Exascale Vault to stop billing related to the Database Storage. See *Managing Exascale Database Storage Vaults* for more information.

## To view details about private DNS configuration

1. Open the navigation menu. Under **Database**, click **Exadata Database Service on Exascale 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](#)

## Adding or Removing a VM From a VM Cluster

You can scale VM Clusters horizontally by adding or removing VMs to or from an existing VM Cluster.

- [Add a VM to a VM Cluster](#)  
Add a Virtual Machine to a VM Cluster
- [Terminate a VM from a VM Cluster](#)  
To remove a virtual machine from a provisioned cluster, use this procedure.

### Add a VM to a VM Cluster

Add a Virtual Machine to a VM Cluster

#### Note:

- This operation is only available with Multi-VM enabled Infrastructure.
- To add a VM to a VM Cluster requires that all TCP ports are open for the client subnet CIDR for ingress and egress.

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Exascale Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to scale.
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 Virtual Machine**.

#### Note:

Add a VM to a VM Cluster is NOT supported using Terraform.

## Terminate a VM from a VM Cluster

To remove a virtual machine from a provisioned cluster, use this procedure.

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Exascale Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to scale.
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 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.

## 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 Cloud 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 5-1 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"> <li>• <code>/var/log/messages</code></li> <li>• OSWatcher archive</li> <li>• <b>Exadata Only:</b> ExaWatcher archive <code>/opt/oracle.ExaWatcher/archive/</code></li> </ul>

**Table 5-1 (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"> <li>• /etc/oracle</li> <li>• GIHOME/crf/db/ HOSTNAME1</li> <li>• GIHOME/crs/log</li> <li>• GIHOME/css/log</li> <li>• GIHOME/cv/log</li> <li>• GIHOME/evm/ admin/log</li> <li>• GIHOME/evm/admin/ logger</li> <li>• GIHOME/evm/log</li> <li>• GIHOME/log/-/client</li> <li>• GIHOME/log/ HOSTNAME1</li> <li>• GIHOME/log/ HOSTNAME1/admin</li> <li>• GIHOME/log/ HOSTNAME1/client</li> <li>• GIHOME/log/ HOSTNAME1/crflogd</li> <li>• GIHOME/log/ HOSTNAME1/crfmond</li> <li>• GIHOME/log/ HOSTNAME1/crsd</li> <li>• GIHOME/log/ HOSTNAME1/cssd</li> <li>• GIHOME/log/ HOSTNAME1/ctssd</li> <li>• GIHOME/log/ HOSTNAME1/diskmon</li> <li>• GIHOME/log/ HOSTNAME1/evmd</li> <li>• GIHOME/log/ HOSTNAME1/gipcd</li> <li>• GIHOME/log/ HOSTNAME1/gnsd</li> <li>• GIHOME/log/ HOSTNAME1/gpnpd</li> <li>• GIHOME/log/ HOSTNAME1/mdnsd</li> <li>• GIHOME/log/ HOSTNAME1/ohasd</li> <li>• GIHOME/log/ HOSTNAME1/racg</li> <li>• GIHOME/log/ HOSTNAME1/srvm</li> <li>• GIHOME/log/ HOSTNAME1/xag</li> </ul>

Table 5-1 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
		<ul style="list-style-type: none"> <li>• GIHOME/log/diag/asmtool</li> <li>• GIHOME/log/diag/clients</li> <li>• GIHOME/log/procwatcher/PRW_SYS_HOSTNAME1</li> <li>• GIHOME/network/log</li> <li>• GIHOME/opmn/logs</li> <li>• GIHOME/racg/log</li> <li>• GIHOME/scheduler/log</li> <li>• GIHOME/srvm/log</li> <li>• GRIDBASE/crsdata/@global/cvu</li> <li>• GRIDBASE/crsdata/HOSTNAME1/core</li> <li>• GRIDBASE/crsdata/HOSTNAME1/crsconfig</li> <li>• GRIDBASE/crsdata/HOSTNAME1/crsdiag</li> <li>• GRIDBASE/crsdata/HOSTNAME1/cvu</li> <li>• GRIDBASE/crsdata/HOSTNAME1/evm</li> <li>• GRIDBASE/crsdata/HOSTNAME1/output</li> <li>• GRIDBASE/crsdata/HOSTNAME1/ovmmwallets</li> <li>• GRIDBASE/crsdata/HOSTNAME1/scripts</li> <li>• GRIDBASE/crsdata/HOSTNAME1/trace</li> <li>• GRIDBASE/diag/crs/-/crs/cdump</li> <li>• GRIDBASE/diag/crs/HOSTNAME1/crs/cdump</li> <li>• GRIDBASE/diag/crs/HOSTNAME1/crs/incident</li> <li>• GRIDBASE/diag/crs/HOSTNAME1/crs/trace</li> </ul>



**Table 5-1 (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 ipspace based on the time range for certain DB incidents.	<ul style="list-style-type: none"> <li>• ORACLE_BASE/diag/rdbms/&lt;dbname&gt;/&lt;instance_name&gt;/cdump</li> <li>• ORACLE_BASE/diag/rdbms/&lt;dbname&gt;/&lt;instance_name&gt;/trace</li> <li>• ORACLE_BASE/diag/rdbms/&lt;dbname&gt;/&lt;instance_name&gt;/incident</li> </ul>

### 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/dbssystem\_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

## Health Metrics

Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.

 **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.

### Guest VM Health Metrics List - Database Metrics

**Table 5-2 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 ECPUs.
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.

**Table 5-2 (Cont.) Guest VM Health Metrics List - Database Metrics**

<b>Metric Name</b>	<b>Metric Display Name</b>	<b>Unit</b>	<b>Aggregation</b>	<b>Interval</b>	<b>Collection Frequency</b>	<b>Description</b>
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.
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.

**Table 5-2 (Cont.) Guest VM Health Metrics List - Database Metrics**

<b>Metric Name</b>	<b>Metric Display Name</b>	<b>Unit</b>	<b>Aggregation</b>	<b>Interval</b>	<b>Collection Frequency</b>	<b>Description</b>
StorageAllocated	Storage Space Allocated	GB	Max	One hour	One hour	Total amount of storage space allocated to the database at the collection time.
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 5-3 Guest VM Health Metrics List - Non-Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
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.

## Using the API to Manage Oracle Exadata Database Service on Exascale Infrastructure Instance

Use these API operations to manage Exadata Cloud Infrastructure virtual machines (VMs) and databases on Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).

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 Oracle Exadata Database Service on Exascale Infrastructure instance components.

### Exascale Database Storage Vault resource

- [ChangeExascaleDbStorageVaultCompartment](#)

- [CreateExascaleDbStorageVault](#)
- [DeleteExascaleDbStorageVault](#)
- [GetExascaleDbStorageVault](#)
- [ListExascaleDbStorageVaults](#)
- [UpdateExascaleDbStorageVault](#)

#### Exadata VM cluster

- [ChangeExadbVmClusterCompartment](#)
- [CreateExadbVmCluster](#)
- [DeleteExadbVmCluster](#)
- [GetExadbVmCluster](#)
- [ListExadbVmClusters](#)
- [RemoveVirtualMachineFromExadbVmCluste](#)
- [UpdateExadbVmCluster](#)

## Manage Exascale Database Vaults on Exadata Database Service on Exascale Infrastructure

You can view, scale, and delete Exascale Database Storage Vaults on Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).

### Viewing Exascale Database Storage Vaults

1. Navigate to Exascale Database Storage Vaults
2. Select **Exascale Database Storage Vaults**.
3. Select the Vault for which you want to view information.

### Deleting Exascale Database Storage Vaults

1. Navigate to Exadata Database Service on Exascale Infrastructure.
2. Select **Exascale Database Storage Vaults**
3. Select the Vault that you want to delete.
4. Select **Delete**.
5. Confirm that you want to delete the Vault.

#### Note:

You can only delete Exascale Database Storage Vaults that no longer have any associated VM Clusters. If you still have associated VM Clusters on the Vault, then you must first delete those VM Clusters, and then return to these steps to delete your Exascale Database Storage Vault.

### Scaling Exascale Database Storage Vaults

1. Navigate to Exadata Database Service on Exascale Infrastructure.

2. Select **Exascale Database Storage Vaults**.
3. Select the Vault that you want to scale.
4. Select **Scale Storage Vault**.
5. On the Scale Storage Vault dialog, enter a number for the capacity for High Capacity storage. This number should be the value for the total storage that you want to have provisioned after the scaling operation completes. In addition to the default flash cache included in the base storage capacity, you can choose to configure additional smart flash cache as a percentage of provisioned storage capacity.
6. Click **Save Changes**. Your Vault will be scaled automatically.

## 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 DB system, 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 DB system, 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 create a database software image](#)  
Follow this procedure to create a database on Oracle Exadata Database Service on Exascale Infrastructure
- [To create a database software image from a Database Home](#)
- [To update database software using custom database software image](#)  
Use the following instructions to update database software using a custom database software image.
- [To delete a software image](#)  
Use the following instructions to delete a software image.

### To create a database software image

Follow this procedure to create a database on Oracle Exadata Database Service on Exascale Infrastructure

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Under **Resources**, click **Database Software Images**.
3. Click **Create Database Software Image**.
4. In the **Display name** field, provide a display name for your image. Avoid entering confidential information.

5. Choose your **Compartment**.
6. Choose the **Database version** for your image. You can create a database software image using any supported Oracle Database release update (RU).

 **Note:**

At the time of initial release of Exadata Database Service on Exascale Infrastructure, only Oracle Database 23ai is supported.

7. Choose the **patch set update, proactive bundle patch, or release update**. For information on Oracle Database patching models, see [Release Update Introduction and FAQ \(Doc ID 2285040.1\)](#)
8. Optionally, you can enter a comma-separated list of one-off (interim) patch numbers.
9. Optionally, you can upload an Oracle Home inventory file from an existing Oracle Database. See [Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Home](#) for instructions on creating an inventory file using OPatch.
10. 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.
11. Click **Create Database Software Image**.

## To create a database software image from a Database Home

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the Database Home: Under **Oracle Exadata Database Service on Exascale 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 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 Database**, then click **Exadata Database Service on Exascale 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. On the resulting VM cluster details page, click the **View updates** link in the **Version** section.
5. On the resulting Updates page, click the **Custom Database Software Images** tab under the **Database Home** section.
6. Choose a **Compartment**.
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 **Run Precheck**.
9. On the resulting Confirm dialog, click **OK** to continue.
10. After running the precheck successfully, select **Apply** from the Actions button (three dots).
11. On the resulting Confirm dialog, click **OK** to continue.

## To delete a software image

Use the following instructions to delete a software image.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale 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 software image**.

## 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 Oracle Exadata Database Service on Exascale Infrastructure System

Learn to create Oracle Database Homes on Oracle Exadata Database Service on Exascale Infrastructure.

- [About Creating Oracle Database Homes on an Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure](#)  
To create an Oracle Database home, review the list of API calls.

## About Creating Oracle Database Homes on an Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn to use the `dbaascli` utility on Oracle Exadata Database Service on Exascale Infrastructure.

## To create a new Database Home in an existing Oracle Exadata Database Service on Exascale 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.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster on which you want to create the new Database Home.

Under **Oracle Exadata Database Service on Exascale 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 Container Database**.
6. In the **Create Container Database** 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 an older 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** and **Select a Database version** selectors to limit the list of custom database software images to a specific compartment or Oracle Database software major release version.

### 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.

- 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 Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the Database Home: Under **Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 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 Oracle Exadata Database Service on Exascale 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 view the details of a Database home, use this procedure.
- [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 add and modify metadata tags to help to manage your Database, use this procedure.
- [Using the Console to Move a Database to Another Database Home](#)  
Learn to move a database to another Database Home.

### To view information about a Database Home

To view the details of a Database home, use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the Database Home.

Under **Oracle Exadata Database Service on Exascale 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.

(Optional) Enter the result of the procedure here.

## 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 Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the Database Home you want to delete:

Under **Oracle Exadata Database Service on Exascale 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.

## To manage tags for your Database Home

To add and modify metadata tags to help to manage your Database, use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the Database Home:

Under **Oracle Exadata Database Service on Exascale 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 (

⋮

) on the row listing the Database Home, and then click **Add Tags**.



## Using the Console to Move a Database to Another Database Home

Learn to move a database to another Database Home.

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure

- [Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Exascale Infrastructure](#)  
Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Exascale Infrastructure.
- [Oracle Database Releases Supported by Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn about the versions of Oracle Database that Oracle Exadata Database Service on Exascale Infrastructure supports.
- [Provisioning and Managing Exadata Databases](#)  
This topic describes creating and managing Oracle Databases on an Oracle Exadata Database Service on Exascale Infrastructure instance.
- [Using the API to manage Databases](#)
- [Create and Manage Exadata Pluggable Databases](#)  
You can create and manage pluggable databases (PDBs) in Oracle Exadata Database Service on Exascale 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.
- [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 Exascale Infrastructure

Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Exascale Infrastructure.

Before you can create and use an Oracle Database on Oracle Exadata Database Service on Exascale Infrastructure, you must:

- Configure a VM cluster
- Create any required backup destinations

You can create one or more databases on each Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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.

## Oracle Database Releases Supported by Oracle Exadata Database Service on Exascale Infrastructure

Learn about the versions of Oracle Database that Oracle Exadata Database Service on Exascale Infrastructure supports.

At the time of this release, Oracle Exadata Database Service on Exascale Infrastructure supports Oracle Database 23ai only.

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

- <https://support.oracle.com/epmos/faces/DocContentDisplay?id=742060.1>

## Provisioning and Managing Exadata Databases

This topic describes creating and managing Oracle Databases on an Oracle Exadata Database Service on Exascale 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).

You can create Database homes, databases, and pluggable databases at any time by using the Console.

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](#).

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 Oracle Exadata Database Service on Exascale Infrastructure](#)  
Customer-managed keys for Oracle Exadata Database Service on Exascale 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 Exascale Infrastructure](#)  
To create or terminate a database, complete procedures using the Oracle Exadata console.

## 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 Oracle Exadata Database Service on Exascale Infrastructure

Customer-managed keys for Oracle Exadata Database Service on Exascale 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 do the following:

- Enable customer-managed keys when you create databases in Oracle Exadata Database Service on Exascale 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. For vaults that are replicated, this policy is needed: `allow dynamic-group dynamic_group_name to read vaults in tenancy`

### Limitations

To enable Oracle Data Guard on Oracle Exadata Database Service on Exascale Infrastructure databases that use customer-managed keys, the primary and standby databases must be in the same realm.

- [To integrate customer-managed key management into Oracle Exadata Database Service on Exascale Infrastructure](#)  
If you choose to encrypt databases in an Oracle Exadata Database Service on Exascale 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

- [Replicating Vaults and Keys](#)
- [Learn About Oracle Cloud Basics](#)

## To integrate customer-managed key management into Oracle Exadata Database Service on Exascale Infrastructure

If you choose to encrypt databases in an Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure instance:

1. Confirm that the LIBKMS package is already installed, as follows:

```
rpm -qa --last | grep libkmstdepkcs11
```

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 Exascale Infrastructure

To create or terminate a database, complete procedures using the Oracle Exadata console.

- [To create a database in an existing Oracle Exadata Database Service on Exascale Infrastructure VM Cluster](#)  
Learn how you can create your first or subsequent databases.
- [Using the Console 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](#)  
Learn how to use a backup to create a database on Exadata Database Service on Exascale Infrastructure.
- [To create a database from the latest backup](#)  
Use this procedure to create a database from the latest backup on Oracle Exadata Database Service on Exascale Infrastructure.
- [To move a database to another Database Home](#)  
To patch a single Oracle Database in your Oracle Exadata Database Service on Exascale Infrastructure instance, you move it to another Database Home.
- [To terminate a database](#)  
Use this procedure to terminate a database on Oracle Exadata Database Service on Exascale Infrastructure.
- [To administer Vault encryption keys](#)  
Use this procedure to rotate the Vault encryption key or or change the encryption management configuration.

## To create a database in an existing Oracle Exadata Database Service on Exascale Infrastructure VM Cluster

Learn how you can create your first or subsequent databases.

### Note:

If IORM is enabled on the Oracle Exadata Database Service on Exascale Infrastructure VM Cluster, 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 Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system you want to create the database in: **Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model)**: Under **Oracle Exadata Database Service on Exascale 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 DB system.
- **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 an older 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:**  
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.

- \* **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** and **Select a Database version**



selectors to limit the list of custom database software images to a specific compartment or Oracle Database software major release version.

- **PDB name:** *(Optional)* 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.
- **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 DB system. 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:
  - **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:**

If you are creating a database in an Exadata Cloud Service VM Cluster, then you can choose to use encryption based on encryption keys that you manage. By default, the database is configured using Oracle-managed encryption keys. To configure the database with encryption based on encryption keys you manage:

- a. Select **Use customer-managed keys**. You must have a valid encryption key in Oracle Cloud Infrastructure Vault service. See [Let security admins manage vaults, keys, and secrets](#).

 **Note:**

You must use AES-256 encryption keys for your database.

- b. Choose a **Vault**.
- c. Select a **Master encryption 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.

- **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**.

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

- [security zone compartment](#)
- [Resource Tags](#)
- [Let security admins manage vaults, keys, and secrets](#)

## Using the Console 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 Database**, then click **Oracle Exadata Database Service on Exascale 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 Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database:  
Under **Exadata at Oracle 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. 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

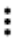
Learn how to use a backup to create a database on Exadata Database Service on Exascale Infrastructure.

Before you begin, note the following:

- When you create a database from a backup, the availability domain is the same as the availability domain that hosts the backup or a different one within the same region.
- 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 [Using the Console](#)

- 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 Database**, then click **Exadata Database Service on Exascale Infrastructure**.
  2. Choose your **Compartment**.
  3. Navigate to a backup.
    - *Standalone backups*: Click **Standalone Backups** under **Oracle Exadata Database Service on Exascale Infrastructure**.
    - *Automatic backups*: Navigate to the Database Details page of the database associated with the backup:
      - Under **Oracle Exadata Database Service on Exascale 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 menu icon (  ) 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 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 DB system** section:
  - 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, then 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, and then choose an image from the list of available images for your selected software version.
  - c. 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
    - Must 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`.

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
- Must not use the following reserved names: grid, ASM

**12. Click Create Database.**

**NOT\_SUPPORTED**

1. Click the Exadata cloud VM cluster or DB system 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 **Exadata Database Service on Exascale 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

Use this procedure to create a database from the latest backup on Oracle Exadata Database Service on Exascale Infrastructure.

Before you begin, note the following:

- When you create a database from a backup, the availability domain is the same as the availability domain that hosts the backup or a different one within the same region.
- 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 Database**, then click **Exadata Database Service on Exascale 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:

Under **Oracle Exadata Database Service on Exascale 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 **Create Database from Last Backup**.



6. In the **Provide basic information for the Exadata infrastructure** section:
  - **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. On the **Create Database from Backup** page, configure the database as follows.
8. In the **Configure your DB system** section: Choose a cloud VM cluster to run the database from the **Select a VM cluster** drop-down list.
9. 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.
  - **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**.
10. In the **Configure database** section:

 **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.
11. 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.
  12. 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`. 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
  13. Click **Create Database**.

## To move a database to another Database Home

To patch a single Oracle Database in your Oracle Exadata Database Service on Exascale Infrastructure instance, you move 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 [Using the Console to Manage Databases on Oracle Exadata Database Service on Exascale Infrastructure](#).

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database you want to move.

Under **Oracle Exadata Database Service on Exascale 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 Database**.
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 run 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

Use this procedure to terminate a database on Oracle Exadata Database Service on Exascale Infrastructure.

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 Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database:

Under **Oracle Exadata Database Service on Exascale 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 the database using Oracle Cloud Infrastructure Object Storage or Oracle Database Autonomous Recovery Service: In the confirmation dialog,
  - Review the message about the backup retention policy.
  - Configure automatic backups as needed.

- Type the name of the database to confirm the termination
5. Click **Terminate Database**.  
The database's status indicates Terminating.

 **Note:**

The database stays in a terminated state with backups listed until all backups are expired.

## To administer Vault encryption keys

Use this procedure to rotate the Vault encryption key or or change the encryption management configuration.

After you provision a database in an Exadata DB system or 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 Database**, then click **Exadata Database Service on Exascale 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 Exascale 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.

## 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 Oracle Exadata Database Service on Exascale 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.

Databases created in Oracle Exadata Database Service on Exascale 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**  
When the CDB is configured with the auto-backup feature, you have the option to take a backup of the PDB during create, clone, or relocate operations. The PDB backup destination will always be the same as the CDB, and the backups cannot be accessed directly or created on demand. Oracle recommends that you immediately back up the PDB after you create or clone them. This is because the PDB will not be recoverable until the next daily auto-backup completes successfully, leading to a possible data loss.
- **Restore**
  - **Oracle Exadata Database Service on Exascale 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, and then selecting a PDB or a subset of them that you want to restore on the new database.

- **Relocate**

You can relocate a PDB from one CDB to another CDB within the same availability domain (AD):

- Across compartments, VM clusters, DB system, or VCNs. 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 an Oracle Data Guard association, a PDB relocated to the primary will be synchronized with the standby as well.

- **Clone**

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, or VCNs. 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, or VCNs. 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 an Oracle 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, then the system displays the same status for all PDBs. If the PDB statuses are different across the nodes, then 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.

## 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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.

- [Using the console to create pluggable database](#)  
To create the PDB, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure



- [Using the console to relocate a pluggable database](#)  
To relocate the PDB, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure
- [Using the API to create pluggable database](#)

## Using the console to create pluggable database

To create the PDB, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure



### Note:

Creating a pluggable database (PDB) is not supported for databases using Data Guard.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the database:  
  
Under **Oracle Exadata Database Service on Exascale 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. Note: For bare metal DB systems, you cannot have two PDBs in the same database that use the same PDB name. You can use the same name for PDBs in different databases within the same DB system.
  - **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, then create and enter a PDB admin password. The password must contain the following:
    - 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 Pluggable Database**.

#### WHAT NEXT?

After creating your PDB, you can get [connection strings](#) for the administrative service using the OCI Console.

## Using the console to relocate a pluggable database

To relocate the PDB, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the database:

Under **Oracle Exadata Database Service on Exascale 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 the following:
      - \* 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 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_pdbs 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, then 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 pluggable database**.

 **Note:**

Relocate will incur downtime during the process. 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 Oracle Exadata Database Service on Exascale 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 start the PDB, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure
- [To stop a pluggable database](#)  
To stop the PDB, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure.
- [To delete a pluggable database](#)
- [To get connection strings for a pluggable database](#)

## To start a pluggable database

To start the PDB, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure

### Note:

The PDB must be available and stopped to use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database.
  - 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 PDB** to confirm the start operation.

## To stop a pluggable database

To stop the PDB, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure.

### Note:

The PDB must be available and running (started) to use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database.

- 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 PDB** to confirm the stop operation.

## To delete a pluggable database

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.

2. Choose your Compartment.

3. Navigate to the database:

*Cloud VM clusters (new resource model)* Under **Oracle Exadata Database Service on Exascale 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.

*DB systems* Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

On the cloud VM cluster or DB system 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 PDB**.

## 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 Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

*Cloud VM clusters (new resource model)* Under **Oracle Exadata Database Service on Exascale 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.

*DB systems* Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

On the cloud VM cluster or DB system 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.

## 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 23ai, the cloned PDB is automatically upgraded to 23ai. For example, if you use refreshable clones to clone to 23ai and then convert it to regular PDB, all necessary upgrade steps are automatically handled, converting the refreshable clone into a fully upgraded 23ai 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.
- [Using the Console to Create a Local Clone of a Pluggable Database \(PDB\)](#)  
Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.
- [Using the Console to Create a Remote Clone of a Pluggable Database \(PDB\)](#)  
Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.
- [Using the Console to Create a Refreshable Clone of a Pluggable Database \(PDB\)](#)  
Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.
- [Using the Console to Refresh a Cloned Pluggable Database \(PDB\)](#)  
Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.

- [Using the Console to Convert a Refreshable Clone to a Regular Pluggable Database \(PDB\)](#)  
Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.
- [Using the API to clone a pluggable database](#)  
Learn how to manage pluggable databases (PDBs) using the pluggable database API endpoints.

## Using the Console to Create a Local Clone of a Pluggable Database (PDB)

Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:  
  
Under **Oracle Exadata Database Service on Exascale 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 the following:
  - 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.

- *Optional.* **Enable thin clone:** Select this option to leverage Exascale redirect-on-write technology to create a thin clone of the PDB. This option results in the reuse of duplicate blocks with the parent PDB, shared with the clone. Deselecting this option results in a traditional, full clone with all blocks copied, and fully independent from the parent.
  - **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 pluggable database**.

## Using the Console to Create a Remote Clone of a Pluggable Database (PDB)

Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:  
Under **Oracle Exadata Database Service on Exascale 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 the following:
  - 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.
- **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.

- *Optional.* **Enable thin clone:** Select this option to leverage Exascale redirect-on-write technology to create a thin clone of the PDB. This option results in the reuse of duplicate blocks with the parent PDB, shared with the clone. Deselecting this option results in a traditional, full clone with all blocks copied, and fully independent from the parent.

- **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 pluggable database**.

## Using the Console to Create a Refreshable Clone of a Pluggable Database (PDB)

Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Under **Oracle Exadata Database Service on Exascale 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.

*DB systems* Under **Bare Metal, VM, and Exadata**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

On the cloud VM cluster or DB system 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.
- **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.

- *Optional.* **Enable thin clone:** Select this option to leverage Exascale redirect-on-write technology to create a thin clone of the PDB. This option results in the reuse of duplicate blocks with the parent PDB, shared with the clone. Deselecting this option results in a traditional, full clone with all blocks copied, and fully independent from the parent.
  - **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 pluggable database**.

## Using the Console to Refresh a Cloned Pluggable Database (PDB)

Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Under **Exadata Database Service on Exascale 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.

## Using the Console to Convert a Refreshable Clone to a Regular Pluggable Database (PDB)

Complete this procedure on Oracle Exadata Database Service on Exascale Infrastructure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Under **Oracle Exadata Database Service on Exascale 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 or DB system 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, then 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

Learn how to manage pluggable databases (PDBs) using the pluggable database API endpoints.

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](#).

## 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 the 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, and then selecting a PDB or a subset of them you want to restore on the new database.
- [Using the Console to Perform an In-Place Restore of a Pluggable Database \(PDB\)](#)  
Complete this procedure for an in-place PDB restore using an RMAN backup on Exadata Database Service on Exascale Infrastructure
- [Using the Console to Perform an Out-of-Place Restore of a Pluggable Database \(PDB\)](#)  
Complete this procedure for an out-of-place PDB restore on Exadata Database Service on Exascale Infrastructure

## Using the Console to Perform an In-Place Restore of a Pluggable Database (PDB)

Complete this procedure for an in-place PDB restore using an RMAN backup on Exadata Database Service on Exascale Infrastructure

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Under **Oracle Exadata Database Service on Exascale 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 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) that 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**.

## Using the Console to Perform an Out-of-Place Restore of a Pluggable Database (PDB)

Complete this procedure for an out-of-place PDB restore on Exadata Database Service on Exascale Infrastructure

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:  
  
Under **Oracle Exadata Database Service on Exascale 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 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**.

## 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure Database

1. Log onto the Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure Database

1. Log onto the Oracle Exadata Database Service on Exascale 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 Exascale Infrastructure

Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Exascale 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.
- [Default Backup Channel Allocation](#)  
These are the default settings for database backup channels when using "Oracle Managed Backup" or "User Configured Backup".
- [Prerequisites for Backups on Oracle Exadata Database Service on Exascale 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 \(RMAN\)](#)  
If you backed up your database using `bkup_api`, then you can manually restore that database backup by using the Oracle Recovery Manager (RMAN) utility.

## 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](#) Backups, configured in the Exadata Cloud Service console, API or `bkup_api` work for a variety of backup and recovery use cases.

## 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 or DB system) 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, then the Database service retries the operation during the next day's backup window. If an on-demand full backup fails, then you can try the operation again when the Oracle Exadata Database Service on Exascale 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, then you can control the retention period. The system automatically deletes backups when the assigned retention period is expired.

### Object Storage Backup retention period

The retention periods (in days) are 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 backup 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.

#### Note:

- **Data Guard**: You can enable the Automatic Backup feature on a database with the standby role in a Data Guard association. However, automatic backups for that database will not be created until it assumes the primary role.
- **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 was made starting with *dbaastools version 21.4.1.1.0*.

## Default Backup Channel Allocation

These are 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 core count of the node. The following table provides the values used and the core range, both the core 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.

Cores Per Node	Formula	Backup Channels Allocation Per Node	Restore Channels Allocation Per Node
Less than or equal to 12	Cores <= 12	2	4
Greater than 12 and less than or equal to 24	Cores > 12 and Cores <= 24	4	8
Greater than 24	Cores > 24	8	16

If needed, a static per node value can be set by using the DBAASCLI `getConfig/configure` to generate a `bkup_cfg` file, 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 core channel based allocation.

## Prerequisites for Backups on Oracle Exadata Database Service on Exascale Infrastructure

### Recovery Service

Ensure that your tenancy is configured to use Recovery Service.

**Table 5-4 Review the prerequisite tasks before you use Recovery Service as the automatic backup destination**

Task	More Information	Required or Optional
Create IAM policies	<a href="#">Policies to Enable Access to Recovery Service and Related Resources</a>	Required
Configure network resources and register a Recovery Service subnet	<a href="#">Creating a Recovery Service Subnet in the Database VCN</a>	Required
Create protection policies	<a href="#">Review Protection Policies for Database Backup Retention</a>	Optional

For more information about Recovery Service, see [Overview of Oracle Database Autonomous Recovery Service](#).

### Object Storage

- Exadata Cloud Service 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 Oracle Exadata Database Service on Exascale 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 configure automatic backups for a database

When you create an Oracle Exadata Database Service on Exascale 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 Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the database you want to configure:

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 **Configure Automatic Backups**.
6. In the Configure Automatic Backups dialog, enter the following details:
  - **Backup Destination:** Your choices are **Autonomous Recovery Service** (default) or **Object Storage**.
    - **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 Changes**.

The Database Details page displays the configuration details, **Health**, **Real-Time Data Protection**, and **Policy information** in the **Backup** section.

#### Related Topics

- [security zone compartment](#)
- [Security Zone Policies](#)

## To create an on-demand backup of a database

 **Note:**

Object Storage creates a full backup of the database while Recovery Service creates an incremental backup.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database you want to back up:  
Under **Oracle Exadata Database Service on Exascale 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**.

## To view backup status

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale 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 Database**, then click **Exadata on Oracle Public Cloud**.
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 Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database backup that you want to delete:  
Under **Oracle Exadata Database Service on Exascale 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 (   
  
    
  
 ) for the backup in which you are interested, and then click **Delete**.
7. Confirm when prompted.

## To delete standalone backups from Object Storage

1. Open the navigation menu. Click **Oracle 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 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 Database**, then click **Exadata on Oracle Public Cloud**.
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.

**DB systems:** Under Oracle Base Database, click **DB Systems**.

In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.

On the cloud **VM cluster** or **DB system** 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 DB system'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 Changes**.

## 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.

You can restore to:

- Restore to latest
- Restore to a timestamp
- Restore to SCN



### 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 Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster or DB system containing the database you want to restore:  
**Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model):** Under **Oracle Exadata Database Service on Exascale 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.  
**DB systems:** Under **Oracle Base Database**, click **DB Systems**. In the list of DB systems, find the Exadata DB system you want to access, and then click its name to display details about it.
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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure instance before performing the following procedures. For more information, see [#unique\\_255](#).

- [Default Backup Configuration](#)  
Oracle best practice guidelines for default backup configuration.
- [To create a backup configuration file](#)
- [To create an on-demand backup](#)
- [To remove the backup configuration](#)
- [To delete a local backup](#)
- [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 create a backup configuration file

 **Note:**

The following procedure must be performed on the first compute node in the Exadata Cloud Infrastructure VM cluster or DB system 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.

1. SSH to one of the database configured nodes in the VM cluster or DB system 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 -
```

3. 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>
```

4. Modify the parameters in the file to meet your requirements.

Parameter	Description
<code>bkup_disk=[yes no]</code>	Whether to back up locally to disk (Fast Recovery Area).
<code>bkup_oss=[yes no]</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> .
<code>bkup_oss_url=&lt;swift_url&gt;</code>	<p>Required if <code>bkup_oss=yes</code>.</p> <p>The Object Storage URL including the tenant and bucket you want to use. The URL is:</p> <pre>https:// swiftobjectstorage.&lt;region_name&gt;.ora clecloud.com/v1/&lt;tenant&gt;/&lt;bucket&gt;</pre> <p>Where:</p> <ul style="list-style-type: none"> <li>• <code>&lt;tenant&gt;</code> - lowercase tenant name (even if it contains uppercase characters) that you specify when signing in to the Console</li> <li>• <code>&lt;bucket&gt;</code> - name of the existing bucket you want to use for backups.</li> </ul>



Parameter	Description
<code>bkup_oss_user=&lt;oci_user_name&gt;</code>	<p>Required if <code>bkup_oss=yes</code>.</p> <p>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.</p> <p>For example, <b>jsmith@example.com</b> for a local user or <code>&lt;identity_provider&gt;/jsmith@example.com</code> for a federated user.</p> <p>To determine which type of user you have, see the following topics:</p> <ul style="list-style-type: none"> <li>• <a href="#">Managing Users</a> (for information on local users)</li> <li>• <a href="#">Federating with Identity Providers</a> (for information on federated users)</li> </ul> <p>Note that the user must be a member of the Administrators group, as described in <a href="#">#unique_262</a>.</p>
<code>bkup_oss_passwd=&lt;auth_token&gt;</code>	<p>Required if <code>bkup_oss=yes</code>.</p> <p>The <b>auth token</b> generated by using the Console or IAM API, as described in <a href="#">Prerequisites</a>.</p> <p>This is not the password for the Oracle Cloud Infrastructure user.</p>
<code>bkup_oss_recovery_window=n</code>	<p>Required if <code>bkup_oss=yes</code>.</p> <p>The number of days for which backups and archived redo logs are maintained in the Object Storage bucket. Specify 7 to 90 days.</p>
<code>bkup_daily_time=hh:mm</code>	<p>The time at which the daily backup is scheduled, specified in hours and minutes (hh:mm), in 24-hour format.</p>
<code>bkup_archlog_cron_entry=[yes no]</code>	<p>When no backups are configured using <code>dbaastools</code>, setting <code>bkup_archlog_cron_entry=no</code> will remove the archive log clean up job from <code>crontab</code>. The default value is <b>“yes”</b>.</p>

For example:

```
bkup_oss=[yes|no]
bkup_oss_url=<swift_url>
bkup_oss_user=<oci_user_name>
bkup_oss_passwd=<auth_token>
bkup_oss_recovery_window=n
bkup_daily_time=hh:mm
bkup_archlog_cron_entry=[yes|no]
bkup_cron_entry=[yes|no]
```

5. 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.

### 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.

- SSH to one of the database configured nodes in the VM cluster or DB system 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.

- Log in as `opc` and then `sudo` to the `root` user.

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

- 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>
```

- Exit the root-user command shell and disconnect from the compute node:

```
# exit
$ exit
```

- 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

- SSH to one of the database configured nodes in the VM cluster or DB system resource.
- Log in as `opc` and then `sudo` to the `root` user.
- Create a `temp` file with following parameters:

- `bkup_oss=no`
- `bkup_cron_entry=no`
- `bkup_archlog_cron_entry=no`

- 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>
```

- 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 local backup



### Note:

`bkup_api` is deprecated. Use `dbaascli database backup` and its option instead.

To delete a backup of a database deployment on the Oracle Exadata Database Service on Exascale Infrastructure instance, use the `bkup_api` utility.

1. Connect to the first compute node in your Exadata VM cluster or DB system resource as the `opc` user.

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. Start a root-user command shell:

```
$ sudo -s#
```

3. List the available backups:

```
# >/var/opt/oracle/bkup_api/bkup_api recover_list --dbname=<database_name>
```

where `dbname` is the database name for the database that you want to act on.

A list of available backups is displayed.

4. Delete the backup you want:

```
# /var/opt/oracle/bkup_api/bkup_api bkup_delete --bkup=<backup-tag> --dbname=<database_name>
```

where `backup-tag` is the tag of the backup you want to delete.

5. Exit the root-user command shell:

```
# exit
$
```

### 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](#)  
Learn how to use the API for database backups on Oracle Exadata Database Service on Exascale Infrastructure.

### Using the API to manage backups

Learn how to use the API for database backups on Oracle Exadata Database Service on Exascale Infrastructure.

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](#) Backups, configured in the Exadata Cloud Service console, API or `bkup_api` work for a variety of backup and recovery use cases.

## Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management

Backups, configured in the Exadata Cloud Service console, API or `bkup_api` work for a variety of backup and recovery use cases.

Backups, configured in the Oracle Exadata Database Service on Exascale Infrastructure console, API or `bkup_api` 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 *Oracle Database Backup and Recovery User's Guide*.

Managing backup and recovery, using RMAN, on Oracle Exadata Database Service on Exascale 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 `bkup_api` utility to disable cloud-managed backups, including disabling the automatic archive log purge job, by following this procedure:

 **Note:**

If you execute these steps, then the automation will no longer purge/backup the archive logs in the FRA for the database.

1. Connect as the `opc` user to the first compute node.

For detailed instructions, see *Connecting to a Compute Node with SSH*.

2. Start a root-user command shell:

```
sudo -s
```

3. Use the `bkup_api get config` command to generate a file containing the current backup settings for the database deployment:

```
/var/opt/oracle/bkup_api/bkup_api get config [--file=filename] --  
dbname=dbname
```

Where:

- `filename` is an optional parameter used to specify a name for the file that is generated
- `dbname` is the database name for the database that you want to act on

4. Edit the parameter values in the generated file to change the following parameters.

This will remove the backup crontab entries and disable all automatic backups. If the values are set to `yes`, then set to `no`.

```
bkup_cron_entry=no  
bkup_archlog_cron_entry=no  
bkup_nfs=no  
bkup_oss=no  
bkup_local=no
```

5. Use the `bkup_api set config` command to update the backup settings using the file containing your updated backup settings:

```
/var/opt/oracle/bkup_api/bkup_api set config --file=filename --  
dbname=dbname
```

Where:

- `filename` is an optional parameter used to specify a name for the file that is generated
- `dbname` is the database name for the database that you want to act on

The job to set the configuration will take several minutes to complete.

6. You can use the `bkup_api configure_status` command to check the status of the configuration update:

```
/var/opt/oracle/bkup_api/bkup_api configure_status --dbname=dbname
```

Where:

- *dbname* is the database name for the database that you want to act on

The **Configure backup status** starts as **running** and then moves to **finished** when complete.

7. Run the `bkup_api get config` command again and verify the settings listed above are set to `no`.

```
/var/opt/oracle/bkup_api/bkup_api get config [--file=filename] --  
dbname=dbname
```

Where:

- *filename* is an optional parameter used to specify a name for the file that is generated
- *dbname* is the database name for the database that you want to act on

#### Note:

After making these changes, no backups, including archive log backups, are made by the cloud automation. Ensure that manual RMAN backups are in place to avoid filling the archive log location.

#### Note:

Changes made using the `bkup_api` command are not reflected in the Oracle Exadata Database Service on Exascale Infrastructure console.

8. 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 Oracle Exadata Database Service on Exascale Infrastructure system by using a Secure Shell (SSH) connection.
- [Oracle Database Backup and Recovery User's Guide for Release 19](#)



## Recovering a Database Using Oracle Recovery Manager (RMAN)

If you backed up your database using `bkup_api`, then you can manually restore that database backup by using the Oracle Recovery Manager (RMAN) utility.

If you backed up your database using `bkup_api`, 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*.

### Note:

While recovering using RMAN is safe, you must not use RMAN to initiate backups or edit backup setting in conjunction with either `backup_api` usage or in conjunction with automated console backups. Doing so could result in conflicting conditions or overwrites of settings, and backups may not run successfully.

### Related Topics

- [Oracle Database Backup and Recovery User's Guide for Release 19](#)

## Patch and Update an Oracle Exadata Database Service on Exascale Infrastructure System

- [User-Managed Maintenance Updates](#)  
Maintaining a secure Oracle Exadata Database Service on Exascale Infrastructure instance in the best working order requires you to perform regular maintenance.
- [Patching and Updating an Oracle Exadata Database Service on Exascale Infrastructure System](#)  
Learn how to perform patching operations on Exadata database virtual machines and Database Homes.

## User-Managed Maintenance Updates

Maintaining a secure Oracle Exadata Database Service on Exascale Infrastructure instance in the best working order requires you to perform regular maintenance.

The following tasks are required

- 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.

# Patching and Updating an Oracle Exadata Database Service on Exascale Infrastructure System

Learn how to perform patching operations on Exadata database virtual machines and Database Homes.

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](#)  
Learn how to perform patching operations on Oracle Exadata Database Service on Exascale Infrastructure resources by using the Console or 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.
- [Upgrading Exadata Databases](#)  
Oracle Database releases on Oracle Exadata Database Service on Exascale Infrastructure can be upgraded using the Console and the API.

## Related Topics

- [Application Checklist for Continuous Service for MAA Solutions](#)

## Patching and Updating VM Cluster's GI and Database Homes

Learn how to perform patching operations on Oracle Exadata Database Service on Exascale Infrastructure resources by using the Console or API.



### 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.

- [About Patching and Updating VM Cluster's GI and Database Homes](#)  
Learn about types of patching performed on an Oracle Exadata Database Service on Exascale Infrastructure instances and how to complete the patching operations.
- [Prerequisites for Patching and Updating an VM Cluster](#)  
The Oracle Exadata Database Service on Exascale 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 Database Service on Exascale Infrastructure VM Clusters](#)  
You can use the Console to view the history of patch operations on Oracle Exadata Database Service on Exascale Infrastructure. Oracle Exadata Database Service on Exascale Infrastructure VM clusters apply patches, and monitor the status of patch operations.
- [To Upgrade the Oracle Grid Infrastructure of a Cloud VM Cluster](#)  
Procedure for upgrading the Oracle Grid Infrastructure of a Cloud VM Cluster.

- [Using the API to Patch an Oracle Exadata Database Service on Exascale 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

Learn about types of patching performed on an Oracle Exadata Database Service on Exascale Infrastructure instances and how to complete the patching operations.

- [Oracle Grid Infrastructure \(GI\) Patching](#)  
Patching an Oracle Exadata Database Service on Exascale Infrastructure instance updates the components on all the compute nodes in the instance. A VM cluster or DB system 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 Oracle Exadata Database Service on Exascale Infrastructure Components](#)

## Oracle Grid Infrastructure (GI) Patching

Patching an Oracle Exadata Database Service on Exascale Infrastructure instance updates the components on all the compute nodes in the instance. A VM cluster or DB system patch updates the Oracle Grid Infrastructure (GI) on the resource.



### Note:

You patch the Grid Infrastructure on the cloud VM cluster resource. VM clusters are used by the databases, which can be easily migrated to the new Grid Infrastructure resource with no system downtime.

## Database Home Patching

A Database Home patch updates the Oracle Database software shared by the databases in that home.

Patching requires moving the database to a new Database Home that has the correct patch version. This affects only the database being moved.

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 Oracle Exadata Database Service on Exascale 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 or an Exadata DB system 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, Exadata DB system, 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 enables you to easily roll back the database to the previous version by moving it back to the old 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, and the three most recent older versions ( $N$  through  $N - 3$ ).
- [dbaascli database runDatapatch](#)  
To patch an Oracle Database, use the `dbaascli database runDatapatch` command.
- [Customer-Managed Keys in Oracle Exadata Database Service on Exascale Infrastructure](#)  
Customer-managed keys for Oracle Exadata Database Service on Exascale 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

#### Frequently Asked Questions

##### Q: What is the purpose of the `dbaascli database runDatapatch` command?

A: The `dbaascli database runDatapatch` command is used to apply patches to an Oracle Database.

##### Q: What must be ensured before running the `dbaascli database runDatapatch` command?

A: Before running the command, ensure that all instances of the database are up and running.

##### Q: How do I specify which database to patch?

A: Use the `--dbname` option followed by the name of the database. For example:

```
--dbname myDatabase
```

##### Q: How do I resume a previously interrupted `runDatapatch` operation?

A: Use the `--resume` option to resume the previous run or the `--sessionID` option to specify a specific session ID. For example:

```
--resume
--sessionID 12345
```

##### Q: How can I skip certain PDBs when running the patch?

A: Use the `--skipPdfs` option followed by a comma-delimited list of PDB names to skip. For example:

```
--skipPdfs pdb1,pdb2
```

##### Q: How can I run the patch only on certain PDBs?

A: Use the `--pdb`s option followed by a comma-delimited list of PDB names to include. For example:

```
--pdbname pdb1,pdb2
```

**Q: How do I apply or roll back a specific set of patches?**

A: Use the `--patchList` option followed by a comma-delimited list of patch names to apply or roll back. For example:

```
--patchList patch1,patch2
```

**Q: What does the --rollback option do?**

A: The `--rollback` option rolls back the patches that were applied during the patching operation.

**Q: What happens if some PDBs are closed during the patching operation?**

A: If some PDBs are closed, you can use the `--skipClosedPdb`s option to skip patching those closed PDBs.

**Q: Can I run prerequisite checks before applying patches?**

A: Yes, use the `--executePrereqs` option to run prerequisite checks before applying the patch.

**Q: How do I find out which session ID to resume a patch?**

A: After a `runDatapatch` operation, the session ID is typically logged. Use the `--sessionID` option to specify that ID when resuming a patch. For example:

```
--sessionID 67890
```

```
dbaascli database runDatapatch --dbname db19
```

## Customer-Managed Keys in Oracle Exadata Database Service on Exascale Infrastructure

Customer-managed keys for Oracle Exadata Database Service on Exascale 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 do the following:

- Enable customer-managed keys when you create databases in Oracle Exadata Database Service on Exascale 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 Oracle Data Guard on Oracle Exadata Database Service on Exascale Infrastructure databases that use customer-managed keys, the primary and standby databases must be in the same [realm](#).

### Related Topics

- [To create a database in an existing Exadata Cloud Service instance](#)
- [To administer Vault encryption keys](#)

### 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

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli database addInstance` command?

A: The `dbaascli database addInstance` command is used to add a new database instance to a specified node in an Oracle Exadata Database Service environment.

#### Q: What are the prerequisites for running the `dbaascli database addInstance` command?

A: The command must be run as the `root` user to have the necessary permissions to add a database instance.

#### Q: What does the `--dbname` option represent in this command?

A: The `--dbname` option specifies the name of the Oracle Database for which you want to add a new instance.

#### Q: What is the `--node` option used for in the `dbaascli database addInstance` command?

A: The `--node` option specifies the name of the node where the new database instance will be added.

#### Q: What is the purpose of the `--newNodeSID` option in this command?

A: The `--newNodeSID` option allows you to specify the SID (System Identifier) for the new database instance that will be created on the specified node.

**Q: Is it mandatory to specify the `--newNodeSID` option when adding a new instance?**

A: The `--newNodeSID` option is optional. If not provided, Oracle will automatically generate an SID for the new database instance.

**Q: When should I use the `dbaascli database addInstance` command?**

A: Use this command when you want to scale your database by adding a new instance to an additional node in a multi-node Oracle Database setup.

**Q: Can I add multiple database instances to different nodes using this command?**

A: Yes, you can run the command multiple times to add database instances to different nodes by specifying the appropriate `--node` and `--dbname` values.

**Q: What happens if the node specified in the `--node` option is not available?**

A: The command will fail if the specified node is not available or reachable. Ensure that the node is properly configured and accessible before running the command.

**Q: Can this command be used in a Data Guard environment?**

A: Yes, you can use the `dbaascli database addInstance` command in a Data Guard environment to add instances, but it is recommended to follow the necessary Data Guard guidelines for such configurations.

**Q: Will this command cause database downtime?**

A: Adding an instance to a new node typically does not cause downtime for the existing database instances, but it's recommended to check your environment for any specific dependencies.

**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 5-2 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.

#### Frequently Asked Questions

##### Q: What is the purpose of the dbaascli database getDetails command?

A: The `dbaascli database getDetails` command shows detailed information about a specified Oracle database, including the database name, node information, and pluggable database (PDB) details.

##### Q: Who can run the dbaascli database getDetails command?

A: The command can be run by the `root` user or the `oracle` user.

##### Q: What does the --dbname option specify in the dbaascli database getDetails command?

A: The `--dbname` option specifies the name of the Oracle database for which detailed information is being retrieved.

**Q: What kind of information does the `dbaascli database getDetails` command provide?**

A: The command provides details such as the database name, node information, and information about pluggable databases (PDBs) associated with the container database.

**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>[--instance <value>] [--backupPrepared] [--
allowBounce]|
--resetParameters <values> [--instance <value>] [--backupPrepared] [--
allowBounce]
}
--responseFile
[--backupPrepared]
[--instance]
[--allowBounce]
[--waitForCompletion]
```

**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.
- `--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.
- `--backupPrepared` acknowledges that a proper database backup is in place prior to modifying critical or sensitive parameters.
- `--allowBounce` grants permission to bounce the database in order to reflect the changes on applicable static parameters.
- `--waitForCompletion` specify `false` to run the operation in background. Valid values : `true|false.`]

**Frequently Asked Questions**

**Q: What is the purpose of the `dbaascli database modifyParameters` command?**

A: The `dbaascli database modifyParameters` command is used to modify or reset Oracle Database initialization parameters.

**Q: How do I specify the database for which I want to modify parameters?**

A: You must use the `--dbname` option to specify the name of the database for which you want to modify or reset parameters.

**Q: How can I modify database parameters using the `modifyParameters` command?**

A: Use the `--setParameters` option followed by a comma-delimited list of parameters and their new values. For example:

```
--setParameters parameter1=valueA,parameter2=valueB
```

**Q: How do I reset parameters to their default values using this command?**

A: Use the `--resetParameters` option followed by a comma-delimited list of parameters to reset to their default values. For example:

```
--resetParameters parameter1,parameter2
```

**Q: Can I modify parameters using a response file?**

A: Yes, you can specify the absolute location of a response JSON file using the `--responseFile` option. The file should contain the parameters you want to modify.

**Q: Is it necessary to take a backup before modifying parameters?**

A: While not mandatory for all changes, if you are modifying critical or sensitive parameters, it's recommended to have a backup in place. You can use the `--backupPrepared` option to acknowledge that a backup has been prepared.

**Q: Can I apply changes only to a specific instance in a multi-instance database?**

A: Yes, you can specify the instance name using the `--instance` option. If this option is not used, the changes will be applied at the database level.

**Q: Will the database need to be bounced (restarted) after modifying parameters?**

A: For some static parameters, a database bounce is required. You can use the `--allowBounce` option to grant permission for the database to bounce if necessary.

**Q: What happens if I don't allow the database to bounce when changing static parameters?**

A: If you do not use the `--allowBounce` option when modifying static parameters, the changes will not take effect until the next manual database restart.

**Q: Can I resume modifying parameters if an earlier session was interrupted?**

A: No, this command does not support session resumption. You will need to re-run the command from the beginning.

**Example 5-3 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.

## Frequently Asked Questions

### Q: What is the purpose of the `dbaascli database upgrade` command?

A: The `dbaascli database upgrade` command is used to upgrade an Oracle Database to a new version.

### Q: What are the prerequisites for using the `dbaascli database upgrade` command?

A: You must run the command as the `root` user and connect to an Exadata Cloud@Customer virtual machine using SSH.

**Q: How do I specify the database that needs to be upgraded?**

A: Use the `--dbname` option followed by the name of the database. For example:

```
--dbname myDatabase
```

**Q: How can I specify the target Oracle home for the upgrade?**

A: You can specify the target Oracle home location with the `--targetHome` option or the name of the target Oracle Database home with the `--targetHomeName` option.

**Q: What does the `--standBy` option do?**

A: The `--standBy` option is used to upgrade standby databases in Data Guard configurations.

**Q: What is the purpose of the `--allStandbyPrepared` flag?**

A: The `--allStandbyPrepared` flag acknowledges that all required operations on standby databases have been performed before upgrading the primary database in a Data Guard configuration.

**Q: What does the `--removeGRP` option do?**

A: The `--removeGRP` option automatically removes the Guaranteed Restore Point (GRP) backup if the database upgrade is successful.

**Q: When should I use the `--increaseCompatibleParameter` option?**

A: Use the `--increaseCompatibleParameter` option to automatically increase the compatible parameter during the database upgrade, provided the upgrade is successful.

**Q: What does the `--executePrereqs` option do?**

A: The `--executePrereqs` option runs only the pre-upgrade checks to ensure that the database is ready for the upgrade.

**Q: How do I handle a failed post-upgrade step?**

A: Use the `--postUpgrade` option to rerun the post-upgrade steps if the initial post-upgrade attempt fails.

**Q: What is the purpose of the `--revert` option?**

A: The `--revert` option reverts the Oracle Database to its original Oracle home, undoing the upgrade.

**Q: How can I pass additional arguments specific to DBUA for the upgrade?**

A: Use the `--upgradeOptions` option to pass DBUA-specific arguments for the Oracle Database upgrade. Refer to the Oracle documentation for supported arguments and options.

**Q: Is it mandatory to specify the target Oracle home for the upgrade?**

A: Yes, you must specify either the `--targetHome` or `--targetHomeName` to indicate the target Oracle home for the upgrade.

**Q: What should I do if I need to perform a pre-upgrade check but not proceed with the upgrade?**

A: Use the `--executePrereqs` option to perform only the pre-upgrade checks without proceeding with the actual upgrade.

**Example 5-4 dbaascli database upgrade pre-upgrade requisite checks**

```
dbaascli database upgrade --dbname dbname --targetHome Target Oracle home
location --executePrereqs
```

**Prerequisites for Patching and Updating an VM Cluster**

The Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Database Service on Exascale Infrastructure VM Clusters**

You can use the Console to view the history of patch operations on Oracle Exadata Database Service on Exascale Infrastructure Oracle Exadata Database Service on Exascale Infrastructure VM clusters apply patches, and monitor the status of patch operations.

- [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 individual Oracle Databases in Oracle Exadata Database Service on Exascale Infrastructure](#)  
You can patch a single Oracle Database in your Oracle Exadata Database Service on Exascale Infrastructure by moving it to another Database Home.
- [Viewing Patch History of Exadata Database Service on Exascale Infrastructure](#)  
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.

**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 Database**, then click **Exadata Database Service on Exascale 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. Under **Version**, click the **View Patches** link beside the **Updates Available** field.
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:
  - **Run Precheck:** Check for any prerequisites to make sure that the patch can be successfully applied.
  - **Update Grid Infrastructure:** 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**.

9. View or Download Logs for the move operation

To view the status of a job or to download the logs for a job, use this procedure.

To view the status of a job or to download the logs for a job, use this procedure.

- a. Go to the **VM Cluster details** page of the cluster by clicking the cluster name for which you want to check the job logs for the Grid Infrastructure patch operation. The **Grid Infrastructure details** page is displayed.
- b. Under **Resources**, click **Associated resources**. There will be a resource with the name **fsujob<\*>**. Click the Actions icon (  ), and then click *View log*. The **View log** page is displayed.

The log is refreshed every two minutes automatically. Click **Refresh log** to refresh the logs on demand. Click *Download log* to download the log.

10. If necessary, you can retry a failed Oracle Grid Infrastructure operation.

The steps to retry a failed Grid Infrastructure operation are as follows:

- a. Click **View or Download Job Logs**. Check the job logs for the Apply ExaDB VM Cluster GI Patch operation to understand the reason for the failure. See: [Incident Logs and Trace Files](#).
- b. Resolve the issue. After addressing the underlying issue, you can retry the failed operation.
- c. Retry the Grid Infrastructure updates Operation: To retry the failed Grid Infrastructure patching, initiate the Apply Grid Infrastructure Patch operation again, selecting the same Grid Infrastructure update as in the previous attempt.

## To patch individual Oracle Databases in Oracle Exadata Database Service on Exascale Infrastructure

You can patch a single Oracle Database in your Oracle Exadata Database Service on Exascale Infrastructure 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 more information, see:

### [To move a database to another Database Home](#)

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database you want to move.:  
Under **Oracle Exadata Database Service on Exascale 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**, and then click **Move to Another Home**.
5. Select the target Database Home.
6. Click **Move Database**.

#### **Note:**

When you update the software release of databases by moving them to a target Database Home, Oracle recommends that you use **Database Homes**, which run the latest (N) to 3 versions from the latest (N-3) release update (RU) versions. Only database homes provisioned with database Release Updates (RUs) that meet this best practice criterion are available as target homes to move your database.


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, the Database Home is updated with the current home. Datapatch is run 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.

8. View or Download Logs for the move operation.

To view the status of a job or to download the logs for a job, use this procedure.



- a. Go to the **VM Cluster details** page of the cluster where the database is configured and click the name of the database for which you want to check the job logs for the move operation. The Database details page is displayed.
  - b. Under **Resources**, click **Work requests**.
  - c. In the **Work requests** section, click the **Update Database Operation** for which you want to view or download logs. The **Update Database Work request details** page is displayed.
  - d. Under **Resources**, click **Associated resources**. There will be a resource with **fsujob<\*> name**.
  - e. Click the Actions icon (  ), and then click **View log**. The **View log** page is displayed. The log is refreshed every two minutes automatically. Click **Refresh log** to refresh the logs on demand. Click **Download log** to download the log.
9. If necessary, you can retry a failed Database Move operation.
- The steps to retry a failed database move operation are as follows:
- a. Click **View or Download Job Logs**. Check the job logs for the Database Move to understand the reason for the failure. See: [Incident Logs and Trace Files](#).
  - b. Resolve the issue. After addressing the underlying issue, you can retry the failed move operation.
  - c. Retry the database move operation. To retry the failed Database Move operation, initiate the **Move Database** step again, selecting the same target home as in the previous attempt.

#### Related Topics

- [To move a database to another Database Home](#)  
To patch a single Oracle Database in your Oracle Exadata Database Service on Exascale Infrastructure instance, you move it to another Database Home.

## Viewing Patch History of Exadata Database Service on Exascale Infrastructure

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.

You can view patch history by navigating to the VM Cluster Details page.

Patch history views in the Console do not show patches that were applied by using command line tools such as `dbaascli`.

- [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](#)

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 Database**, then click **Exadata Database Service on Exascale 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. Under **Version**, click the **View Patches** link beside the **Updates Available** field.
6. 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. When 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 Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster that contains the Database Home.

Under **Oracle Exadata Database Service on Exascale 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. Click the name of the Database Home you want to view to display the Database Home details.
6. Under **Database Software Version**, click **View** by the **Latest Patch Available** field.
7. 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.

## 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 23ai, 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 23ai:
  - 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 Database**, then click **Exadata Database Service on Exascale 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. Under **Version**, click the **View Patches** link beside the **Updates Available** field.
6. Click **Updates** to view the list of available patches and upgrades.
7. Click the Actions icon (three dots) at the end of the row listing the Oracle Grid Infrastructure (GI) upgrade, then click **Upgrade Grid Infrastructure**.
8. 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 Patch an Oracle Exadata Database Service on Exascale 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:

- [GetExadbVmClusterUpdate'](#)
- [ListExadbVmClusterUpdates](#)

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, Exadata DB system, 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.

- [Updating the Operating System using the Console](#)

## Updating the Operating System using the Console

 **Note:**

After 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 Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Under **Oracle Exadata Database Service on Exascale 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. In the **Version** section, to the right of the **Updates Available**, click **View Updates** to display the **Updates** page.
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:
  - **Run 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 Databases

Oracle Database releases on Oracle Exadata Database Service on Exascale Infrastructure can be upgraded 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.

- [Prerequisites to Upgrade Oracle Databases](#)  
Review the list of prerequisites to upgrade an Oracle Exadata Database Service on Exascale 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:

#### Related Topics

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

## Prerequisites to Upgrade Oracle Databases

Review the list of prerequisites to upgrade an Oracle Exadata Database Service on Exascale Infrastructure Oracle Database instance.

- You must have an available Oracle Database Home that uses the four most recent versions of Oracle Database available. See *To Create a new Oracle Database Home in an existing Oracle Exadata Database Service on Exascale 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.
  - Redo apply needs to be disabled during the upgrade of both the primary and standby.
  - If you have configured an observer, then the observer needs to be disabled prior to upgrade.

Before you start the upgrade, your Oracle Database configuration must be configured with the following settings:

- The database must be in archive log mode.
- The database must have flashback enabled.

To learn more about these settings, see the Oracle Database documentation for your database release.

#### Related Topics

- [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.
- [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:

## 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 new Database home. These variables should be automatically updated during the upgrade process.

- If you are upgrading a non-container database, 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, then 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 \(Doc ID 2288024.1\)](#)

## 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 the upgrade history of a database](#)  
To view upgrade history for databases on Exadata Database Service on Exascale Infrastructure, use this procedure.

## 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.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Under **Oracle Exadata Database Service on Exascale 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.
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 **More Actions**, then **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 Database**, click **Exadata Database Service on Exascale 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 the upgrade history of a database

To view upgrade history for databases on Exadata Database Service on Exascale Infrastructure, use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.
3. Under **Oracle Exadata Database Service on Exascale 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.
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. On the Database Details page, under **Database Version**, click the **View** link that is displayed for databases that have been upgraded. This link does not appear for databases that have not been updated.  
The **Updates History** page is displayed. The table displayed on this page shows precheck and upgrade operations performed on the database.

## 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 Oracle Exadata Database Service on Exascale Infrastructure database, you must specify `DB_HOME` as the upgrade source.

## Manual Software Updates

For authorized environments, learn how to download manual 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 manual 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 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 Software Update

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
2. Under **Resources**, click **Manual software updates**.  
Manual software update page is displayed.
3. Click **Create software update**.  
Create 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. **One-off patch number**: Optionally, enter a one-off (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 Update details page, click **Download**.  
The system starts downloading the patch.
2. You can also download a patch from the Interim Software Updates 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 Update details page, click **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 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 Update details page, click **Move Resource**.
2. In the resulting dialog, choose a new compartment, and click **Move Resource**.
3. You can also move a patch resource from the Interim Software Updates 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 Oracle Exadata Database Service on Exascale Infrastructure

Learn to configure and manage Data Guard groups in your VM cluster.

- [About Using Oracle Data Guard with Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure](#)  
An 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 Data Guard standby database.
- [Working with Oracle Data Guard](#)  
Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.
- [Using the Console to Manage Oracle Data Guard Associations](#)  
Learn how to enable a Data Guard association between databases, change the role of a database in a Data Guard association 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 Oracle Exadata Database Service on Exascale Infrastructure instance:

## About Using Oracle Data Guard with Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure

An 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 Data Guard standby database.

When enabling Oracle Data Guard, you can create a new Database Home on the standby Exadata instance to house the new standby database during the enable Data Guard operation. Alternately, you can choose to provision the standby database in an existing Database Home on the standby 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.

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. :

If you are creating an Oracle Data Guard Association 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](#)  
Ensure that you meet the requirements for using Oracle Exadata Database Service on Exascale Infrastructure with Oracle Data Guard.
- [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](#)  
If node addition is done either on the standby database or the primary database, the metadata must be updated manually on the database other than the one where the node was added.
- [Removing a Node from a VM Cluster](#)  
If node removal is done either on the standby database or the primary database, the metadata must be updated manually on the database other than the one where the node was removed.

## Network Requirements for Data Guard

Ensure that you meet the requirements for using Oracle Exadata Database Service on Exascale Infrastructure with Oracle Data Guard.

Ensure that your environment meets the following network requirements:

- The primary and standby databases can be part of VM clusters in different compartments.
- The primary and standby databases must, however, be part of the same VCN within the same region.
- 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 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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

- [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

If node addition is done either on the standby database or the primary database, the metadata must be updated manually on the database other than the one where the node was added.

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.

## Removing a Node from a VM Cluster

If node removal is done either on the standby database or the primary database, the metadata must be updated manually on the database other than the one where the node was removed.

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.

## Working with Oracle Data Guard

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.

The Data Guard implementation requires two databases, one in a primary role and one in a standby role. The two databases compose a Data Guard association. Most of your applications access the primary database. The 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 switch or fail over the standby database to the primary role.

- [Switchover](#)  
A switchover reverses the primary and standby database roles.
- [Failover](#)  
With Oracle Data Guard, a failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable.
- [Reinstat](#)  
The `reinstat` command reinstates a database into the standby role in an Oracle Data Guard association.

## 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

With Oracle Data Guard, a failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable.

A failover might result in some data loss when you use **Maximum Performance** protection mode.

## Reinstate

The `reinstate` command reinstates a database into the standby role in an Oracle Data Guard association.

You can use the `reinstate` command to return a failed database into service after correcting the cause of failure.



### Note:

You can't terminate a primary database that has a Data Guard association with a peer (standby) database. Delete the standby database first. Alternatively, you can switch over the primary database to the standby role, and then terminate the former primary.

You can't terminate a VM cluster that includes Data Guard-enabled databases. You must first remove the Data Guard association by terminating the standby database.

## Using the Console to Manage Oracle Data Guard Associations

Learn how to enable a Data Guard association between databases, change the role of a database in a Data Guard association using either a switchover or a failover operation, and reinstate a failed database.

When you enable Data Guard, a separate Data Guard association is created for the primary and the standby database.

- [To enable Data Guard on Exadata Database Service on Exascale Infrastructure](#)  
Learn how to enable Data Guard association between databases.
- [To view Data Guard associations of databases in a Cloud VM Cluster](#)  
To view the role of each database in a Data Guard association 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 initiate a switchover operation by using the Data Guard association of the primary database.
- [To edit the Oracle Data Guard association](#)  
You edit the Oracle Data Guard association to configure the Data Guard protection for the primary database.
- [To perform a database failover](#)  
You initiate a failover operation by using the Data Guard association of the standby database.
- [To reinstate a database](#)
- [To terminate a Data Guard association on an Oracle Exadata Database Service on Exascale Infrastructure instance](#)  
On an Oracle Exadata Database Service on Exascale Infrastructure instance, you remove a Data Guard association by terminating the standby database.

## To enable Data Guard on Exadata Database Service on Exascale Infrastructure

Learn how to enable Data Guard association between databases.



### Note:

When you enable Data Guard, replication of data happens only over the client network.

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment** that contains the Oracle Exadata Database Service on Exascale 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:
4. Under **Exadata Database Service on Exascale Infrastructure**, click **Exadata VM clusters**. In the list of VM clusters, find the VM cluster that you want to access and click its highlighted name to view the details page for the cluster.
5. On the VM cluster details page, in the **Databases** section, click the name of the database that you want to make primary.
6. On the Database Details page, under **Resources**, click **Data Guard Associations**.
7. In the **Data Guard Associations** section, click **Enable Data Guard**.
8. On the Enable Data Guard page, configure your Data Guard association.
  - In the **Select 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:
    - **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.
- **Data Guard peer resource type:** Select **VM Cluster**.  
Select a VM cluster from the drop-down list.
- **Data Guard association 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 association.  
See [Redo Transport Services](#) for information on these options.
- 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 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**.

 **Note:**

Oracle recommends applying the same list of patches to the Database Homes of the primary and standby databases.

- 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.

9. *Optional.* **Enable thin clone:** Select this option to leverage Exascale redirect-on-write technology to create a thin clone of the PDB. This option results in the reuse of duplicate blocks with the parent PDB, shared with the clone. Deselecting this option results in a traditional, full clone with all blocks copied, and fully independent from the parent.
10. 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`.

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
11. Click **Enable Data Guard**. 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.

When the association is created, the details for a database and its peer display their respective roles as **Primary** or **Standby**.

- [View Data Guard Provisioning Progress](#)  
View the progress of Data Guard Provisioning tasks using the Work Requests page.

**Related Topics**

- [Network Setup for Oracle Exadata Database Service on Exascale Infrastructure Instances](#)  
This topic describes the recommended configuration for the VCN and several related requirements for the Oracle Exadata Database Service on Exascale Infrastructure instance.
- [Changing the Database Passwords](#)  
To change the SYS password, or to change the TDE wallet password, use this procedure.

## View Data Guard Provisioning Progress

View the progress of Data Guard Provisioning tasks using the Work Requests page.

After you have completed the task To Enable Data Guard, multiple work requests are issued to complete the provisioning of the Data Guard group. To view the progress of these work requests:

1. Navigate to the **Work Requests Details** page. On the **Work Requests Details** page there is a bar in the Work Request Information tab that shows the overall progress of the Data Guard Provisioning
2. Under **Resources**, select **Log Messages**. The table shows a message for each task that is completed or in progress.

## To view Data Guard associations of databases in a Cloud VM Cluster

To view the role of each database in a Data Guard association in an Cloud VM Cluster, follow this procedure.

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Exascale 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.

## To enable automatic backups on a standby database

Learn to enable automatic backups on a standby database.

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Exascale 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 or DB system that contains the primary database. Under **Oracle Exadata Database Service on Exascale 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 page, in the **Databases** section, click the name of the primary database.
5. On the Database Details page, under **Resources**, click **Data Guard Associations**.
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. If your database is in a security zone compartment, you must enable automatic backups.
  - **Backup Scheduling:**
    - **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 database backup includes all datafiles, control file, and parameter files associated with the target database. Archive backups are separate and decoupled and executed every 30 minutes. You can choose to execute the first full backup immediately or defer to the assigned full backup scheduling time. If you defer to the latter, the database will not be recoverable until the first backup completes.
  - **Backup Destination:** Object Storage is selected by default and you cannot change it.

 **Note:**

- If automatic backup is enabled on the primary database and the backup destination is Autonomous Recovery Service, then you cannot enable backup on the standby database.
- If automatic backup is enabled on the primary database and the backup destination is Object Storage, then you can enable backup on the standby database. Note that you can only select Object Storage as the backup destination.
- If automatic backup is disabled on the primary database, you can still enable backup on the standby database by selecting Object Storage as the backup destination.

9. Click **Save Changes**.

## To perform a database switchover

You initiate a switchover operation by using the Data Guard association of the primary database.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose the **Compartment** that contains the Oracle Exadata Database Service on Exascale Infrastructure instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster or DB system that contains the Data Guard association: **Oracle Exadata Database Service on Exascale 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 Associations**.
5. For the Data Guard association 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 **OK**.

This database should now assume the role of the standby, and the standby should assume the role of the primary in the Data Guard association.

## To edit the Oracle Data Guard association

You edit the Oracle Data Guard association to configure the Data Guard protection for the primary database.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale 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 or DB system that contains the Data Guard association:
 

Under **Oracle Exadata Database Service on Exascale 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 Associations**.
5. For the Data Guard association you want to manage, click the Actions menu (
 

⋮

 ), and then click **Edit Protection Mode**.
6. In the **Edit Data Guard Association** panel, configure the Data Guard association:
  - **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 Oracle Data Guard association.
  - **Database admin password:** Enter the ADMIN password for the database.
7. Click **Save**.

### Related Topics

- [Oracle Data Guard Protection Modes](#)



## To perform a database failover

You initiate a failover operation by using the Data Guard association of the standby database.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose the **Compartment** that contains the Oracle Exadata Database Service on Exascale 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:  
Under **Oracle Exadata Database Service on Exascale 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 Associations**.
5. For the Data Guard association on which you want to perform a failover, click **Failover**.
6. In the **Failover Database** dialog box, enter the database admin password, and then click **OK**.

This database should now assume the role of the primary, and the old primary's role should display as **Disabled Standby**.

## 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 by using its Data Guard association.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**
2. Choose the **Compartment** that contains the Oracle Exadata Database Service on Exascale Infrastructure with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster or DB system that contains the Data Guard association:  
Under **Oracle Exadata Database Service on Exascale 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 Associations**.
5. For the Data Guard association 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 **OK**.

This database should now be reinstated as the standby in the Data Guard association.

## To terminate a Data Guard association on an Oracle Exadata Database Service on Exascale Infrastructure instance

On an Oracle Exadata Database Service on Exascale Infrastructure instance, you remove a Data Guard association by terminating the standby database.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata Database Service on Exascale Infrastructure**.
2. Choose the **Compartment** that contains the Oracle Exadata Database Service on Exascale Infrastructure VM cluster with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster that contains the standby database:  
Under **Oracle Exadata Database Service on Exascale Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster that 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 (  ), and then click **Terminate**.
5. In the **Terminate Database** dialog box, enter the name of the database, and then click **OK**.

## Using the API to manage Data Guard associations

Use these API operations to manage Data Guard associations on an Oracle Exadata Database Service on Exascale 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](#).

- [CreateDataGuardAssociation](#)
- [ListDataGuardAssociations](#)
- [GetDataGuardAssociation](#)
- [UpdateDataGuardAssociation](#)
- [SwitchoverDataGuardAssociation](#)
- [FailoverDataGuardAssociation](#)
- [ReinstateDataGuardAssociation](#)
- [DeleteDatabase](#) - To terminate an Oracle Exadata Database Service on Exascale Infrastructure instance Data Guard association, you delete the standby database.

For the complete list of APIs for the Database service, see [Database Service API](#).

## Configure Oracle Database Features for Oracle Exadata Database Service on Exascale Infrastructure

Learn how to configure Oracle Multitenant, tablespace encryption, and other options for your Oracle Exadata Database Service on Exascale Infrastructure instance.

- [Using Oracle Multitenant on an Oracle Exadata Database Service on Exascale Infrastructure Instance](#)  
Learn about requirements for different features when using Multitenant environments in Oracle Exadata Database Service on Exascale Infrastructure.

- [Managing Tablespace Encryption](#)  
Learn about how tablespace encryption is implemented in Oracle Exadata Database Service on Exascale Infrastructure

## Using Oracle Multitenant on an Oracle Exadata Database Service on Exascale Infrastructure Instance

Learn about requirements for different features when using Multitenant environments in Oracle Exadata Database Service on Exascale Infrastructure.

When you create an Oracle Exadata Database Service on Exascale Infrastructure Instance, an Oracle Multitenant environment is created.

The multitenant architecture enables 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.

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 Oracle Exadata Database Service on Exascale 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 "Using Transparent Data Encryption with Other Oracle Features" in *Oracle Database Advanced Security Guide*.

See "ADMINISTER KEY MANAGEMENT" in *Oracle Database SQL Language Reference*.

- [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](#)

### Related Topics

- [Using Transparent Data Encryption with Other Oracle Features in Oracle Database Advanced Security Guide](#)
- [ADMINISTER KEY MANAGEMENT in Oracle Database SQL Language Reference](#)

## 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

Learn about how tablespace encryption is implemented in Oracle Exadata Database Service on Exascale Infrastructure

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 \(Doc ID 2359020.1\)](#).

### 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*.

### 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.

**Related Topics**

- [Oracle Database Advanced Security Guide Release 19c](#)
- [Oracle Database Advanced Security Guide Release 18c](#)
- [Oracle Database Advanced Security Guide Release 12c \(12.2\)](#)

## Migrate to Oracle Exadata Database Service on Exascale Infrastructure

For general guidance on methods and tools to migrate databases to Oracle Cloud Infrastructure database services, including Oracle Exadata Database Service on Exascale Infrastructure see "Migrating Databases to the Cloud".

A recommended approach for migrating to Oracle Exadata Database Service on Exascale Infrastructure is using Zero Downtime Migration

**Related Topics**

- [Migrating Databases to the Cloud](#)

## Connect Identity and Access Management (IAM) Users to Oracle Exadata Database Service on Exascale Infrastructure

You can configure Exadata Database Service on Exascale 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 Exascale 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.
- [Enable, Disable, and Re-enable Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Authentication on Oracle Database](#)  
Learn to enable, disable, and re-enable Identity and Access Management (IAM) Authentication on Oracle Database.



- [Manage Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Groups and Policies, Users, Roles, and Database Passwords](#)  
Your Oracle Exadata Database Service on Exascale Infrastructure system provides several different methods of service management.
- [Configuring Client Connection](#)  
Configure various clients to use IAM authentication.

## 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 Exascale 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 Database 23ai 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 users 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 Database 23ai 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 use 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 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:
    - **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](#)

## Enable, Disable, and Re-enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database

Learn to enable, disable, and re-enable Identity and Access Management (IAM) Authentication on Oracle Database.

- [Enable Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Authentication on Oracle Database](#)  
Review the steps to enable or re-enable IAM user access to Oracle Database.
- [Disable Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Authentication on Oracle Database](#)  
Describes the steps to disable IAM external authentication user access for Oracle Database.

- [Using Oracle Database Tools with Identity and Access Management \(IAM\) Authentication](#)  
Review the notes for using Oracle Database tools with IAM authentication enabled.

## Enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database

Review the steps to enable or re-enable IAM user access to Oracle Database.

### 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.

1. Perform the prerequisites for IAM authorization and authentication on Oracle Database. See *Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database* for more information.
2. Enable Oracle Cloud Infrastructure (IAM) Authentication and Authorization using the `ALTER SYSTEM` command.

```
ALTER SYSTEM SET IDENTITY_PROVIDER_TYPE=OCI_IAM SCOPE=BOTH;
```

3. Verify the value of `IDENTITY_PROVIDER_TYPE` system parameter.

```
SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME='identity_provider_type';
```

NAME	VALUE
-----	-----
identity_provider_type	OCI_IAM

### Related Topics

- [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.
- [Disable External Authentication Scheme](#)  
Review the prerequisites for enabling IAM user access to Oracle Database.

## Disable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database

Describes the steps to disable IAM external authentication user access for Oracle Database.

To disable IAM user access on your Oracle Database instance:



1. Disable IAM integration using the `ALTER SYSTEM` command.

```
ALTER SYSTEM RESET IDENTITY_PROVIDER_TYPE SCOPE=BOTH;
```

2. If you also want to remove the IAM policy to allow database access, you may need to review and either modify or remove the IAM groups and the policies you set up to allow access to the database by IAM users.

## Using Oracle Database Tools with Identity and Access Management (IAM) Authentication

Review the notes for using Oracle Database tools with IAM authentication enabled.

- Oracle APEX is not supported for IAM users with Oracle Database.
- Database Actions is not supported for IAM users with Oracle Database. See *Provide Database Actions Access to Database Users* for information on using regular database users with Oracle Database.
- Oracle Machine Learning Notebooks and other components are not supported for IAM Authorized users with Oracle Database. See *Add Existing Database User Account to Oracle Machine Learning Components* for information on using regular database users with Oracle Database.

### Related Topics

- [Provide Database Actions Access to Database Users](#)
- [Add Existing Database User Account to Oracle Machine Learning Components](#)

## Manage Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies, Users, Roles, and Database Passwords

Your Oracle Exadata Database Service on Exascale Infrastructure system provides several different methods of service management.

- [Create Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Groups and Policies for IAM Users](#)  
Review the steps to write policy statements for an IAM group to enable IAM user access to Oracle Cloud Infrastructure resources, specifically Oracle Database instances using IAM database tokens.
- [Authorize Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Users on Oracle Database](#)  
Review the steps to authorize IAM users on an Oracle Database instance.
- [To Exclusively Map a Local IAM User to an Oracle Database Global User](#)  
You can map a local IAM user exclusively to an Oracle Database global user.
- [Add Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Roles on Oracle Database](#)  
Optionally, create global roles to provide additional database roles and privileges to IAM users when multiple IAM users are mapped to the same shared global user.
- [Create Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Database Password for IAM Users](#)  
To add an IAM user and allow the IAM user to login to Oracle Database by supplying a username and password, you must create an IAM database password.

## Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies for IAM Users

Review the steps to write policy statements for an IAM group to enable IAM user access to Oracle Cloud Infrastructure resources, specifically Oracle Database instances using IAM database tokens.

A policy is a group of statements that specifies who can access particular resources, and how. Access can be granted for the entire tenancy, databases in a compartment, or individual databases. This means you write a policy statement that gives a specific group a specific type of access to a specific type of resource within a specific compartment.

**Note:** Defining a policy is required to use IAM tokens to access Oracle Database. A policy is not required when using IAM database password verifiers to access Oracle Database.

1. Create an IAM group for IAM users that will access the database. Review OCI IAM documentation for creating groups and adding IAM users to a group. For example, create the group *DBUsers*. For more information, see *Managing Groups*.
2. Write policy statements to enable access to Oracle Cloud Infrastructure resources.
  - a. In the Oracle Cloud Infrastructure console, click **Identity and Security**, and then click **Policies**.
  - b. To write a policy, click **Create Policy**, and then enter a **Name** and a **Description**.
  - c. Use the **Policy Builder** to create a policy. For example, to create a policy to allow users in IAM group *DBUsers* to access any Oracle Database in their tenancy:

```
allow group DBUsers to use database-connections in tenancy
```

Where, *database-connections* is the OCI resource name to connect to the database. *Use* is the minimum verb to allow access to the database. Both *use* and *manage* can be used.

For example to create a policy that limits members of *DBUsers* group to access Oracle Databases in the compartment *testing\_compartment* only:

```
allow group DBUsers to use database-connections in compartment
testing_compartment
```

For example, to create a policy that limits group access to a single database in a compartment:

```
allow group DBUsers to use database-connections in compartment
testing_compartment where target.database.id =
'ocid1.database.oc1.iad.aaaabbbbcccc'
```

- d. Click **Create**.  
For more information about policies, see *Managing Policies*.

Notes for creating policies for use with IAM users on Oracle Database:

- Policies can allow IAM users to access Oracle Database instances across the entire tenancy, in a compartment, or can limit access to a single Oracle Database instance.
- You must use dynamic groups for Instance Principals and Resource Principals. You can create Dynamic Groups and reference dynamic groups in the policies you create to access

Oracle Cloud Infrastructure. See *Accessing Cloud Resources by Configuring Policies and Roles and Managing Dynamic Groups* for details.

### Related Topics

- [Managing Groups](#)
- [Accessing Cloud Resources by Configuring Policies and Roles](#)
- [Managing Dynamic Groups](#)

## Authorize Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users on Oracle Database

Review the steps to authorize IAM users on an Oracle Database instance.

To authorize IAM users to allow access to Oracle Database, map database global users to IAM groups or directly to IAM users with `CREATE USER` or `ALTER USER` statements with `IDENTIFIED GLOBALLY AS` clause.

The authorization of IAM users to an Oracle Database instance works by mapping IAM global users (schemas) to IAM users (exclusive mapping) or IAM groups (shared schema mapping).

To authorize IAM users on a database instance:

1. Log in as a user with DBA privileges to the database that is enabled to use IAM. A user with the DBA role will need the required `CREATE USER` and `ALTER USER` system privileges for these steps.
2. Create a mapping between the Oracle Database user (schema) with `CREATE USER` or `ALTER USER` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the IAM group name. Use the following syntax to map a global user to an IAM group:

```
CREATE USER global_user IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=IAM_GROUP_NAME';
```

For example, to map an IAM group named `db_sales_group` to a shared database global user named `sales_group`:

```
CREATE USER sales_group IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=db_sales_group';
```

This creates a shared global user mapping. The mapping, with the global user `sales_group` is effective for all users in the IAM group. Thus, anyone in the `db_sales_group` can log in to the database using their IAM credentials through the shared mapping of the `sales_group` global user.

If you want to create additional global user mappings for other IAM groups or users, follow these steps for each IAM group or user.

### Note:

Database users that are not `IDENTIFIED GLOBALLY` can continue to login as before, even when the Oracle Database is enabled for IAM authentication.

## To Exclusively Map a Local IAM User to an Oracle Database Global User

You can map a local IAM user exclusively to an Oracle Database global user.

1. Log in as a user with DBA privileges to the database that is enabled to use IAM. A user with the DBA role has will need the required `CREATE USER` and `ALTER USER` system privileges that you need for these steps.
2. Create a mapping between the Oracle Database user (schema) with `CREATE USER` or `ALTER USER` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the IAM local IAM user name. For example, to create a new database global user named `peter_fitch` and map this user to an existing local IAM user named `peterfitch`:

```
CREATE USER peter_fitch IDENTIFIED GLOBALLY AS
'IAM_PRINCIPAL_NAME=peterfitch'
```

You can use either instance principal or resource principal to retrieve database tokens to establish a connection from your application to an Oracle Database instance.

If you are using an instance principal or resource principal, you must map a dynamic group. Thus, you cannot exclusively map instance and resource principals. You only can map them through a shared mapping and putting the instance or resource instance in an IAM dynamic group

## Add Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Roles on Oracle Database

Optionally, create global roles to provide additional database roles and privileges to IAM users when multiple IAM users are mapped to the same shared global user.

Creating global roles is optional, but useful when assigning users to a shared schema.

Use a global role to optionally differentiate users who use the same shared schema. For example, a set of users can all have the same shared schema and the shared schema could have the `CREATE SESSION` privilege. Then global roles can be used to provide differentiated privileges and roles assigned to different groups of users who all use the same shared schema.

Granting additional roles to IAM users in Oracle Database works by mapping Oracle Database global roles to IAM groups.

1. Log in as a user with DBA privileges to the database that is enabled to use IAM. A user with the DBA privileges `CREATE ROLE` and `ALTER ROLE` system privileges is needed for these steps.
2. Set database authorization for Oracle Database roles with `CREATE ROLE` or `ALTER ROLE` statements and include the `IDENTIFIED GLOBALLY AS` clause, specifying the IAM group name. Use the following syntax to map a global role to an IAM group:

```
CREATE ROLE global_role IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=IAM_GROUP_of_WHICH_the_IAM_USER_IS_a_MEMBER';
```

For example, to map an IAM group named `ExporterGroup` to a shared database global role named `export_role`:

```
CREATE ROLE export_role IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=ExporterGroup';
```

3. Use the `GRANT` statements to grant the required privileges or other roles to the global role.

```
GRANT CREATE SESSION TO export_role;
GRANT DWROLE TO export_role;
```

4. If you want an existing database role to be associated with an IAM group, then use the `ALTER ROLE` statement to alter the existing database role to map the role to an IAM group. Use the following syntax to alter an existing database role to map it to an IAM group:

```
ALTER ROLE existing_database_role IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=IAM_Group_Name';
```

Follow these steps for each IAM group to add additional global role mappings for other IAM groups.

## Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password for IAM Users

To add an IAM user and allow the IAM user to login to Oracle Database by supplying a username and password, you must create an IAM database password.

For more information, see *Working with IAM Database Passwords*.

### Related Topics

- [Working with IAM Database Passwords](#)

## Configuring Client Connection

Configure various clients to use IAM authentication.

- [Configure a Client Connection for SQL\\*Plus that Uses an IAM Database Password Verifier](#)  
You can configure SQL\*Plus to use an IAM database password verifier.
- [Configure Client Connection for SQL\\*Plus that Uses an IAM Token](#)  
You can configure a client connection for SQL\*Plus that uses an IAM token.
- [Client Connections That Use a Token Requested by an IAM User Name and Database Password](#)  
You can create a client connection that uses a token requested by an IAM user name and database password.
- [Use Instance Principal to Access Database with IAM Authentication](#)  
After the ADMIN user enables OCI IAM on the database, an application can access the database through an OCI IAM database token using an instance principal.
- [Configure Proxy Authentication](#)  
Proxy authentication allows an IAM user to proxy to a database schema for tasks such as application maintenance.

- **Use Database Link with IAM Authenticated Users**  
You can use a database link to connect from one database instance to another as an OCI IAM user.

## Configure a Client Connection for SQL\*Plus that Uses an IAM Database Password Verifier

You can configure SQL\*Plus to use an IAM database password verifier.

As the IAM user, log in to the database by using the following syntax:

```
CONNECT user_name@db_connect_string
Enter password: password
```

In this specification, `user_name` is the IAM user name. There is a limit of 128 bytes for the combined `domain_name/user_name`.

The following example shows how IAM user `peter_fitck` can log in to a database instance.

```
sqlplus /nolog
connect peter_fitck@db_connect_string
Enter password: password
```

Some special characters will require double quotation marks around `user_name` and `.`. For example:

```
"peter_fitck@example.com"@db_connect_string

"IAM database password"
```

## Configure Client Connection for SQL\*Plus that Uses an IAM Token

You can configure a client connection for SQL\*Plus that uses an IAM token.

1. Ensure you have an IAM user account.
2. Check with an IAM administrator and the database administrator to ensure you have a policy allowing you to access the database in the compartment or your tenancy and that you are mapped to a global schema in the database.
3. If your application or tool does not support direct IAM integration, then download, install, and configure the OCI CLI. (See [OCI Command Line Interface Quickstart](#).) Set up an API key as part of the OCI CLI configuration and select default values.
  - a. Set up the API key access for the IAM user.
  - b. Retrieve the `db-token`. For example:
    - Retrieve a `db-token` with an `API-key` using the OCI CLI:

```
oci iam db-token get
```

- Retrieve `db-token` with a security (or session) token:

```
oci iam db-token get --auth security_token
```

- Retrieve `db-token` with a delegation token: When you log in to the cloud shell, the delegation token is automatically generated and placed in the `/etc` directory. To get this token, execute the following command in the OCI CLI:

```
oci iam db-token get
```

- Using an instance principal to retrieve a `db-token` using OCI CLI:

```
oci iam db-token get --auth instance_principal
```

If the security token has expired, a window will appear so the user can log in to OCI again. This generates the security token for the user. OCI CLI will use this refreshed token to get the `db-token`.

See [Required Keys and OCIDs](#) for more information.

4. Ensure that you are using the latest release updates for the Oracle Database client releases 19c.  
This configuration only works with the Oracle Database client release 19c.
5. Follow the existing process to download the wallet from the database and then follow the directions for configuring it for use with SQL\*Plus.
  - a. Confirm that DN matching is enabled by looking for `SSL_SERVER_DN_MATCH=ON` in `sqlnet.ora`.
  - b. Configure the database client to use the IAM token by adding `TOKEN_AUTH=OCI_TOKEN` to the `sqlnet.ora` file. Because you will be using the default locations for the database token file, you do not need to include the token location.

The `TOKEN_AUTH` and `TOKEN_LOCATION` values in the `tnsnames.ora` connect strings take precedence over the `sqlnet.ora` settings for that connection. For example, for the connect string, assuming that the token is in the default location (`~/oci/db-token` for Linux):

```
(description=
  (retry_count=20) (retry_delay=3)
  (address=(protocol=tcps) (port=1522)
  (host=example.us-phoenix-1.oraclecloud.com) )

  (connect_data=(service_name=aaabbbccc_exampledb_high.example.oraclecloud.com))
  (security=(ssl_server_cert_dn="CN=example.uscom-east-1.oraclecloud.com,
    OU=Oracle BMCS US, O=Example Corporation,
    L=Redwood City, ST=California, C=US")
  (TOKEN_AUTH=OCI_TOKEN)))
```

After the connect string is updated with the `TOKEN_AUTH` parameter, the IAM user can log in to the database instance by running the following command to start SQL\*Plus. You can include the connect descriptor itself or use the name of the descriptor from the `tnsnames.ora` file.

```
connect /@exampledb_high
```

Or:

```
connect /@(description=
  (retry_count=20) (retry_delay=3)
  (address=(protocol=tcps) (port=1522)
```

```
(host=example.us-phoenix-1.oraclecloud.com)

(connect_data=(service_name=aaabbbccc_exampledb_high.example.oraclecloud.com)
  (security=(ssl_server_cert_dn="CN=example.uscom-east-1.oraclecloud.com,
    OU=Oracle BMCS US, O=Example Corporation,
    L=Redwood City, ST=California, C=US")
  (TOKEN_AUTH=OCI_TOKEN)))
```

The database client is already configured to get a `db-token` because `TOKEN_AUTH` has already been set, either through the `sqlnet.ora` file or in a connect string. The database client gets the `db-token` and signs it using the private key and then sends the token to the database. If an IAM user name and IAM database password are specified instead of slash /, then the database client will connect using the password instead of using the `db-token`.

## Client Connections That Use a Token Requested by an IAM User Name and Database Password

You can create a client connection that uses a token requested by an IAM user name and database password.

- IAM users can connect to the Oracle DBaaS instance by using an IAM token that was retrieved using an IAM user name and IAM database password.  
For more information, see *About Client Connections That Use a Token Requested by an IAM User Name and Database Password*
- To set these parameters, you modify either the `sqlnet.ora` file or the `tnsnames.ora` file.  
For more information, see *Parameters to Set for Client Connections That Use a Token Requested by an IAM User Name and Database Password*
- You can configure the database client to retrieve the IAM database token using the provided IAM user name and IAM database password.  
For more information, see *Configuring the Database Client to Retrieve a Token Using an IAM User Name and Database Password*
- You can enable an IAM user name and a secure external password store (SEPS) to request the IAM database token.  
For more information, see *Configuring a Secure External Password Store Wallet to Retrieve an IAM Token*

### Related Topics

- [Client Connections That Use a Token Requested by an IAM User Name and Database Password](#)
- [About Client Connections That Use a Token Requested by an IAM User Name and Database Password](#)
- [Parameters to Set for Client Connections That Use a Token Requested by an IAM User Name and Database Password](#)
- [Configuring the Database Client to Retrieve a Token Using an IAM User Name and Database Password](#)
- [Configuring a Secure External Password Store Wallet to Retrieve an IAM Token](#)



## Use Instance Principal to Access Database with IAM Authentication

After the ADMIN user enables OCI IAM on the database, an application can access the database through an OCI IAM database token using an instance principal.

For more information, see *Accessing the Oracle Cloud Infrastructure API Using Instance Principals*.

For more information, see *Accessing the Database Using an Instance Principal or a Resource Principal*.

### Related Topics

- [Accessing the Oracle Cloud Infrastructure API Using Instance Principals](#)
- [Accessing the Database Using an Instance Principal or a Resource Principal](#)

## Configure Proxy Authentication

Proxy authentication allows an IAM user to proxy to a database schema for tasks such as application maintenance.

Proxy authentication is typically used to authenticate the real user and then authorize them to use a database schema with the schema privileges and roles in order to manage an application. Alternatives such as sharing the application schema password are considered insecure and unable to audit which actual user performed an action.

A use case can be in an environment in which a named IAM user who is an application database administrator can authenticate by using their credentials and then proxy to a database schema user (for example, `hrapp`). This authentication enables the IAM administrator to use the `hrapp` privileges and roles as user `hrapp` in order to perform application maintenance, yet still use their IAM credentials for authentication. An application database administrator can sign in to the database and then proxy to an application schema to manage this schema.

You can configure proxy authentication for both the password authentication and token authentication methods.

### Configuring Proxy Authentication for the IAM User

To configure proxy authentication for an IAM user, the IAM user must already have a mapping to a global schema (exclusive or shared mapping). A separate database schema for the IAM user to proxy to must also be available.

After you ensure that you have this type of user, alter the database user to allow the IAM user to proxy to it.

1. Log in to the database instance as a user who has the `ALTER USER` system privileges.
2. Grant permission for the IAM user to proxy to the local database user account. An IAM user cannot be referenced in the command so the proxy must be created between the database global user (mapped to the IAM user) and the target database user. In the following example, `hrapp` is the database schema to proxy to, and `peterfitch_schema` is the database global user exclusively mapped to user `peterfitch`.

```
ALTER USER hrapp GRANT CONNECT THROUGH peterfitch_schema;
```

At this stage, the IAM user can log in to the database instance using the proxy. For example:

- To connect using a password verifier:

```
CONNECT peterfitch[hrapp]@connect_string
Enter password: password
```

- To connect using a token:

```
CONNECT [hrapp]/@connect_string
```

### Validating the IAM User Proxy Authentication

You can validate the IAM user proxy configuration for both password and token authentication methods.

1. Connect as the IAM user and proxied to the database user. Run the `SHOW USER` and `SELECT SYS_CONTEXT` commands.

For example, suppose you want to check the proxy authentication of the IAM user *peterfitch* when they proxy to database user *hrapp*. You will need to connect to the database using the different types of authentication methods shown here, but the output of the commands that you execute will be the same for all types.

- For password authentication:

```
CONNECT peterfitch[hrapp]/password\!@connect_string SHOW USER;
```

```
--The output should be USER is "HRAPP"
SELECT SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD') FROM DUAL;
--The output should be "PASSWORD_GLOBAL"
SELECT SYS_CONTEXT('USERENV','PROXY_USER') FROM DUAL;
--The output should be "PETERFITCH_SCHEMA"
SELECT SYS_CONTEXT('USERENV','CURRENT_USER') FROM DUAL;
--The output should be "HRAPP"
```

- For token authentication:

```
CONNECT [hrapp]/@connect_string
SHOW USER;
```

```
--The output should be USER is "HRAPP "
SELECT SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD') FROM DUAL;
--The output should be "TOKEN_GLOBAL"
SELECT SYS_CONTEXT('USERENV','PROXY_USER') FROM DUAL;
--The output should be "PETERFITCH_SCHEMA"
SELECT SYS_CONTEXT('USERENV','CURRENT_USER') FROM DUAL;
--The output should be "HRAPP"
```

## Use Database Link with IAM Authenticated Users

You can use a database link to connect from one database instance to another as an OCI IAM user.

You can use either connected user or fixed user database link to connect to a database as an OCI IAM user.

**Note:**

Current user database link is not supported for connecting to a database in Exadata Cloud Infrastructure as an OCI IAM user.

- **Connected User Database Link:** For a connected user database link, an IAM user must be mapped to a schema in both the source and target databases connected by a database link. You can use a database password verifier or an IAM database token to use a connected user database link.
- **Fixed User Database Link:** A fixed user database link can be created using a database user or an IAM user. When using an IAM user as a fixed user database link, the IAM user must have a schema mapping in the target database. The IAM user for a database link can be configured with a password verifier only.

# 6

## Reference Guides for Oracle Exadata Database Service on Exascale Infrastructure

- [Using the dbaascli Utility on Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn to use the `dbaascli` utility on Oracle Exadata Database Service on Exascale Infrastructure.
- [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.
- [Monitor Metrics for VM Cluster Resources](#)
- [Metrics for Oracle Exadata Database Service on Exascale Infrastructure in the Monitoring Service](#)  
learn about 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 Exascale Infrastructure Events](#)  
Oracle Exadata Database Service on Exascale Infrastructure resources emit events, which are structured messages that indicate changes in resources.
- [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.
- [Policy Details for Oracle Exadata Database Service on Exascale Infrastructure](#)  
This topic covers details for writing policies to control access to Oracle Exadata Database Service on Exascale Infrastructure resources.
- [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 Oracle Cloud (Bare Metal, Virtual Machine VM, and Exadata Cloud Infrastructure).
- [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.
- [Troubleshooting Oracle Exadata Database Service on Exascale Infrastructure Systems](#)  
These topics cover some common issues you might run into and how to address them.

### Using the dbaascli Utility on Oracle Exadata Database Service on Exascale Infrastructure

Learn to use the `dbaascli` utility on Oracle Exadata Database Service on Exascale Infrastructure.

- [About Using the dbaascli Utility on Oracle Exadata Database Service on Exascale Infrastructure](#)  
You can use the `dbaascli` utility to perform various database lifecycle and administration operations on Oracle Exadata Database Service on Exascale Infrastructure
- [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 Oracle Exadata Database Service on Exascale Infrastructure Software Images Using the Dbaascli Utility](#)  
You can list and download the Oracle database software images on an Oracle Exadata Database Service on Exascale Infrastructure instance, which can then be used for provisioning a database home.
- [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 Exascale Infrastructure, complete this procedure.
- [Creating a Duplicate Database](#)
- [dbaascli Command Reference](#)  
You use `dbaascli` to create databases and integrate them with the cloud automation framework.

## About Using the dbaascli Utility on Oracle Exadata Database Service on Exascale Infrastructure

You can use the `dbaascli` utility to perform various database lifecycle and administration operations on Oracle Exadata Database Service on Exascale Infrastructure

For example, with `dbaascli`, you can change the password of a database user, start a database, or manage 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 Oracle Exadata Database Service on Exascale Infrastructure virtual machine.

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 *dbasscli Command Reference* in the document for commands and command specific information.

## 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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:

```
dbascli 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:

```
dbascli database create --dbName database name --oracleHome Oracle Home Path
```

To create a non-CDB, run the following command:

```
dbascli 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 `dbascli 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure Database

1. Log onto the Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure Database

1. Log onto the Oracle Exadata Database Service on Exascale Infrastructure virtual machine as `opc`.
2. Run the following command:

```
sudo dbaascli tde changepassword --dbname database_name
```

## Managing Oracle Exadata Database Service on Exascale Infrastructure Software Images Using the Dbaascli Utility

You can list and download the Oracle database software images on an Oracle Exadata Database Service on Exascale Infrastructure instance, which can then be used for provisioning a database home.

### Note:

You can create custom database software images for your Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale Infrastructure system by using a Secure Shell (SSH) connection.

## 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 `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--BeVuXaipSEEMISrSCFrVEolG/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 Exascale Infrastructure, complete this procedure.

Cloud-specific tooling is used on the Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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](#)

## 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 prerequisites 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

```
dbascli database duplicate --dbName <database name> --oracleHome <Oracle
Home Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod FILE --tdeConfigMethod FILE
```



#### Note:

If source database is using OKV for TDE keystore management, current duplicate database operation does not support this configuration.

## 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

```
dbascli 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

```
dbascli 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.

## dbaascli Command Reference

You 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 Oracle Exadata Database Service on Exascale Infrastructure compute node.

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 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 Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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.
- [dbaascli database modifyParameters](#)  
To modify or reset initialization parameters for an Oracle Database, use the `dbaascli database modifyParameters` command.
- [dbaascli database move](#)  
To move the database from one home to another, use the `dbaascli database move` command.
- [dbaascli database recover](#)  
To recover a database, use the `dbaascli database recover` command.
- [dbaascli database runDatapatch](#)  
To patch an Oracle Database, use the `dbaascli database runDatapatch` command.



- [dbascli database createTemplate](#)  
To create database templates (DBCA templates) that can subsequently be used to create databases, use the `dbascli database createTemplate` command.
- [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 dbhome create](#)  
To create an Oracle Database home of desired version, use the `dbascli dbhome create` command.
- [dbascli dbHome delete](#)  
To delete a given Oracle Database home, use the `dbascli dbHome delete` command.
- [dbascli dbhome getDatabases](#)  
To view information about all Oracle Databases running from a given database Oracle home, use the `dbascli dbHome getDatabases` command. Specify either the Oracle home location or Oracle home name.
- [dbascli dbHome getDetails](#)  
To view information about a specific Oracle home, use the `dbascli dbHome getDetails` command. Specify either the Oracle home location or Oracle home name.
- [dbascli dbHome patch](#)  
To patch Oracle home from one patch level to another, use the `dbascli dbHome patch` command.
- [dbascli dbimage purge](#)  
The `dbimage purge` command removes the specified software image from your Oracle Exadata Database Service on Exascale Infrastructure environment.
- [dbascli diag collect](#)  
To collect diagnostics, use the `dbascli diag collect` command.
- [dbascli diag healthCheck](#)  
To run diagnostic health checks, use the `dbascli diag healthCheck` command.
- [dbascli grid configureTCPS](#)  
To configure TCPS for the existing cluster, use the `dbascli grid configureTCPS` command.
- [dbascli grid patch](#)  
To patch Oracle Grid Infrastructure to the specified minor version, use the `dbascli 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 createSnapshot](#)  
To create a snapshot of a given pluggable database (PDB), use the `dbaascli pdb createSnapshot` command.
- [dbaascli pdb configureSnapshot](#)  
To configure automatic snapshots for a given pluggable database (PDB), use the `dbaascli pdb configureSnapshot` command.
- [dbaascli pdb delete](#)  
To delete a pluggable database (PDB) run the `dbaascli pdb delete` command.
- [dbaascli pdb deleteSnapshot](#)  
To delete a snapshot of a given pluggable database (PDB), use the `dbaascli pdb deleteSnapshot` command.
- [dbaascli pdb getDetails](#)  
To view details of a pluggable database (PDB), use the `dbaascli pdb getDetails` command.
- [dbaascli pdb getSnapshot](#)  
To obtain details of a given pluggable database (PDB) snapshot, use the `dbaascli pdb getSnapshot` command.
- [dbaascli pdb list](#)  
To view the list of pluggable databases (PDB) in a container database, use the `dbaascli pdb list` command.

- [dbaascli pdb listSnapshots](#)  
To list the snapshots of a given pluggable database (PDB), use the `dbaascli pdb listSnapshots` 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 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](#)  
To rotate the master key for database encryption, use the `dbaascli tde rotateMasterKey` command.
- [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 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 Oracle Exadata Database Service on Exascale 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

### Frequently Asked Questions

**Q: What is the dbaascli admin updateStack command used for?**

A: The `dbaascli admin updateStack` command is used to install or update a `dbaastools` RPM on Exadata Cloud Infrastructure.

**Q: What are the prerequisites for using the `dbaascli admin updateStack` command?**

A: You must run the command as the root user and connect to an Exadata Cloud Infrastructure virtual machine.

**Q: What does the `--resume` option do?**

A: The `--resume` option resumes the previous execution of the `updateStack` command if it was interrupted or incomplete.

**Q: What is the purpose of the `--prechecksOnly` option?**

A: The `--prechecksOnly` option runs only the prechecks for the operation without actually performing the installation or update.

**Q: How is the `--nodes` parameter used?**

A: The `--nodes` parameter specifies a comma-delimited list of nodes on which the RPM should be installed. If not provided, the RPM will be installed on all cluster nodes.

**Q: What should I do if I encounter issues with the `dbaascli admin updateStack` command?**

A: Ensure you are running the command as the root user and that you are connected to an Exadata Cloud Infrastructure virtual machine. Check if there are any specific error messages and consult the command documentation or Oracle support if needed.

**Q: How do I connect to an Exadata Cloud Infrastructure virtual machine to use the `dbaascli admin updateStack` command?**

A: You need to use SSH to connect to the virtual machine. Refer to the section on "Connecting to a Virtual Machine with SSH" in the documentation for detailed instructions.

### Example Use Cases

**Example 1: Installing or updating the `dbaastools` RPM on all nodes**

```
dbaascli admin updateStack
```

Installs or updates the `dbaastools` RPM on all nodes of the Exadata Cloud@Customer environment.

**Example 2: Running prechecks only before installing or updating the RPM**

```
dbaascli admin updateStack --prechecksOnly
```

Runs only the prechecks for the `dbaastools` RPM update, without actually performing the installation. It ensures that all prerequisites are satisfied before proceeding with the update.

**Example 3: Resuming a previously interrupted `updateStack` operation**

```
dbaascli admin updateStack --resume
```

Resumes a previous `dbaastools` RPM update operation that was interrupted or did not complete successfully.

**Example 4: Installing or updating `dbaastools` on specific nodes**

```
dbaascli admin updateStack --nodes node1,node2
```

Installs or updates the `dbaastools` RPM on the specified nodes `node1` and `node2` only, without affecting other nodes in the cluster.

#### Example 5: Resuming the updateStack process on specific nodes

```
dbaascli admin updateStack --resume --nodes node3,node4
```

Resumes the update process for `dbaastools` on the specific nodes `node3` and `node4` only, if the previous execution was interrupted.

#### Related Topics

- [Connecting to a Virtual Machine with SSH](#)  
You can connect to the virtual machines in an Oracle Exadata Database Service on Exascale 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

#### Frequently Asked Questions

##### Q: What is the purpose of the dbaascli cswlib deleteLocal command?

A: The `dbaascli cswlib deleteLocal` command is used to delete a local Oracle home image from the system.

##### Q: What are the prerequisites for running the dbaascli cswlib deleteLocal command?

A: The command must be run as the `root` user to ensure the necessary permissions are available to delete the local image.

##### Q: How do I specify which local image to delete?

A: Use the `--imageTag` option to specify the Oracle home image tag that you want to delete.

##### Q: What does the --imageTag option represent in the command?

A: The `--imageTag` option represents the identifier or tag associated with the Oracle home image that you want to delete.

##### Q: Can I delete multiple local images at once using this command?

A: No, the `dbaascli cswlib deleteLocal` command allows you to delete only one local image at a time, specified by its image tag.

##### Q: What happens if I run the dbaascli cswlib deleteLocal command without specifying the --imageTag?

A: The command will fail because the `--imageTag` is required to identify which local image should be deleted.

**Q: Is it possible to recover a local image after it has been deleted using this command?**

A: No, once the local image is deleted using the `dbaascli cswlib deleteLocal` command, it cannot be recovered. Make sure to verify the image tag before proceeding.

**Q: When would I need to use the dbaascli cswlib deleteLocal command?**

A: You would use this command when you need to remove an unused or outdated Oracle home image from the local system to free up space or clean up your environment.

**Example 6-3 dbaascli cswlib deletelocal**

```
dbaascli cswlib deletelocal --imagetag 19.15.0.0.0
DBAAS CLI version MAIN
Executing command cswlib deletelocal --imagetag 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 Oracle Exadata Database Service on Exascale Infrastructure system by using a Secure Shell (SSH) connection.

## dbaascli cswlib download

To download available software images and make them available in your Oracle Exadata Database Service on Exascale 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 Oracle Exadata Database Service on Exascale 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`

**Frequently Asked Questions**

**Q: What is the purpose of the dbaascli cswlib download command?**

A: The `dbaascli cswlib download` command is used to download available software images and make them available in your Exadata Cloud Infrastructure.

**Q: What are the prerequisites for running the dbaascli cswlib download command?**

A: You must run the command as the `root` user. Additionally, you need to be connected to an Exadata Cloud Infrastructure virtual machine.

**Q: How do I connect to the virtual machine required for this command?**

A: You need to use SSH to connect to the Exadata Cloud Infrastructure virtual machine. Detailed instructions can be found in the documentation under "Connecting to a Virtual Machine with SSH."

**Q: What does the --version option specify in the command?**

A: The `--version` option specifies the Oracle home image version that you want to download.

**Q: How do I use the --imageTag option in the dbaascli cswlib download command?**

A: The `--imageTag` option is used to specify the image tag of the software image you want to download.

**Q: What is the purpose of the --product option in the command?**

A: The `--product` option specifies the type of image you want to download. The valid values are `database` or `grid`.

**Q: Can I download both database and grid images simultaneously?**

A: No, you must specify either `database` or `grid` using the `--product` option, so each download operation is specific to one type of image.

**Q: What happens if I do not specify a version or image tag?**

A: The command will likely fail or prompt you for the required information since the `--version` or `--imageTag` options are necessary to identify the specific software image to download.

**Q: Is it necessary to specify both --version and --imageTag together?**

A: No, you typically specify either `--version` or `--imageTag` depending on how you want to identify the image to download, but not both at the same time.

**Q: When would I use the dbaascli cswlib download command?**

A: You would use this command when you need to download Oracle home software images for `database` or `grid` environments in your Exadata Cloud Infrastructure setup.

**Example 6-4 dbaascli cswlib download --product --imageTag**

```
dbaascli cswlib download --product database --imageTag 19.14.0.0.0
```

**Example 6-5 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 Oracle Exadata Database Service on Exascale 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`.

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli cswlib listLocal` command?

A: The `dbaascli cswlib listLocal` command is used to view the list of locally available Database and Grid Infrastructure images on your system.

#### Q: What are the prerequisites for running the `dbaascli cswlib listLocal` command?

A: The command must be run as the `root` user to have the necessary permissions to access and list the available images.

#### Q: How do I specify which type of images to list using this command?

A: Use the `--product` option to specify the type of Oracle home images you want to list. The valid values are `database` or `grid`.

#### Q: What does the `--product` option represent in the `dbaascli cswlib listLocal` command?

A: The `--product` option identifies the Oracle home product type, allowing you to filter the list of available images to either `database` or `grid` types.

#### Q: Can I list both database and grid images simultaneously?

A: No, the `--product` option allows you to list either `database` or `grid` images at a time. You need to run the command twice with different `--product` values to see both lists.

#### Q: What happens if I do not specify the `--product` option in the command?

A: If the `--product` option is not specified, the command might list all locally available images or it could require you to specify the product type. The behavior may depend on your environment setup.

#### Q: When should I use the `dbaascli cswlib listLocal` command?

A: You should use this command when you want to check which Database or Grid Infrastructure images are currently available locally on your system.

#### Q: How can I differentiate between database and grid images in the list?

A: The `--product` option lets you filter the list, so by specifying `database` or `grid`, you will only see the images relevant to that product type, making it easier to differentiate.

#### Q: Is there any risk associated with running the `dbaascli cswlib listLocal` command?

A: No, this command is non-destructive and only displays information about locally available images. It does not modify or delete any files.

**Q: Does this command display remote or cloud-stored images?**

A: No, the `dbaascli cswlib listLocal` command only displays images that are available locally on your system, not those stored remotely or in the cloud.

**Example 6-6 dbaascli cswlib listlocal**

```
dbaascli cswlib listlocal
DBAAS CLI version MAIN
Executing command cswlib listlocal
Job id: bc4f047c-0a34-4d4d-a1ea-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 Oracle Exadata Database Service on Exascale 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`.

**Frequently Asked Questions**

**Q: What is the purpose of the `dbaascli cswlib showImages` command?**

A: The `dbaascli cswlib showImages` command is used to view the list of available Database and Grid Infrastructure images that can be downloaded or managed within your Oracle Exadata Database Service environment.

**Q: What are the prerequisites for running the dbaascli cswlib showImages command?**

A: The command must be run as the `root` user to ensure you have the necessary permissions to view the available images.

**Q: How do I filter the images listed by this command?**

A: You can filter the images by specifying the `--product` option with either `database` or `grid` to list only the images related to that product type.

**Q: What does the --product option represent in the dbaascli cswlib showImages command?**

A: The `--product` option identifies the Oracle home product type, allowing you to filter the list of images to either `database` or `grid`.

**Q: Can I view both database and grid images in a single command execution?**

A: No, you need to run the command twice with different `--product` values (`database` and `grid`) to view both types of images.

**Q: What happens if I do not specify the --product option in the command?**

A: If the `--product` option is not specified, the command may list all available images or it may prompt you to specify the product type, depending on your environment configuration.

**Q: When should I use the dbaascli cswlib showImages command?**

A: Use this command when you want to view the list of Database or Grid Infrastructure images that are available for download or deployment in your Oracle Exadata Database Service environment.

**Q: Is there any difference between dbaascli cswlib showImages and dbaascli cswlib listLocal commands?**

A: Yes, `dbaascli cswlib showImages` lists all available images that you can download or manage, while `dbaascli cswlib listLocal` lists only the images that are already downloaded and available locally on your system.

**Q: Can this command be used to view images stored in the cloud?**

A: Yes, this command can show images that are available for download from Oracle's repositories, not just those that are stored locally.

**Q: What type of images can I expect to see with this command?**

A: You can expect to see images related to Oracle Database and Grid Infrastructure, which are essential components for managing and running Oracle databases on Exadata platforms.

**Example 6-7 dbaascli cswlib showImages**

```
dbaascli cswlib showImages
```

**Related Topics**

- [Connecting to a Virtual Machine with SSH](#)  
You can connect to the virtual machines in an Oracle Exadata Database Service on Exascale 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

### Frequently Asked Questions

#### Q: What is the purpose of the dbaascli database addInstance command?

A: The `dbaascli database addInstance` command is used to add a new database instance to a specified node in an Oracle Exadata Database Service environment.

#### Q: What are the prerequisites for running the dbaascli database addInstance command?

A: The command must be run as the `root` user to have the necessary permissions to add a database instance.

#### Q: What does the --dbname option represent in this command?

A: The `--dbname` option specifies the name of the Oracle Database for which you want to add a new instance.

#### Q: What is the --node option used for in the dbaascli database addInstance command?

A: The `--node` option specifies the name of the node where the new database instance will be added.

#### Q: What is the purpose of the --newNodeSID option in this command?

A: The `--newNodeSID` option allows you to specify the SID (System Identifier) for the new database instance that will be created on the specified node.

#### Q: Is it mandatory to specify the --newNodeSID option when adding a new instance?

A: The `--newNodeSID` option is optional. If not provided, Oracle will automatically generate an SID for the new database instance.

#### Q: When should I use the dbaascli database addInstance command?

A: Use this command when you want to scale your database by adding a new instance to an additional node in a multi-node Oracle Database setup.

#### Q: Can I add multiple database instances to different nodes using this command?

A: Yes, you can run the command multiple times to add database instances to different nodes by specifying the appropriate `--node` and `--dbname` values.

**Q: What happens if the node specified in the `--node` option is not available?**

A: The command will fail if the specified node is not available or reachable. Ensure that the node is properly configured and accessible before running the command.

**Q: Can this command be used in a Data Guard environment?**

A: Yes, you can use the `dbaascli database addInstance` command in a Data Guard environment to add instances, but it is recommended to follow the necessary Data Guard guidelines for such configurations.

**Q: Will this command cause database downtime?**

A: Adding an instance to a new node typically does not cause downtime for the existing database instances, but it's recommended to check your environment for any specific dependencies.

## 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>
  | --getBackupReport
  {
    --tag <value>
    | --latest
  }
  --json <value>
  | --configure
  {
    --configFile <value>
    | --enableRTRT
    | --disableRTRT
  }
  | --getConfig [--configFile <value>]
```

```

    | --validate [--untilTime <value>]
    | --showHistory [--all]
}

```

**Where:**

```

--dbname: Oracle Database name.
--list | --start | --delete | --status | --getBackupReport | --configure | --
getConfig
--list: Returns database backup information.
    [--json: Specify the file name for JSON output.]
--start: Begins database backup.
    [--level0 | --level1 | --archival]
    [--level0: Creates a Level-0 (full) backup. ]
    [--level1: Creates a Level-1 (incremental) backup. ]
    [--archival: Creates an Archival full backup. ]
    --tag: Specify backup tag.
--delete: Deletes Archival backup.
    --backupTag <value>
--status
    --uuid <value>
--getBackupReport: Returns backup report.
    --tag: Specify backup tag.
    --latest: Returns latest backup report (all types of database backup).
    --json: Specify the file name for JSON output.
--configure: Configures database for backup.
    --configFile | --enableRTRT | --disableRTRT
    --configFile: Specify database backup configuration file.
    --enableRTRT: Enables Real Time Redo Transport.
    --disableRTRT: Disables Real Time Redo Transport.
--getConfig: Returns database backup configuration.
    [--configFile: Specify the database backup configuration file.]
--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 the history of backup operations.
    [--all: Displays all backup operations.]

```

 **Note:**

enableRTRT and disableRTRT are applicable only for ZDLRA backup destination on Exadata Database Service on Cloud@Customer.

**Frequently Asked Questions****Q: What is the purpose of the dbascli database backup command?**

A: The dbascli database backup command is used to configure Oracle Database backup storage destinations, take backups, query backups, and delete existing backups.

**Q: What are the prerequisites for running the dbascli database backup command?**

A: The command must be run as the root user to have the necessary permissions for backup management.

**Q: How do I start a full backup of the Oracle Database using this command?**

A: To start a full backup (Level-0), use the following syntax:

```
dbaascli database backup --dbname <value> --start --level0
```

**Q: How do I perform an incremental backup using the dbaascli database backup command?**

A: To perform a Level-1 incremental backup, use this syntax:

```
dbaascli database backup --dbname <value> --start --level1
```

**Q: What is the difference between Level-0 and Level-1 backups?**

A: A Level-0 backup is a full backup of the database, while a Level-1 backup is an incremental backup that only captures changes made since the last Level-0 or Level-1 backup.

**Q: Can I take an archival backup using this command?**

A: Yes, you can create an archival backup by using the `--archival` option along with the `--start` command:

```
dbaascli database backup --dbname <value> --start --archival --tag <backup_tag>
```

**Q: How do I delete an existing archival backup?**

A: To delete an archival backup, use the following syntax:

```
dbaascli database backup --dbname <value> --delete --backupTag <tag_value>
```

**Q: How can I check the status of a specific backup using the command?**

A: You can check the status of a backup by using the `--status` option with the `--uuid` parameter, like this:

```
dbaascli database backup --dbname <value> --status --uuid <backup_uuid>
```

**Q: How do I list all backups for a database?**

A: To list all available backups for a specific database, use the `--list` option:

```
dbaascli database backup --dbname <value> --list
```

For JSON output, add the `--json` option:

```
dbaascli database backup --dbname <value> --list --json <file_name>
```

**Q: How can I retrieve a backup report?**

A: You can get a backup report using the `--getBackupReport` option, either for a specific tag or for the latest backup:

```
dbaascli database backup --dbname <value> --getBackupReport --tag <backup_tag> --json <file_name>
```

Or to retrieve the latest report:

```
dbaascli database backup --dbname <value> --getBackupReport --latest --json <file_name>
```

**Q: How do I configure database backup settings?**

A: Use the `--configure` option to specify the backup configuration file or to enable/disable Real-Time Redo Transport (RTRT):

```
dbaascli database backup --dbname <value> --configure --configFile <config_file>
```

To enable RTRT:

```
dbaascli database backup --dbname <value> --configure --enableRTRT
```

**Q: How do I check the current backup configuration for my database?**

A: To view the current database backup configuration, use the `--getConfig` option:

```
dbaascli database backup --dbname <value> --getConfig
```

**Q: What does the `--validate` option do in the dbaascli database backup command?**

A: The `--validate` option checks if backups are complete and free from corruption. You can specify a time range using the `--untilTime` option:

```
dbaascli database backup --dbname <value> --validate --untilTime "DD-MON-YYYY
HH24:MI:SS"
```

**Q: How do I view the history of all backup operations for a database?**

A: Use the `--showHistory` option to display the history of all backup operations:

```
dbaascli database backup --dbname <value> --showHistory
```

For complete history, including all operations:

```
dbaascli database backup --dbname <value> --showHistory --all
```

**Q: What are RTRT (Real-Time Redo Transport) options and when should I use them?**

A: The RTRT options (`--enableRTRT` and `--disableRTRT`) are used to control Real-Time Redo Transport, applicable only for ZDLRA (Zero Data Loss Recovery Appliance) backup destinations in Exadata Cloud@Customer environments. Enable RTRT to ensure real-time redo log shipping.

**Example 6-8 Examples**

- To change the archive log retention period follow the below steps:

```
dbaascli database backup --getConfig --dbname <dbname>
```

This will generate a backup config file `.cfg`.

Update `bkup_archlog_fra_retention` value in this config file.

Run the configure command:

```
dbaascli database backup --configure --dbname <dbname> --configfile
<config file generated above>
```

- To get backup configuration for a database `myTestDB`:

```
dbaascli database backup --dbName myTestDB --getConfig --configFile /tmp/
configfile_1.txt
```

- To set backup configuration for a database `myTestDB` by modifying the config file with configuration details:

```
dbaascli database backup --dbName myTestDB --configure --configFile /tmp/
configfile_1_modified.txt
```



- To take backup of the database *myTestDB*:

```
dbaascli database backup --dbName myTestDB --start
```

- To query the status of backup request submitted with `uuid` *58fdcae0bd1c11eb92bc020017075151*:

```
dbaascli database backup --dbName myTestDB --status --uuid  
58fdcae0bd1c11eb92bc020017075151
```

- To enable RTRT for the database *myTestDB*:

```
dbaascli database backup --dbName myTestDB --configure --enableRTRT
```

## dbaascli database bounce

To shut down and restart a specified Oracle Exadata Database Service on Exascale 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.

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli database bounce` command?

A: The `dbaascli database bounce` command is used to shut down and restart an Oracle Database in Exadata Cloud Infrastructure. It supports restarting the database in a rolling manner, ensuring minimal disruption.

#### Q: What are the prerequisites for running the `dbaascli database bounce` command?

A: The command must be run as the `oracle` user, which has the required privileges to shut down and restart the database.

#### Q: What does the `--dbname` option specify in this command?

A: The `--dbname` option specifies the name of the Oracle Database that you want to shut down and restart.

#### Q: What is the `--rolling` option used for in the `dbaascli database bounce` command?

A: The `--rolling` option specifies whether to bounce (restart) the database in a rolling manner. If set to true, the database instances are restarted one by one, ensuring minimal downtime. The default value is `false`, which restarts all instances at once.

**Q: What does "bouncing the database" mean?**

A: Bouncing the database refers to shutting it down and then restarting it. This operation can be used for maintenance, applying changes, or recovering from certain types of issues.

**Q: Does the dbascli database bounce command perform a graceful shutdown?**

A: Yes, the command performs a shutdown in "immediate" mode, which closes the database and rolls back uncommitted transactions without waiting for users to disconnect.

**Q: Will this command automatically open all PDBs in an Oracle 12c or later database?**

A: Yes, if the database is running Oracle Database 12c or later, the command will automatically open all Pluggable Databases (PDBs) after restarting the database.

**Q: Can the dbascli database bounce command be used in a multi-node or RAC (Real Application Clusters) environment?**

A: Yes, in a multi-node or RAC environment, you can use the `--rolling` option to restart the database instances one by one, minimizing downtime.

**Q: What happens if I do not specify the --rolling option?**

A: If the `--rolling` option is not specified, or if it's set to `false`, the command will shut down and restart all database instances at the same time, which may cause a brief downtime.

**Q: Is there a default value for the --rolling option in the dbascli database bounce command?**

A: Yes, the default value for the `--rolling` option is `false`, meaning the database will be bounced in a non-rolling fashion unless otherwise specified.

**Q: How do I restart a database in rolling mode?**

A: To restart the database in rolling mode, use the following syntax:

```
dbascli database bounce --dbname <value> --rolling true
```

**Q: Is it safe to run the dbascli database bounce command during active sessions?**

A: While the command uses an immediate shutdown, which rolls back uncommitted transactions, it is always recommended to ensure there are no critical or active sessions before bouncing the database.

**Q: Can this command be used for specific PDBs in a multitenant database?**

A: No, the `dbascli database bounce` command operates on the entire database. In Oracle 12c or later, it will bounce the Container Database (CDB) and open all PDBs, but it does not allow for bouncing individual PDBs.

**Q: What should I do if the database does not come back online after bouncing it?**

A: If the database fails to restart, check the logs for any errors during the shutdown or startup process. Investigating the Oracle alert logs may provide insight into what caused the issue.

**Example 6-9 dbascli database bounce**

```
dbascli database bounce --dbname dbname
```

## dbascli database changepassword

To change the password of a specified Oracle Database user, use the `dbascli 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

```
dbascli 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

### Frequently Asked Questions

#### Q: What does the `dbascli database changePassword` command do?

A: The `dbascli database changePassword` command is used to change the password of a specified Oracle Database user. You will be prompted to enter the username and then the new password.

#### Q: What are the prerequisites for using the `dbascli database changePassword` command?

A: You must run the command as either the `root` or `oracle` user to change the password for a database user.

#### Q: How do I specify the database when using this command?

A: Use the `--dbname` option to specify the name of the Oracle Database you want to act on. For example:

```
dbascli database changePassword --dbname <db_name>
```

#### Q: How do I specify the user whose password I want to change?

A: Use the `--user` option to specify the username whose password needs to be changed. For example:

```
dbascli database changePassword --user <username>
```

**Q: What is the purpose of the `--prepareStandbyBlob` option in the `dbascli database changePassword` command?**

A: The `--prepareStandbyBlob` option is used in Data Guard environments to generate a blob file that contains the artifacts required for the password change on the standby database. This ensures password synchronization across the Data Guard environment.

**Q: What does the `--blobLocation` option specify?**

A: The `--blobLocation` option allows you to specify a custom path where the standby blob file should be generated. If not provided, the file will be saved in the default location.

**Q: How do I use the blob generated from the primary database to change the password on the standby?**

A: You can use the `--standbyBlobFromPrimary` option to specify the blob file prepared from the primary database to apply the password change to the standby database. For example:

```
dbascli database changePassword --standbyBlobFromPrimary <blob_file_path>
```

**Q: What is the `--resume` option used for in this command?**

A: The `--resume` option is used to resume a previously interrupted password change operation. You can specify the session ID if needed using the `--sessionID` option.

**Q: Can I resume a specific session with the `dbascli database changePassword` command?**

A: Yes, you can use the `--resume` option along with `--sessionID` to resume a specific password change session by specifying the session ID.

**Q: Is the `dbascli database changePassword` command applicable in a Data Guard environment?**

A: Yes, it is. The `--prepareStandbyBlob` option can be used to ensure that password changes are propagated to the standby database in a Data Guard setup.

**Q: What happens if I don't provide a `--blobLocation` when using `--prepareStandbyBlob`?**

A: If no `--blobLocation` is provided, the blob file containing the password change artifacts will be saved to the default location.

**Q: How do I check the status of a resumed session using `dbascli database changePassword`?**

A: You can specify the session ID using the `--sessionID` option to resume a specific session. The system will pick up where it left off in changing the password.

**Q: Can this command be used for both regular databases and those in a Data Guard environment?**

A: Yes, the command works for both regular Oracle Databases and databases in a Data Guard environment. In Data Guard environments, additional options like `--prepareStandbyBlob` can be used to manage password changes on both primary and standby databases.

### Example 6-10 dbascli database changePassword

```
dbascli database changepassword --dbname db19
```

## dbascli database convertToPDB

To convert the specified non-CDB database to PDB, use the `dbascli database convertToPDB` command.

### Syntax

```
dbascli 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:

```
dbaascli database convertToPDB --dbname tst19 --cdbname cdb19 --copyDatafiles
```

## dbaascli database create

To create Oracle Database, use the `dbaascli 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

```
dbaascli 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` specifies `true` or `false` to create database as CDB or Non-CDB
- `--pdbName` specifies the name of the PDB
- `--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` resumes or reverts to 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` specifies `true` or `false` to lock the PDB admin user account. Default value is `true`
- `--dbcaTemplateFilePath` specifies 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`

### Frequently Asked Questions

#### Q: What does the dbascli database create command do?

A: The `dbascli database create` command is used to create a new Oracle Database instance. It supports creating Oracle Database version 12.1.0.2 or higher with the release update JAN 2021 or higher.

#### Q: How do I specify the name of the Oracle Database to create?

A: Use the `--dbName` option to specify the name of the Oracle Database. For example:

```
dbascli database create --dbName <db_name>
```

**Q: How can I create a Container Database (CDB)?**

A: Use the `--createAsCDB` option and specify `true` to create the database as a CDB. For example:

```
dbascli database create --dbName <db_name> --createAsCDB true
```

**Q: How do I specify the Oracle Home for the database?**

A: You can use either the `--oracleHome` option to specify the location of the Oracle home or the `--oracleHomeName` option to specify the name of the Oracle home.

**Q: How do I specify a unique database name or SID?**

A: Use the `--dbUniqueName` option to specify a unique name for the database and the `--dbSID` option to specify the SID of the database.

**Q: How do I create a Pluggable Database (PDB) along with a CDB?**

A: You can use the `--pdbName` option to specify the name of the PDB, and the `--pdbAdminUserName` option to set the PDB admin username. For example:

```
dbascli database create --dbName <db_name> --createAsCDB true --pdbName
<pdb_name> --pdbAdminUserName <admin_user>
```

**Q: How can I specify the database character set and national character set?**

A: Use the `--dbCharset` option to specify the database character set and the `--dbNCharset` option to specify the national character set. For example:

```
dbascli database create --dbName <db_name> --dbCharset AL32UTF8 --dbNCharset
AL16UTF16
```

**Q: How do I set the memory settings (SGA and PGA) for the database?**

A: Use the `--sgaSizeInMB` option to specify the SGA size and the `--pgaSizeInMB` option to specify the PGA size, both in megabytes.

**Q: How do I specify the destination for datafiles and Fast Recovery Area (FRA)?**

A: Use the `--datafileDestination` option to specify the ASM disk group for datafiles, and the `--fraDestination` option to specify the ASM disk group for the FRA. You can also set the FRA size with the `--fraSizeInMB` option.

**Q: Can I configure Transparent Data Encryption (TDE) during database creation?**

A: Yes, you can configure TDE using the `--tdeConfigMethod` option. Valid values are `FILE` (for file-based encryption) or `KMS` (for using Oracle Key Management Service). If using KMS, provide the KMS key OCID with the `--kmsKeyOCID` option.

**Q: How do I create the database on a specific list of nodes?**

A: Use the `--nodeList` option to specify a comma-separated list of nodes where the database should be created.

**Q: How can I resume or revert a previous database creation attempt?**

A: Use the `--resume` option to resume the previous execution or the `--revert` option to roll back the previous run. You can also specify a `--sessionID` to resume or revert a specific session.



**Q: What does the --executePrereqs option do?**

A: The --executePrereqs option only runs the prerequisites for the database creation operation, without actually creating the database. Use `yes` or `no` to enable or disable this option.

**Q: Can I specify a custom DBCA template for the database creation?**

A: Yes, use the --dBCATemplateFilePath option to provide the absolute path of the DBCA template file that should be used to create the database.

**Q: Can I run the database creation operation in the background?**

A: Yes, you can use the --waitForCompletion option to specify whether the command should wait for the database creation to complete (`true`) or run the operation in the background (`false`).

**Q: What happens if I don't specify the --dbUniqueName option?**

A: If you don't specify a unique name for the database using --dbUniqueName, the system will automatically generate one based on the provided --dbName.

**Q: Can I lock the PDB admin account during the creation of a CDB?**

A: Yes, you can use the --lockPDBAdminAccount option and set it to `true` to lock the PDB admin account after database creation. By default, this value is set to `true`.

**Example 6-12 dbascli database create**

```
dbascli database create --dbName db19 --oracleHomeName myhome19 --dbSid db19sid --nodeList node1,node2 --createAsCDB true
```

## dbascli database delete

To delete an Oracle Database, use the `dbascli database delete` command.

### Prerequisite

Run the command as the `root` user.

### Syntax

```
dbascli 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.

### Frequently Asked Questions

#### Q: What is the purpose of the dbascli database delete command?

A: The `dbascli database delete` command is used to delete an Oracle Database on Exadata Cloud Infrastructure.

#### Q: How do I specify the database I want to delete?

A: Use the `--dbname` option to specify the name of the Oracle Database you want to delete. For example:

```
dbascli database delete --dbname <db_name>
```

#### Q: How do I delete the archive logs when deleting a database?

A: You can delete the archive logs by setting the `--deleteArchiveLogs` option to `true`. For example:

```
dbascli database delete --dbname <db_name> --deleteArchiveLogs true
```

#### Q: Can I also delete backups when deleting the database?

A: Yes, use the `--deleteBackups` option and set it to `true` to delete all associated backups. For example:

```
dbascli database delete --dbname <db_name> --deleteBackups true
```

#### Q: How can I run only the prechecks for the delete operation without actually deleting the database?

A: You can use the `--precheckOnly` option and set it to `yes` to run the prechecks without deleting the database. For example:

```
dbascli database delete --dbname <db_name> --precheckOnly yes
```

#### Q: How do I force the deletion of a database?

A: To force the deletion of a database, use the `--force` flag. This bypasses checks and forces the deletion process. For example:

```
dbascli database delete --dbname <db_name> --force
```

#### Q: How do I run the delete operation in the background?

A: Use the `--waitForCompletion` option and set it to `false` to run the operation in the background. For example:

```
dbascli database delete --dbname <db_name> --waitForCompletion false
```

#### Q: Can I specify the SID of the database I want to delete?

A: Yes, you can specify the SID of the database using the `--dbSID` option. For example:

```
dbaascli database delete --dbname <db_name> --dbSID <sid>
```

**Q: How do I resume a previously interrupted delete operation?**

A: To resume a previous delete execution, use the `--resume` option. You can also specify a session ID using the `--sessionID` option if needed. For example:

```
dbaascli database delete --dbname <db_name> --resume --sessionID <session_id>
```

**Q: What user privileges are required to run the dbaascli database delete command?**

A: The command must be run as the `root` user.

**Q: What does the `--precheckOnly` option do in the dbaascli database delete command?**

A: The `--precheckOnly` option allows you to run only the prechecks for the delete operation without actually deleting the database. It ensures that all checks pass before proceeding with the actual deletion.

**Q: Can I delete a database without waiting for the operation to complete?**

A: Yes, by setting the `--waitForCompletion` option to `false`, the delete operation will run in the background, and you don't have to wait for it to complete.

**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

**Frequently Asked Questions**

**Q: What is the purpose of the dbaascli database deleteInstance command?**

A: The `dbaascli database deleteInstance` command is used to delete a specific Oracle Database instance on a specified node in an Exadata Cloud Infrastructure environment.

**Q: How do I specify which Oracle Database instance to delete?**

A: You can specify the Oracle Database instance to delete by using the `--dbname` option to provide the database name and the `--node` option to provide the node name. For example:

```
dbaascli database deleteInstance --dbname <db_name> --node <node_name>
```

**Q: Can I delete the instance even if the node is unreachable?**

A: Yes, you can use the `--continueOnUnreachableNode` option to proceed with the deletion, even if the specified node is unreachable. For example:

```
dbaascli database deleteInstance --dbname <db_name> --node <node_name> --  
continueOnUnreachableNode
```

**Q: What happens if the node specified is unreachable during the delete instance operation?**

A: If the node is unreachable and the `--continueOnUnreachableNode` option is not used, the operation will fail. If the option is used, the operation will continue even if the node cannot be accessed.

**Q: How do I delete a database instance from a specific node?**

A: Use the following command to delete a database instance from a specific node:

```
dbaascli database deleteInstance --dbname <db_name> --node <node_name>
```

**Q: What user privileges are required to run the dbaascli database deleteInstance command?**

A: The command must be run as the `root` user.

**Q: Can I delete an instance without specifying the node?**

A: No, the `--node` option is required to specify which node the database instance should be deleted from.

**Q: What does the --continueOnUnreachableNode option do?**

A: The `--continueOnUnreachableNode` option allows the operation to proceed even if the specified node cannot be reached, ensuring that the instance deletion continues in scenarios where the node might be down.

**Q: Is it possible to delete multiple database instances at once using this command?**

A: No, the `dbaascli database deleteInstance` command is used to delete a single database instance on a specified node at a time. You would need to run the command separately for each instance you want to delete.

**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...`

## Frequently Asked Questions

### Q: What is the purpose of the dbascli database duplicate command?

A: The `dbascli database duplicate` command is used to create a new Oracle Database by duplicating an existing active database.

### Q: What are the prerequisites for using the dbascli database duplicate command?

A: You must run the command as the `root` user.

### Q: How do I specify the source database for duplication?

A: Use the `--sourceDBConnectionString` option to provide the source database connection string in the format `<scan_name>:<scan_port>/<database_service_name>`. For example:

```
--sourceDBConnectionString <scan_name>:<scan_port>/<database_service_name>
```

### Q: How do I specify the location of the Oracle Home for the new database?

A: You can specify the Oracle Home location using the `--oracleHome` option or the Oracle Home name using the `--oracleHomeName` option. For example:

```
--oracleHome <value>
```

or

```
--oracleHomeName <value>
```

**Q: What is the purpose of the --sourceDBWalletLocation option?**

A: The --sourceDBWalletLocation option specifies the location of the source database TDE wallet file, which is required to duplicate the database from an active source database.

**Q: Can I skip duplicating specific PDBs from the source database?**

A: Yes, you can use the --skipPDBs option to specify a comma-delimited list of PDB names that should be excluded from the duplicate operation. For example:

```
--skipPDBs pdb1,pdb2
```

**Q: How do I configure TDE for the new database?**

A: Use the --tdeConfigMethod option to specify the TDE configuration method (FILE or KMS). If you choose KMS, you can provide the KMS key OCID using the --kmsKeyOCID option. For example:

```
--tdeConfigMethod FILE
```

or

```
--tdeConfigMethod KMS --kmsKeyOCID <value>
```

**Q: What does the --executePrereqs option do?**

A: The --executePrereqs option specifies whether to run only the prerequisite checks for the operation. Valid values are yes to run prereqs only, or no to proceed with the full operation.

**Q: How can I resume a previously interrupted duplicate operation?**

A: Use the --resume option along with the --sessionID option to resume a previously interrupted duplicate operation. For example:

```
--resume --sessionID <value>
```

**Q: What does the --waitForCompletion option do?**

A: The --waitForCompletion option specifies whether to wait for the operation to complete. Setting this to true waits for completion, while false runs the operation in the background. For example:

```
--waitForCompletion true
```

**Q: What is the purpose of the --rmanParallelism option?**

A: The --rmanParallelism option specifies the parallelism value for RMAN (Recovery Manager) during the duplication process. This can improve the speed of the duplication operation by using multiple parallel processes.

**Q: How do I specify the size of the SGA and PGA for the new database?**

A: Use the --sgaSizeInMB and --pgaSizeInMB options to specify the sizes of the SGA (System Global Area) and PGA (Program Global Area) in megabytes, respectively. For example:

```
--sgaSizeInMB <value>
```

```
--pgaSizeInMB <value>
```

**Q: What does the --revert option do?**

A: The `--revert` option is used to roll back a previous duplicate operation. You need to provide the `--sessionID` to specify which session to revert.

**Example 6-15 dbaascli database duplicate**

```
dbaascli database duplicate --sourceDBConnectionString test-user-
scan.dbaastoolslrqsu.dbaastoolslrqvc.oraclevcn.com:1521/
mynew.dbaastoolslrqsu.dbaastoolslrqvc.oraclevcn.com --oracleHome /u02/app/
oracle/product/19.0.0.0/dbhome_2 --dbName newdup --
sourceDBWalletLocation /var/opt/oracle/dbaas_acfs/tmp/prim_wallet
```

## 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.

**Frequently Asked Questions****Q: What is the purpose of the dbaascli database getDetails command?**

A: The `dbaascli database getDetails` command shows detailed information about a specified Oracle database, including the database name, node information, and pluggable database (PDB) details.

**Q: Who can run the dbaascli database getDetails command?**

A: The command can be run by the `root` user or the `oracle` user.

**Q: What does the --dbname option specify in the dbaascli database getDetails command?**

A: The `--dbname` option specifies the name of the Oracle database for which detailed information is being retrieved.

**Q: What kind of information does the dbaascli database getDetails command provide?**

A: The command provides details such as the database name, node information, and information about pluggable databases (PDBs) associated with the container database.



## dbaascli database getPDBs

To view the list of all pluggable databases in a container database, use the `dbaascli database getPDBs` command.

Run the command as the `root` or `oracleuser`.

### Syntax

```
dbaascli database getPDBs --dbname <value>
```

Where:

- `--dbname` specifies the name of the container database

### Frequently Asked Questions

**Q: What is the purpose of the dbaascli database getPDBs command?**

A: The `dbaascli database getPDBs` command is used to list all the pluggable databases (PDBs) within a specified container database (CDB).

**Q: How do I specify the container database for the getPDBs command?**

A: You use the `--dbname` option to specify the name of the container database. For example:

```
--dbname <value>
```

**Q: Do I need to run the dbaascli database getPDBs command as a specific user?**

A: Yes, you must run the command as either the `root` user or the `oracle` user.

**Q: Can I view PDBs in a non-CDB database using the getPDBs command?**

A: No, the `getPDBs` command is only applicable to container databases (CDBs). You cannot use this command for non-CDB databases.

**Q: What is the format of the output from the dbaascli database getPDBs command?**

A: The command returns a list of all PDBs within the specified container database. The output typically includes PDB names, statuses, and other relevant details about each pluggable database.

**Q: Can this command be used for multiple databases at once?**

A: No, the `dbaascli database getPDBs` command works with a single container database at a time, specified by the `--dbname` option.

**Q: Is it necessary to shut down the database to use the getPDBs command?**

A: No, the `getPDBs` command does not require the database to be shut down. It can be run while the container database is operational.

**Example 6-16 dbaascli database getPDBs --dbname**

```
dbaascli database getPDBs --dbname apr_db1
```

## 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>[--instance <value>] [--backupPrepared] [--
allowBounce]|
--resetParameters <values> [--instance <value>] [--backupPrepared] [--
allowBounce]
}
--responseFile
[--backupPrepared]
[--instance]
[--allowBounce]
[--waitForCompletion]
```

### 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.
- `--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.
- `--backupPrepared` acknowledges that a proper database backup is in place prior to modifying critical or sensitive parameters.
- `--allowBounce` grants permission to bounce the database in order to reflect the changes on applicable static parameters.
- `--waitForCompletion` specify `false` to run the operation in background. Valid values : `true|false.`]

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli database modifyParameters` command?

A: The `dbaascli database modifyParameters` command is used to modify or reset Oracle Database initialization parameters.

#### Q: How do I specify the database for which I want to modify parameters?

A: You must use the `--dbname` option to specify the name of the database for which you want to modify or reset parameters.

**Q: How can I modify database parameters using the modifyParameters command?**

A: Use the `--setParameters` option followed by a comma-delimited list of parameters and their new values. For example:

```
--setParameters parameter1=valueA,parameter2=valueB
```

**Q: How do I reset parameters to their default values using this command?**

A: Use the `--resetParameters` option followed by a comma-delimited list of parameters to reset to their default values. For example:

```
--resetParameters parameter1,parameter2
```

**Q: Can I modify parameters using a response file?**

A: Yes, you can specify the absolute location of a response JSON file using the `--responseFile` option. The file should contain the parameters you want to modify.

**Q: Is it necessary to take a backup before modifying parameters?**

A: While not mandatory for all changes, if you are modifying critical or sensitive parameters, it's recommended to have a backup in place. You can use the `--backupPrepared` option to acknowledge that a backup has been prepared.

**Q: Can I apply changes only to a specific instance in a multi-instance database?**

A: Yes, you can specify the instance name using the `--instance` option. If this option is not used, the changes will be applied at the database level.

**Q: Will the database need to be bounced (restarted) after modifying parameters?**

A: For some static parameters, a database bounce is required. You can use the `--allowBounce` option to grant permission for the database to bounce if necessary.

**Q: What happens if I don't allow the database to bounce when changing static parameters?**

A: If you do not use the `--allowBounce` option when modifying static parameters, the changes will not take effect until the next manual database restart.

**Q: Can I resume modifying parameters if an earlier session was interrupted?**

A: No, this command does not support session resumption. You will need to re-run the command from the beginning.

**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

```
dbascli database move
{
  --oracleHome <value> | --oracleHomeName <value>
}
--dbname <value>
[--executePrereqs]
[--resume [--sessionID <value>]]
[--rollback [--sessionID <value>]]
[--skipDatapatch]
[--skipPDBs <value>]
[--skipClosedPDBs]
[--continueWithDbDowntime]
[--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
- `--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.
- `--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

## Frequently Asked Questions

### Q: What is the dbascli database move command used for?

A: The `dbascli database move` command is used to move a database from one Oracle home to another.

**Q: What are the prerequisites for using the dbascli database move command?**

A: Before performing a move operation, ensure that all database instances associated with the database are up and running. Additionally, the command must be run as the `root` user.

**Q: What does the `--oracleHome` parameter specify?**

A: The `--oracleHome` parameter specifies the path of the Oracle home to which the database will be moved.

**Q: What does the `--oracleHomeName` parameter specify?**

A: The `--oracleHomeName` parameter specifies the name of the Oracle home to which the database will be moved.

**Q: What is the purpose of the `--dbname` parameter?**

A: The `--dbname` parameter specifies the name of the database that you want to move.

**Q: What does the `--executePrereqs` option do?**

A: The `--executePrereqs` option runs the prerequisite checks and reports the results.

**Q: What is the `--resume` option used for?**

A: The `--resume` option resumes a previously interrupted or incomplete move operation.

**Q: How is `--sessionID` used in the command?**

A: The `--sessionID` specifies the session ID to resume a previous run or rollback.

**Q: What does the `--rollback` option do?**

A: The `--rollback` option rolls the database back to its previous Oracle home.

**Q: What does the `--skipDatapatch` option do?**

A: The `--skipDatapatch` option skips running the datapatch on the databases during the move operation.

**Q: What is the function of the `--skipPDBs` option?**

A: The `--skipPDBs` option skips running the datapatch on a specified comma-delimited list of PDBs (e.g., `pdb1, pdb2`).

**Q: What does the `--skipClosedPDBs` option do?**

A: The `--skipClosedPDBs` option skips patching of closed PDBs.

**Q: What does `--continueWithDbDowntime` mean?**

A: The `--continueWithDbDowntime` option allows the move operation to proceed even if there is only one active instance up, allowing for downtime during the process.

**Q: What is the purpose of the `--allowParallelDBMove` option?**

A: The `--allowParallelDBMove` option allows the database move to be performed in parallel, potentially speeding up the process.

**Q: What does `--waitForCompletion` specify?**

A: The `--waitForCompletion` option specifies whether to wait for the operation to complete. Setting it to `false` runs the operation in the background.

**Q: How do I use the `--nodeList` parameter?**

A: The `--nodeList` parameter specifies a comma-delimited list of nodes on which the move operation will be performed, if it is not to be applied to all nodes.

**Q: What should I do if I encounter issues with the dbascli database move command?**

A: Ensure all database instances are running and verify that you are running the command as the `root` user. If issues persist, consult the detailed command documentation or open a support ticket with Oracle.

**Q: Can I perform a move operation if one of the database instances is down?**

A: No, all associated database instances must be up and running before performing the move operation.

**Q: What happens if the move operation is interrupted?**

A: You can use the `--resume` option to continue the move operation from where it left off by using the same session or specifying the `--sessionID`.

**Q: What does the --allowParallelDBMove option do?**

A: It enables the database move to be performed in parallel, which can reduce the time it takes to complete the operation, especially on larger environments.

**Q: How do I monitor the progress of a move operation that is running in the background?**

A: When using `--waitForCompletion false`, the command does not wait for the operation to complete. You can check the status of the operation manually using appropriate Oracle logs or status commands.

**Q: What is the significance of the --skipClosedPDBs option?**

A: It skips patching for PDBs that are closed, reducing the operation time if there are PDBs that don't need to be patched.

**Q: Can the database move be rolled back at any time?**

A: Yes, the move operation can be rolled back using the `--rollback` option, either by specifying the session ID or simply rolling back to the previous Oracle home.

**Q: What is the role of --nodeList in a multi-node environment?**

A: In a multi-node environment, you can restrict the move operation to specific nodes by providing a comma-delimited list of node names with `--nodeList`.

**Q: Can I move the database to a new Oracle home while skipping specific nodes in a multi-node environment?**

A: Yes, you can use the `--nodeList` option to specify which nodes to include in the move operation. Any nodes not listed will be skipped.

**Q: What is the maximum number of nodes I can specify with the --nodeList parameter?**

A: The `--nodeList` parameter allows you to specify a comma-delimited list of as many nodes as needed, limited only by your environment configuration. Ensure all nodes are valid and reachable.

**Q: How do I know which PDBs are closed before using the --skipClosedPDBs option?**

A: You can query the `v$pdbs` view to check the status of your PDBs. Any PDBs with a status of "MOUNTED" or "CLOSED" will be skipped when using `--skipClosedPDBs`.

**Q: How do I verify if a rollback has completed successfully?**

A: After running the rollback command, you can review the database logs or use the Oracle alert logs to verify that the database has been successfully rolled back to its previous Oracle home.

**Q: Is there a way to force the move operation if some prerequisites fail?**

A: The move command enforces prerequisite checks for system stability. You cannot bypass critical prerequisite failures. Address any issues reported by the `--executePrereqs` option before proceeding with the move.

**Example Use Cases**

**Example 1: Basic Database Move by Oracle Home Path**

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL
```

Moves the database *ORCL* to the Oracle home located at `/u01/app/oracle/product/19.0.0/dbhome_1`.

**Example 2: Database Move by Oracle Home Name**

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL
```

Moves the database *ORCL* to the Oracle home named `DB_HOME_NAME`.

**Example 3: Running Prerequisite Checks Before Moving**

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --executePrereqs
```

Moves the database *ORCL* to the Oracle home while running the prerequisite checks beforehand.

**Example 4: Resuming a Previous Move Operation**

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --resume
```

Resumes a previous move operation for the *ORCL* database.

**Example 5: Resuming a Move Operation with a Specific Session ID**

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL --resume --  
sessionID 12345
```

Resumes the move operation for the *ORCL* database using session ID 12345.

**Example 6: Rolling Back a Move Operation**

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --rollback
```

Rolls back the move operation for the *ORCL* database, restoring it to the previous Oracle home.

**Example 7: Rolling Back a Move Operation with a Session ID**

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL --rollback --  
sessionID 67890
```

Rolls back the move operation for *ORCL* using session ID 67890.

### Example 8: Skipping Datapatch

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --skipDatapatch
```

Moves the database *ORCL* without running datapatch on the databases.

### Example 9: Skipping Specific PDBs During Datapatch

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --skipPDBs pdb1,pdb2
```

Moves the *ORCL* database to a new Oracle home but skips running datapatch on the specified PDBs (pdb1 and pdb2).

### Example 10: Skipping Datapatch on Closed PDBs

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --skipClosedPDBs
```

Moves the *ORCL* database and skips running datapatch on any closed PDBs.

### Example 11: Allowing Database Downtime During Move

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL --  
continueWithDbDowntime
```

Moves the *ORCL* database to the specified Oracle home while allowing downtime of the database during the move process.

### Example 12: Moving Database in Parallel

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --allowParallelDBMove
```

Moves the database *ORCL* to the specified Oracle home with the option to run the move in parallel for better performance.

### Example 13: Running the Operation in the Background

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --waitForCompletion false
```

Moves the database *ORCL* to a new Oracle home but runs the operation in the background.

### Example 14: Specifying Nodes for the Move

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --nodeList node1,node2
```

Moves the database *ORCL* to the specified Oracle home but performs the operation only on node1 and node2.

### Example 15: Combination of Move with Prerequisite Checks, Skipping Specific PDBs, and Allowing Downtime

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL --  
executePrereqs --skipPDBs pdb1 --continueWithDbDowntime
```

Moves the *ORCL* database to the specified Oracle home, runs prerequisite checks, skips running datapatch on pdb1, and allows database downtime during the operation.

### Example 16: Combination of Move in Parallel and Running in the Background



```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --allowParallelDBMove --waitForCompletion false
```

Moves the *ORCL* database to a new Oracle home, runs the move in parallel, and executes the operation in the background.

#### **Example 17: Combining Move with Parallel Execution and Skipping Closed PDBs**

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --  
dbname TESTDB --allowParallelDBMove --skipClosedPDBs
```

Moves the *TESTDB* database to the new Oracle home `/u02/app/oracle/product/19.0.0/dbhome_2`, while running the operation in parallel and skipping datapatch on closed PDBs.

#### **Example 18: Running Prerequisites Check Only**

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --  
dbname PRODDB --executePrereqs
```

Checks the prerequisites for moving the *PRODDB* database to the Oracle home located at `/u02/app/oracle/product/19.0.0/dbhome_2` without actually performing the move.

#### **Example 19: Skipping Datapatch for Specific PDBs**

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --  
dbname HRDB --skipPDBs pdb1,pdb3
```

Moves the *HRDB* database to the new Oracle home, but skips running datapatch for *pdb1* and *pdb3*.

#### **Example 20: Running the Move on Specific Nodes**

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --  
dbname FINDB --nodeList node1,node3
```

Moves the *FINDB* database to the new Oracle home only on *node1* and *node3*.

#### **Example 21: Database Move with Downtime Allowed**

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
dbname ORCL --continueWithDbDowntime
```

Moves the *ORCL* database to the specified Oracle home while allowing downtime during the move operation.

#### **Example 22: Combination of Parallel Move and Skipping Datapatch**

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --  
dbname CRMDB --allowParallelDBMove --skipDatapatch
```

Moves the *CRMDB* database in parallel, skipping the datapatch process.

#### **Example 23: Move Operation in Background with a Node List**

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --  
dbname SALESDB --waitForCompletion false --nodeList node2,node3
```

Moves the *SALESDB* database to the specified Oracle home in the background, and the operation is applied only on *node2* and *node3*.

#### **Example 24: Database Move with Prerequisite Check and Allowing Parallel Move**

```
dbascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_2 --
dbname ORCL --executePrereqs --allowParallelDBMove
```

Moves the *ORCL* database to the new Oracle home after performing the prerequisite checks and running the move operation in parallel.

#### Example 25: Rolling Back a Move Operation and Skipping Closed PDBs

```
dbascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname DEVDB --rollback --skipClosedPDBs
```

Rolls back the move operation for the *DEVDB* database, skipping any closed PDBs.

#### Example 26: Moving the Database with Specific Downtime and Parallel Execution

```
dbascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname FINDB --allowParallelDBMove --continueWithDbDowntime
```

Moves the *FINDB* database to the specified Oracle home while allowing database downtime and enabling parallel execution to speed up the process.

#### Example 27: Checking Database Move Prerequisites without Executing the Move

```
dbascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname HRDB --executePrereqs
```

Runs prerequisite checks to validate that the *HRDB* database can be moved to the specified Oracle home without executing the move itself.

#### Example 28: Moving Database and Running the Command in the Background on Specific Nodes

```
dbascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_3 --
dbname PRODDB --waitForCompletion false --nodeList node1,node4
```

Moves the *PRODDB* database to a new Oracle home, executing the operation in the background, and applying it only on node1 and node4.

#### Example 29: Combining Prerequisite Checks, Skipping Closed PDBs, and Allowing Parallel Execution

```
dbascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname CRMDB --executePrereqs --skipClosedPDBs --allowParallelDBMove
```

Performs prerequisite checks before moving the *CRMDB* database to the new Oracle home, skips patching closed PDBs, and allows the operation to run in parallel for faster execution.

#### Example 30: Database Move with Rollback on Specific Session ID and Skipping Datapatch

```
dbascli database move --oracleHomeName DB_HOME_NAME --dbname DEVDB --rollback --
sessionID 45678 --skipDatapatch
```

Rolls back a previously executed move operation for the *DEVDB* database to its prior Oracle home using session ID 45678, skipping the datapatch process during the rollback.

#### Example 31: Moving Database with Allow Parallel Execution and Specifying Datapatch Skipping for PDBs

```
dbascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_3 --
dbname ANALYTICDB --allowParallelDBMove --skipPDBs pdb2,pdb4
```

Moves the *ANALYTICDB* database in parallel to the specified Oracle home and skips the datapatch process for pdb2 and pdb4.

## 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>
```

### Frequently Asked Questions

#### Q: What is the purpose of the dbaascli database recover command?

A: The `dbaascli database recover` command is used to recover an Oracle Database from backups stored in a backup storage destination.

#### Q: How do I specify which database to recover?

A: You can specify the database you want to recover using the `--dbname` option followed by the database name. For example:

```
--dbname <database_name>
```

#### Q: What are the recovery options available with the dbaascli database recover command?

A: The recovery options are:

--untilTime: Recover the database to a specific time.

--untilSCN: Recover the database to a specific System Change Number (SCN).

--latest: Recover the database to the last known state.

--tag: Recover the database using an archive tag.

**Q: How do I recover the database to a specific time?**

A: Use the --untilTime option followed by the time in the format DD-MON-YYYY HH24:MI:SS. For example:

```
--untilTime 05-SEP-2024 15:30:00
```

**Q: How do I recover the database to a specific SCN?**

A: Use the --untilSCN option followed by the SCN value. For example:

```
--untilSCN 123456789
```

**Q: How can I recover the database to the latest known state?**

A: Use the --latest option to recover the database to the most recent state possible. For example:

```
--latest
```

**Q: What is the use of the --tag option in the recovery process?**

A: The --tag option allows you to recover the database using an archival tag associated with the backups. For example:

```
--tag <backup_tag>
```

**Q: How can I check the status of a recovery operation?**

A: Use the --status option along with the --uuid value to check the status of an ongoing or previous recovery operation. For example:

```
--status --uuid <recovery_uuid>
```

**Q: What does the --start option do in the recovery process?**

A: The --start option initiates the recovery operation based on the selected recovery method (--untilTime, --untilSCN, --latest, or --tag).

**Q: Is there a way to recover the database without specifying a time or SCN?**

A: Yes, you can recover the database to its last known state using the --latest option, which doesn't require specifying a time or SCN.

**Q: Can I perform a partial recovery?**

A: Yes, you can recover the database to a specific point in time or SCN using the --untilTime or --untilSCN options, respectively.

**Example 6-18 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`:

```
dbascli database recover --dbname myTestDb --status --uuid
2508ea18be2911eb82d0020017075151
```

## dbascli database runDatapatch

To patch an Oracle Database, use the `dbascli 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

```
dbascli database runDatapatch --dbname
[--resume]
    [--sessionID]
[--skipPdb | --pdb]
[--executePrereqs]
[--patchList]
[--skipClosedPdb]
[--rollback]
```

### Where:

- `--dbname` specifies the name of the database
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--skipPdb` skips running the datapatch on a specified comma-delimited list of PDBs. For example: `pdb1,pdb2...`
- `--pdb` 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...`
- `--skipClosedPdb` skips running the datapatch on closed PDBs
- `--rollback` rolls back the patches applied

### Frequently Asked Questions

#### Q: What is the purpose of the `dbascli database runDatapatch` command?

A: The `dbascli database runDatapatch` command is used to apply patches to an Oracle Database.

#### Q: What must be ensured before running the `dbascli database runDatapatch` command?

A: Before running the command, ensure that all instances of the database are up and running.

**Q: How do I specify which database to patch?**

A: Use the `--dbname` option followed by the name of the database. For example:

```
--dbname myDatabase
```

**Q: How do I resume a previously interrupted runDatapatch operation?**

A: Use the `--resume` option to resume the previous run or the `--sessionID` option to specify a specific session ID. For example:

```
--resume  
  
--sessionID 12345
```

**Q: How can I skip certain PDBs when running the patch?**

A: Use the `--skipPdfs` option followed by a comma-delimited list of PDB names to skip. For example:

```
--skipPdfs pdb1,pdb2
```

**Q: How can I run the patch only on certain PDBs?**

A: Use the `--pdfs` option followed by a comma-delimited list of PDB names to include. For example:

```
--pdfs pdb1,pdb2
```

**Q: How do I apply or roll back a specific set of patches?**

A: Use the `--patchList` option followed by a comma-delimited list of patch names to apply or roll back. For example:

```
--patchList patch1,patch2
```

**Q: What does the --rollback option do?**

A: The `--rollback` option rolls back the patches that were applied during the patching operation.

**Q: What happens if some PDBs are closed during the patching operation?**

A: If some PDBs are closed, you can use the `--skipClosedPdfs` option to skip patching those closed PDBs.

**Q: Can I run prerequisite checks before applying patches?**

A: Yes, use the `--executePrereqs` option to run prerequisite checks before applying the patch.

**Q: How do I find out which session ID to resume a patch?**

A: After a `runDatapatch` operation, the session ID is typically logged. Use the `--sessionID` option to specify that ID when resuming a patch. For example:

```
--sessionID 67890
```

```
dbaascli database runDatapatch --dbname db19
```

## dbaascli database createTemplate

To create database templates (DBCA templates) that can subsequently be used to create databases, use the `dbaascli database createTemplate` command.

### Prerequisites:

Run the command as the `root` user.

### Syntax

Create a new DBCA template from the specified database.

```
dbaascli database createTemplate --dbname <value> --templateLocation <value>
[--templateName <value>]
[--rmanParallelism <value>]
```

### Where:

- `--dbname` specifies the name of the database.
- `--templateLocation` specifies the template name.
- `--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.

### Frequently Asked Questions

#### Q: What is the purpose of the dbaascli database start command?

A: The `dbaascli database start` command is used to start an Oracle Database.

#### Q: What must be done before running the dbaascli database start command?

A: The command must be run as the `root` user.

**Q: How do I specify the database I want to start?**

A: Use the `--dbname` option followed by the name of the database. For example:

```
--dbname myDatabase
```

**Q: What are the possible modes in which I can start the database?**

A: You can start the database in `mount` or `nomount` mode using the `--mode` option. For example:

```
--mode mount
```

**Q: What is the default mode if I don't specify one?**

A: If you don't specify a mode, the database will start in the default `open` mode.

**Q: Will this command open all PDBs in Oracle Database 12c or later?**

A: Yes, when starting the database in Oracle Database 12c or later, all pluggable databases (PDBs) will also be opened.

**Q: How can I start a database in nomount mode?**

A: Use the `--mode` option and set it to `nomount`. For example:

```
--mode nomount
```

**Q: How can I start a database in mount mode?**

A: Use the `--mode` option and set it to `mount`. For example:

```
--mode mount
```

**Q: Is it mandatory to specify a database name when running the dbascli database start command?**

A: Yes, it is recommended to specify the database name using the `--dbname` option to ensure the correct database is started.

**Example 6-19 dbascli database start**

```
dbascli database start --dbname dbname --mode mount
```

## dbascli database status

To check the status of an Oracle Database, use the `dbascli database status` command.

### Prerequisites

Run the command as the `root` user.

### Syntax

```
dbascli 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 6-20 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.

### **Frequently Asked Questions**

#### **Q: What is the purpose of the dbaascli database stop command?**

A: The `dbaascli database stop` command is used to stop an Oracle Database.

#### **Q: What are the prerequisites for using the dbaascli database stop command?**

A: You must run the command as the `root` user, and you must connect to an Exadata Cloud@Customer virtual machine using SSH.

#### **Q: How do I specify which database to stop?**

A: You can specify the database by using the `--dbname` option followed by the name of the database. For example:

```
--dbname myDatabase
```

#### **Q: What are the valid shutdown modes for the dbaascli database stop command?**

A: The valid shutdown modes are:

```
abort
```

```
immediate
normal
transactional
```

**Q: What is the default shutdown mode if no mode is specified?**

A: If no mode is specified, the database will be shut down in `immediate` mode by default.

**Q: What happens in immediate shutdown mode?**

A: In `immediate` mode, no new connections or transactions are permitted, active transactions are rolled back, and all connected users are disconnected.

**Q: How can I stop the database in abort mode?**

A: To stop the database in abort mode, use the `--mode` option with `abort`. For example:

```
--mode abort
```

**Q: What does normal mode do when stopping the database?**

A: In normal mode, the database allows current user sessions to complete and then stops without affecting active transactions.

**Q: What is transactional mode used for in the dbascli database stop command?**

A: In `transactional` mode, the database stops only after all active transactions are completed, but no new transactions are allowed.

**Q: Is it mandatory to specify the shutdown mode in the dbascli database stop command?**

A: No, specifying a `shutdown` mode is optional. If not provided, the default `immediate` mode will be used.

**Example 6-21 dbascli database stop**

```
dbascli database stop --dbname db19
```

## dbascli database upgrade

To upgrade an Oracle Database, use the `dbascli database upgrade` command.

**Prerequisite**

Run the command as the `root` user.

**Syntax**

```
dbascli 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.

### Frequently Asked Questions

**Q: What is the purpose of the dbascli database upgrade command?**

A: The `dbascli database upgrade` command is used to upgrade an Oracle Database to a new version.

**Q: What are the prerequisites for using the dbascli database upgrade command?**

A: You must run the command as the `root` user and connect to an Exadata Cloud@Customer virtual machine using SSH.

**Q: How do I specify the database that needs to be upgraded?**

A: Use the `--dbname` option followed by the name of the database. For example:

```
--dbname myDatabase
```

**Q: How can I specify the target Oracle home for the upgrade?**

A: You can specify the target Oracle home location with the `--targetHome` option or the name of the target Oracle Database home with the `--targetHomeName` option.

**Q: What does the --standBy option do?**

A: The `--standBy` option is used to upgrade standby databases in Data Guard configurations.

**Q: What is the purpose of the `--allStandbyPrepared` flag?**

A: The `--allStandbyPrepared` flag acknowledges that all required operations on standby databases have been performed before upgrading the primary database in a Data Guard configuration.

**Q: What does the `--removeGRP` option do?**

A: The `--removeGRP` option automatically removes the Guaranteed Restore Point (GRP) backup if the database upgrade is successful.

**Q: When should I use the `--increaseCompatibleParameter` option?**

A: Use the `--increaseCompatibleParameter` option to automatically increase the compatible parameter during the database upgrade, provided the upgrade is successful.

**Q: What does the `--executePrereqs` option do?**

A: The `--executePrereqs` option runs only the pre-upgrade checks to ensure that the database is ready for the upgrade.

**Q: How do I handle a failed post-upgrade step?**

A: Use the `--postUpgrade` option to rerun the post-upgrade steps if the initial post-upgrade attempt fails.

**Q: What is the purpose of the `--revert` option?**

A: The `--revert` option reverts the Oracle Database to its original Oracle home, undoing the upgrade.

**Q: How can I pass additional arguments specific to DBUA for the upgrade?**

A: Use the `--upgradeOptions` option to pass DBUA-specific arguments for the Oracle Database upgrade. Refer to the Oracle documentation for supported arguments and options.

**Q: Is it mandatory to specify the target Oracle home for the upgrade?**

A: Yes, you must specify either the `--targetHome` or `--targetHomeName` to indicate the target Oracle home for the upgrade.

**Q: What should I do if I need to perform a pre-upgrade check but not proceed with the upgrade?**

A: Use the `--executePrereqs` option to perform only the pre-upgrade checks without proceeding with the actual upgrade.

**Example 6-22 dbaascli database upgrade pre-upgrade requisite checks**

```
dbaascli database upgrade --dbbname 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

```
dbascli 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

## Frequently Asked Questions

### Q: What is the purpose of the dbascli dataguard prepareStandbyBlob command?

A: The `dbascli dataguard prepareStandbyBlob` command is used to generate a blob file containing various files required on the standby site in a Data Guard environment.

### Q: What are the prerequisites for running the dbascli dataguard prepareStandbyBlob command?

A: The command should be run as the `root` or `oracle` user.

### Q: How do I specify the name of the Oracle Database for which I want to prepare the standby blob?

A: Use the `--dbname` option followed by the name of the Oracle Database. For example:

```
--dbname myDatabase
```

### Q: How do I specify the location where the standby blob file will be generated?

A: Use the `--blobLocation` option to specify the custom directory path where the standby blob file will be generated. For example:

```
--blobLocation /path/to/standby_blob
```

### Q: What does the --dbname option do in the command?

A: The `--dbname` option specifies the name of the Oracle Database for which the standby blob file is being prepared.

### Q: What is the purpose of the --blobLocation option?

A: The `--blobLocation` option defines the custom directory path where the standby blob file will be created.

### Q: Can I run the dbascli dataguard prepareStandbyBlob command as a user other than root or oracle?

A: No, the command must be run as either the `root` or `oracle` user.

### Q: Is it possible to use a relative path for the --blobLocation option?

A: It is recommended to use an absolute path for the `--blobLocation` option to ensure the standby blob file is created in the correct directory.

### Q: What should I do if I want to change the location where the standby blob file is generated?

A: Modify the `--blobLocation` option to specify a new directory path for the standby blob file.

**Q: Do I need to perform any additional steps after generating the standby blob file?**

A: Yes, after generating the standby blob file, you need to transfer it to the standby site and use it for the Data Guard configuration.

## 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.

**Frequently Asked Questions****Q: What is the purpose of the dbaascli dataguard updateDGConfigAttributes command?**

A: The `dbaascli dataguard updateDGConfigAttributes` command is used to update Data Guard automation attributes across all cluster nodes.

**Q: What are the prerequisites for running the dbaascli dataguard updateDGConfigAttributes command?**

A: The command must be run as either the `root` or `oracle` user.

**Q: How do I specify the attributes that I want to update using this command?**

A: Use the `--attributes` option followed by the attributes to be modified. The attributes should be in a comma-delimited format, such as `attribute=value`. For example:

```
--attributes attribute1=value1,attribute2=value2
```

**Q: What format should the --attributes option values be in?**

A: The `--attributes` option values should be in a comma-delimited format with each attribute specified as `attribute=value`.

**Q: Can I specify multiple attributes in the --attributes option?**

A: Yes, you can specify multiple attributes by separating them with commas. For example:

```
--attributes attribute1=value1,attribute2=value2
```

**Q: What happens if I provide an attribute that is not predefined in the Data Guard configuration file?**

A: If you provide an attribute that is not predefined, the command may fail or ignore the unrecognized attribute. Ensure that all attributes are predefined in the Data Guard configuration file.

**Q: Do I need to restart any services after updating Data Guard automation attributes?**

A: In most cases, you do not need to restart services after updating attributes. However, check the specific attributes and their impact to determine if a restart is required.

**Q: How can I verify if the Data Guard attributes have been successfully updated?**

A: After running the command, you can verify the updated attributes by checking the Data Guard configuration or using appropriate verification commands/tools specific to your setup.

**Q: What should I do if the command fails to update the attributes?**

A: Check the error messages for details on what went wrong. Ensure that you have specified the correct attributes and that they are predefined in the Data Guard configuration file. Verify user permissions and command syntax.

**Q: Is it possible to update attributes for only specific nodes using this command?**

A: No, the `dbaascli dataguard updateDGConfigAttributes` command updates attributes across all cluster nodes. If you need to update attributes for specific nodes, you may need to use different methods or commands.

## 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`.

### Frequently Asked Questions

**Q: What is the purpose of the dbaascli dbhome create command?**

A: The `dbaascli dbhome create` command is used to create a new Oracle Database home with the desired version.

**Q: What are the prerequisites for running the `dbaascli dbhome create` command?**

A: The command must be run as the `root` user.

**Q: How do I specify the Oracle Database version while creating a new Oracle Home?**

A: Use the `--version` option followed by the Oracle Database version in the format of five numeric segments separated by periods, such as 19.11.0.0.0.

**Q: What does the `--oracleHome` option specify?**

A: The `--oracleHome` option specifies the location where you want to install the Oracle Home. If not provided, the default location will be used.

**Q: Can I assign a custom name to the new Oracle Home?**

A: Yes, you can use the `--oracleHomeName` option to specify a user-defined name for the Oracle Home. If not specified, a default name will be used.

**Q: How do I enable or disable Unified Auditing in the new Oracle Home?**

A: Use the `--enableUnifiedAuditing` option and specify `true` to enable or `false` to disable Unified Auditing for the Oracle Home.

**Q: What does the `--imageTag` option do?**

A: The `--imageTag` option specifies the Oracle Home image tag, which can be used in cases where the image tag differs from the version.

**Q: What is an example of using the `dbaascli dbhome create` command with version and image tag?**

A: An example of the command with version and image tag is:

```
dbaascli dbhome create --version 19.8.0.0.0 --imageTag 19.8.0.0.0
```

This creates an Oracle Home for version 19.8.0.0.0 with the corresponding image tag.

**Q: What happens if I don't provide the `--oracleHome` or `--oracleHomeName` options?**

A: If `--oracleHome` is not provided, the Oracle Home will be installed in the default location. If `--oracleHomeName` is not specified, a default name will be assigned to the Oracle Home.

**Q: How can I verify if the Oracle Home creation was successful?**

A: After running the command, check the output logs for any success messages or errors. You can also verify the Oracle Home by navigating to the specified location or using Oracle tools like `oraInstRoot.sh`.

**Q: Is it possible to create multiple Oracle Homes with different versions on the same system?**

A: Yes, you can create multiple Oracle Homes with different versions by specifying different values for the `--version` and `--oracleHomeName` options.

**Q: What should I do if the Oracle Home creation fails?**

A: Check the output logs for detailed error messages. Verify that you have the correct version format, required permissions, and sufficient disk space. Correct any issues and try running the command again.



### Example 6-23 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

### Frequently Asked Questions

#### Q: What is the purpose of the dbaascli dbHome delete command?

A: The `dbaascli dbHome delete` command is used to delete a specified Oracle Database home from the system.

#### Q: What are the prerequisites for running the dbaascli dbHome delete command?

A: The command must be run as the `root` user.

#### Q: How do I specify the Oracle Home to be deleted?

A: You can specify the Oracle Home to be deleted using one of the following options:

`--oracleHome <value>` to provide the absolute path of the Oracle Home location.

`--oracleHomeName <value>` to provide the name of the Oracle Home.

#### Q: What is the difference between --oracleHome and --oracleHomeName options?

A:

`--oracleHome` specifies the physical location or path of the Oracle Home to be deleted.

`--oracleHomeName` specifies the user-defined name of the Oracle Home to be deleted.

#### Q: How can I resume a previously interrupted deletion process?

A: You can use the `--resume` option to resume a previous deletion process. If you know the specific session ID of the process, you can include it with the `--sessionID` option.

**Q: What is the `--sessionID` option used for in the `dbaascli dbHome delete` command?**

A: The `--sessionID` option is used to resume a specific session that was previously interrupted or failed during the deletion process.

**Q: What happens if I do not provide the `--resume` or `--sessionID` options?**

A: If the `--resume` or `--sessionID` options are not provided, the command will initiate a new deletion process instead of resuming an interrupted one.

**Q: Is there any way to confirm the deletion of the Oracle Home after running the command?**

A: You can verify the deletion by checking the output logs for success messages and ensuring the Oracle Home directory is no longer present at the specified location.

**Q: Can I delete an Oracle Home that is currently in use by a running database?**

A: No, the Oracle Home should not be in use by any running databases or services during the deletion process. Make sure to stop any related databases before running the delete command.

**Q: What should I do if the `dbaascli dbHome delete` command fails?**

A: Review the output logs for any error messages. Ensure that the Oracle Home is not in use, verify the correct Oracle Home location or name, and confirm that you have the necessary permissions. After resolving any issues, rerun the command or use the `--resume` option if needed.

**Q: Can I delete multiple Oracle Homes at once using the `dbaascli dbHome delete` command?**

A: No, the command only allows you to delete one Oracle Home at a time by specifying either the `--oracleHome` or `--oracleHomeName` option.

**Q: What is an example of deleting an Oracle Home by its name?**

A: Here is an example of deleting an Oracle Home by name:

```
dbaascli dbHome delete --oracleHomeName myOracleHome
```

This command deletes the Oracle Home with the name `myOracleHome`.

**Q: What is an example of deleting an Oracle Home by its location?**

A: Here is an example of deleting an Oracle Home by specifying its location:

```
dbaascli dbHome delete --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1
```

This command deletes the Oracle Home located at `/u01/app/oracle/product/19.0.0/dbhome_1`.

**Q: Can I cancel the deletion process once it has started?**

A: No, once the deletion process has started, it cannot be canceled. Ensure that the Oracle Home is ready for deletion before executing the 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.

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

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli dbHome getDatabases` command?

A: The `dbaascli dbHome getDatabases` command is used to view information about all Oracle Databases running from a specified Oracle Database home.

#### Q: How can I specify the Oracle Database home to check?

A: You can specify the Oracle Database home by using one of the following options:

`--oracleHomeName <value>` to specify the user-defined name of the Oracle home.

`--oracleHome <value>` to specify the full location (path) of the Oracle home.

#### Q: What is the difference between `--oracleHomeName` and `--oracleHome` options?

A:

`--oracleHomeName` refers to a user-defined name for the Oracle home.

`--oracleHome` refers to the physical location (or directory path) of the Oracle home on the system.

#### Q: How do I run the `dbaascli dbHome getDatabases` command?

A: To run the command, use the following syntax:

```
dbaascli dbHome getDatabases --oracleHomeName <value>
```

or

```
dbaascli dbHome getDatabases --oracleHome <value>
```

Ensure to run the command as the `root` user.

#### Q: Can I specify both the Oracle home name and Oracle home location in the same command?

A: No, you can only specify either `--oracleHomeName` or `--oracleHome` in a single command execution. Choose one option based on how you identify the Oracle home.

**Q: What kind of information does the dbascli dbHome getDatabases command return?**

A: The command returns information about all Oracle Databases running from the specified Oracle home. This includes details such as database names and statuses.

**Q: What is an example of using dbascli dbHome getDatabases with the Oracle home location?**

A: Here is an example of using the command with the Oracle home location:

```
dbascli dbHome getDatabases --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1
```

This command retrieves the list of databases running from the Oracle home located at /u01/app/oracle/product/19.0.0/dbhome\_1.

**Q: What is an example of using dbascli dbHome getDatabases with the Oracle home name?**

A: Here is an example of using the command with the Oracle home name:

```
dbascli dbHome getDatabases --oracleHomeName myOracleHome
```

This command retrieves the list of databases running from the Oracle home named myOracleHome.

**Q: Do I need any special permissions to run this command?**

A: Yes, you must run the command as the root user to view the information about Oracle Databases running from a specified Oracle home.

**Q: What should I check if the dbascli dbHome getDatabases command returns no databases?**

A: Ensure that you have specified the correct Oracle home name or location and that there are databases running from that Oracle home. Additionally, confirm that the Oracle home is properly configured and active.

**Q: Can I use the dbascli dbHome getDatabases command on multiple Oracle homes at once?**

A: No, the command works on a single Oracle home at a time. You must run the command separately for each Oracle home you want to query.

**Q: Is there a way to verify that the Oracle home specified in the command is correct?**

A: You can verify the Oracle home by checking the directory structure or the configuration details in your system to ensure the path or name provided matches the actual Oracle home.

**Q: What happens if I run the command without specifying an Oracle home or Oracle home name?**

A: The command requires either the `--oracleHome` or `--oracleHomeName` option to be specified. If neither option is provided, the command will fail to execute.

**Q: Can this command retrieve databases that are currently stopped?**

A: Yes, the command will list all databases associated with the specified Oracle home, regardless of whether they are currently running or stopped.

**Example 6-24 dbascli dbHome getDatabases --oracleHome**

```
dbascli 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

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli dbHome getDetails` command?

A: The `dbaascli dbHome getDetails` command is used to view detailed information about a specific Oracle home on the system.

#### Q: How do I specify the Oracle home I want to get details about?

A: You can specify the Oracle home using one of the following options:

`--oracleHomeName <value>` to specify the user-defined name of the Oracle home.

`--oracleHome <value>` to specify the full location (path) of the Oracle home.

#### Q: What is the difference between `--oracleHomeName` and `--oracleHome`?

A:

`--oracleHomeName` is the user-defined name for an Oracle home.

`--oracleHome` refers to the full directory path where the Oracle home is located.

#### Q: How do I run the `dbaascli dbHome getDetails` command?

A: To run the command, use the following syntax:

```
dbaascli dbHome getDetails --oracleHomeName <value>
```

or

```
dbaascli dbHome getDetails --oracleHome <value>
```

Make sure to run the command as the `root` user.

#### Q: Can I specify both `--oracleHomeName` and `--oracleHome` in the same command?

A: No, you can only use one option per command execution. You must either specify the Oracle home name or the Oracle home location, not both.

#### Q: What information does the `dbaascli dbHome getDetails` command return?

A: The command provides detailed information about the specified Oracle home, such as its version, status, and any other configuration details associated with the Oracle home.

**Q: What is an example of using the dbascli dbHome getDetails command with an Oracle home location?**

A: Here is an example:

```
dbascli dbHome getDetails --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1
```

This command retrieves detailed information about the Oracle home located at `/u01/app/oracle/product/19.0.0/dbhome_1`.

**Q: What is an example of using the dbascli dbHome getDetails command with an Oracle home name?**

A: Here is an example:

```
dbascli dbHome getDetails --oracleHomeName myOracleHome
```

This command retrieves detailed information about the Oracle home named `myOracleHome`.

**Q: Do I need any special permissions to run this command?**

A: Yes, you must run the command as the root user to view details about the Oracle home.

**Q: What should I do if the dbascli dbHome getDetails command returns no information?**

A: Ensure that you have correctly specified the Oracle home name or location, and that the Oracle home is properly configured and exists on the system.

**Q: Can I use the dbascli dbHome getDetails command on multiple Oracle homes simultaneously?**

A: No, the command only works on a single Oracle home at a time. You must run the command separately for each Oracle home.

**Q: Is it possible to verify the Oracle home name before running the command?**

A: Yes, you can verify the Oracle home name by checking your system's configuration files or by listing all Oracle homes available on your system.

**Q: What happens if I do not specify an Oracle home name or location in the command?**

A: The command requires either the `--oracleHome` or `--oracleHomeName` option to be specified. If neither is provided, the command will fail to execute.

**Q: Can I retrieve information about Oracle homes that are currently not in use?**

A: Yes, the `dbascli dbHome getDetails` command provides details about Oracle homes regardless of whether they are in use or idle.

**Q: What should I check if the command returns an error?**

A: Ensure that the Oracle home name or location is correct, the Oracle home exists, and that you are running the command as the `root` user. Double-check for typos or incorrect paths.

**Example 6-25 dbascli dbHome getDetails - using Oracle home location**

```
dbascli dbHome getDetails --oracleHome /u02/app/home_db19c/
```

### Example 6-26 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>]
```

### 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
- `--skipClosedPDBs` 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

### Frequently Asked Questions

#### Q: What is the dbaascli dbHome patch command used for?

A: The `dbaascli dbHome patch` command is used to patch Oracle home from one patch level to another.

#### Q: Do I need special permissions to run the dbaascli dbHome patch command?

A: Yes, you need to run the command as the `root` user.

#### Q: How do I specify the Oracle home path or name for the patch?

A: Use the `--oracleHome` option to specify the path of the Oracle home, or `--oracleHomeName` to specify the name of the Oracle home.

#### Q: How can I define the target version for the patch?

A: Use the `--targetVersion` option followed by the version number in the format `19.12.0.0.0`.

#### Q: What does the --resume option do?

A: The `--resume` option allows you to resume a previous patching session.

#### Q: How do I specify a particular session ID when resuming a patch?

A: Use the `--sessionID` option to specify the session ID of the patching session you want to resume.

#### Q: What is the --continueWithDbDowntime option used for?

A: The `--continueWithDbDowntime` option allows patching to continue even if there is database downtime, useful in environments with only one active instance.

#### Q: How can I skip patching on unreachable nodes?

A: Use the `--skipUnreachableNodes` option to skip operations on nodes that are unreachable.

#### Q: How do I patch only specific nodes in a cluster?

A: Use the `--nodes` option followed by a comma-delimited list of node names to patch a subset of nodes.

#### Q: What is the --executePrereqs option for?

A: The `--executePrereqs` option runs prerequisite checks before applying the patch.

#### Q: How can I skip running datapatch on the databases?



A: Use the `--skipDatapatch` option to skip the datapatch process during patching.

**Q: Can I specify a custom location for the database image?**

A: Yes, use the `--imageLocation` option to specify a custom location for the database image.

**Q: What does the `--skipPDBs` option do?**

A: The `--skipPDBs` option allows you to skip running datapatch on a specified comma-delimited list of pluggable databases (PDBs).

**Q: How can I skip datapatch on closed PDBs?**

A: Use the `--skipClosedPDBs` option to skip datapatch on PDBs that are closed.

**Q: What happens if I use the `--rollback` option?**

A: The `--rollback` option will revert the Oracle home to its previous state before the patch was applied.

**Q: How do I specify the Oracle home path for patching?**

A: Use the `--oracleHome` option followed by the path to the Oracle home directory.

**Q: How can I patch an Oracle home by its name rather than the path?**

A: Use the `--oracleHomeName` option followed by the name of the Oracle home.

**Q: How do I resume a patching operation if it was interrupted?**

A: Use the `--resume` option along with the `--sessionID` option to resume a specific interrupted session.

**Q: Can I continue the patching process if the database is down?**

A: Yes, use the `--continueWithDbDowntime` option to continue patching even if the database is down.

**Q: What should I do if some nodes are unreachable during the patching process?**

A: Use the `--skipUnreachableNodes` option to bypass the unreachable nodes.

**Q: How can I apply the patch to only certain nodes?**

A: Specify the nodes you want to patch using the `--nodes` option with a comma-separated list of node names.

**Q: How do I check prerequisites before applying the patch?**

A: Use the `--executePrereqs` option to run prerequisite checks before applying the patch.

**Q: What should I do if I want to avoid applying datapatch during the patching process?**

A: Use the `--skipDatapatch` option to skip the datapatch step.

**Q: How can I specify a different location for the database image used in the patching process?**

A: Use the `--imageLocation` option to specify a custom location for the image.

**Q: What if I need to skip datapatch on certain PDBs?**

A: Use the `--skipPDBs` option to skip datapatch on a specified comma-delimited list of PDBs.

**Q: Can I skip datapatch on PDBs that are not currently open?**

A: Yes, use the `--skipClosedPDBs` option to skip datapatch on closed PDBs.

**Q: What should I do if the patching fails midway?**

A: You can use the `--rollback` option to revert to the previous state or try resuming the patching process with the `--resume` option.

**Q: How can I check if all prerequisites are met before applying the patch?**

A: Run the patch command with the `--executePrereqs` option to ensure all prerequisites are satisfied.

**Q: What if the patching operation does not complete successfully and I need to retry?**

A: Use the `--resume` option to retry the patching operation from where it left off. If needed, you can specify a `--sessionID` to resume a specific session.

**Q: How can I verify if the patch was successfully applied?**

A: You can verify the patching process by checking the Oracle home version using the `opatch lsinventory` command after the patch is completed.

**Q: Can I run the patching command in a dry run mode to preview actions?**

A: No, the `dbascli dbHome patch` command does not have a dry-run feature. However, you can use the `--executePrereqs` option to run prerequisite checks before actually applying the patch.

**Q: Is it possible to apply multiple patches in one run?**

A: The `dbascli dbHome patch` command only allows for one target version at a time. You would need to run the command separately for each patch version.

**Q: How do I handle patching if the environment uses multiple Oracle homes?**

A: You can specify the Oracle home you want to patch using either the `--oracleHome` or `--oracleHomeName` options, depending on whether you're specifying the path or the name of the Oracle home.

**Q: Can I skip both PDB and CDB datapatching in one command?**

A: Yes, you can combine the `--skipPDBs` and `--skipDatapatch` options to skip datapatching for both PDBs and the CDB in a single patch run.

**Q: Can I apply a patch and rollback immediately if it causes issues?**

A: Yes, after applying a patch, you can use the `--rollback` option to revert to the previous patch level if any issues arise.

**Q: Can I patch multiple Oracle homes simultaneously?**

A: No, you need to run the `dbascli dbHome patch` command individually for each Oracle home.

**Q: How do I track the progress of the patching operation?**

A: During the patching process, the command provides output messages that show the progress. You can also check the log files for detailed information.

**Q: Can I run patching in parallel on a clustered environment?**

A: Patching operations can be applied to a subset of nodes using the `--nodes` option. However, simultaneous patching should be handled carefully, and you should ensure no overlapping sessions.

**Q: How can I identify which patches are available for my Oracle home?**

A: You can check the available patches via the Oracle support portal or by running the `opatch lsinventory` command to see the current patches applied to your Oracle home.

**Q: Can I specify a timeout for draining resources when stopping the database during patching?**

A: Yes, you can use the `--drainTimeoutInSeconds` option to specify the time in seconds for resource draining when stopping the database.

**Q: What happens if the patch fails on one of the nodes in a multi-node environment?**

A: You can use the `--skipUnreachableNodes` option to skip the failed node and continue the patching process on the remaining nodes. You can then address the issue on the failed node separately.

**Q: How can I make the patching process run in the background?**

A: Use the `--waitForCompletion` option with a value of `false` to allow the patching process to run in the background. This way, you don't need to wait for the process to complete interactively.

**Q: Can I perform a rollback operation on a subset of nodes in a clustered environment?**

A: Yes, you can use the `--nodes` option along with the `--rollback` option to roll back the patching on a specific set of nodes.

**Q: What if I need to update the image location after starting the patch process?**

A: The `--resume` option does not allow changing the image location. However, you can stop the session and start a new patch process with the updated `--imageLocation`.

**Q: Is there a way to check which session IDs are available for resuming a patch?**

A: You can check the log files or use Oracle Cloud tools to identify active or paused patching sessions and their session IDs.

**Q: Can I limit the downtime during patching?**

A: If you need to limit downtime, use the `--continueWithDbDowntime` option carefully. This allows you to proceed even when downtime is expected but requires planning for minimal service impact.

### Example Use Cases

#### Example 1: Basic Oracle Home Patching by Oracle Home Path

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0
```

Patches the Oracle home located at `/u01/app/oracle/product/19.0.0/dbhome_1` to version 19.12.0.0.0.

#### Example 2: Patching by Oracle Home Name

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0
```

Patches Oracle home named `DB_HOME_NAME` to version 19.12.0.0.0.

### Example 3: Resuming a Previous Patch Operation

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --resume
```

Resumes the previous patching operation for the Oracle home located at /u01/app/oracle/product/19.0.0/dbhome\_1.

### Example 4: Resuming a Patch with a Specific Session ID

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --resume --sessionID 12345
```

Resumes the patching operation for Oracle home DB\_HOME\_NAME using session ID 12345.

### Example 5: Patching with Database Downtime Allowed

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0 --continueWithDbDowntime
```

Patches the Oracle home located at /u01/app/oracle/product/19.0.0/dbhome\_1 to version 19.12.0.0.0 while allowing database downtime.

### Example 6: Skipping Unreachable Nodes

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0 --skipUnreachableNodes
```

Patches Oracle home to version 19.12.0.0.0 while skipping any unreachable nodes.

### Example 7: Patching a Subset of Nodes

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0 --nodes node1,node2
```

Patches Oracle home to version 19.12.0.0.0 on node1 and node2 only.

### Example 8: Running Prerequisite Checks Before Patching

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0 --executePrereqs
```

Patches Oracle home DB\_HOME\_NAME to version 19.12.0.0.0 after running prerequisite checks.

### Example 9: Skipping the Datapatch Step

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0 --skipDatapatch
```

Patches Oracle home to version 19.12.0.0.0 without running datapatch on the databases.

### Example 10: Using an Image File for Patching

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0 --imageFilePath /path/to/image/file.zip
```

Patches Oracle home to version 19.12.0.0.0 using an image file located at /path/to/image/file.zip.

### Example 11: Skipping Specific PDBs During Datapatch

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0 --skipPDBs cdb1:pdb1,cdb2:pdb2
```

Patches Oracle home to version 19.12.0.0.0 and skips running datapatch on the specified PDBs (pdb1 in cdb1 and pdb2 in cdb2).

**Example 12: Skipping Datapatch on Closed PDBs**

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0 --skipClosedPDBs
```

Patches Oracle home to version 19.12.0.0.0 while skipping running datapatch on any closed PDBs.

**Example 13: Rolling Back Oracle Home**

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --rollback
```

Rolls back the last applied patch on the Oracle home named DB\_HOME\_NAME.

**Example 14: Combination of Patching with Prerequisite Checks and Specific Nodes**

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0 --executePrereqs --nodes node1,node2
```

Patches Oracle home to version 19.12.0.0.0, runs prerequisite checks, and applies the patch only on node1 and node2.

**Example 15: Skipping Unreachable Nodes and Specific PDBs**

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0 --skipUnreachableNodes --skipPDBs cdb1:pdb1
```

Patches Oracle home DB\_HOME\_NAME to version 19.12.0.0.0 while skipping unreachable nodes and avoiding running datapatch on pdb1 within cdb1.

**Example 16: Checking Oracle Home Version Post-Patch**

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0
```

```
opatch lsinventory
```

This example shows how to check the Oracle home version after a successful patch by running `opatch lsinventory`.

**Example 17: Rolling Back Patching with Specific Session ID**

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --rollback --sessionID 67890
```

Rolls back the Oracle home patching for a session ID of 67890.

**Example 18: Patching with Skipping Prerequisite Checks**

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --targetVersion 19.12.0.0.0 --skipPrereqs
```

Patches the Oracle home but skips the prerequisite checks before applying the patch.

**Example 19: Applying a Patch to a Custom Oracle Home Image**

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0 --imageLocation /custom/location/image.zip
```

Patches Oracle home using a custom image file located at /custom/location/image.zip.

**Example 20: Skipping Specific Nodes and Running Prereqs**

```
dbascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
targetVersion 19.12.0.0.0 --skipUnreachableNodes --executePrereqs
```

Skips patching unreachable nodes and runs prerequisite checks before applying the patch.

#### Example 21: Skipping Datapatch on All PDBs in Multiple CDBs

```
dbascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
targetVersion 19.12.0.0.0 --skipPDBs cdb1:pdb1, cdb2:pdb2, cdb3:pdb3
```

Patches Oracle home but skips datapatch on the specified PDBs in multiple CDBs.

#### Example 22: Continuing Patching with Downtime on Multiple Nodes

```
dbascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0  
--continueWithDbDowntime --nodes node3, node4
```

Continues patching on node3 and node4 with database downtime allowed.

#### Example 23: Skipping Datapatch on PDBs and Closed PDBs

```
dbascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
targetVersion 19.12.0.0.0 --skipDatapatch --skipClosedPDBs
```

Patches Oracle home while skipping both the datapatch and closed PDBs.

#### Example 24: Rolling Back and Reapplying Patch

```
dbascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
rollback
```

```
dbascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
targetVersion 19.12.0.0.0
```

Rolls back the current patch and then reapplies the patch to the Oracle home.

#### Example 25: Skipping Datapatch and Allowing Downtime on a Specific Node

```
dbascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
targetVersion 19.13.0.0.0 --skipDatapatch --continueWithDbDowntime --nodes node1
```

Patches Oracle home to version 19.13.0.0.0 on node1, skipping the datapatch step and allowing downtime.

#### Example 26: Specifying Drain Timeout During Database Shutdown

```
dbascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.13.0.0.0  
--drainTimeoutInSeconds 300
```

Patches the Oracle home `DB_HOME_NAME` to version 19.13.0.0.0 and allows a 5-minute timeout (300 seconds) for draining resources during shutdown.

#### Example 27: Running Patching in the Background

```
dbascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
targetVersion 19.13.0.0.0 --waitForCompletion false
```

Patches Oracle home to version 19.13.0.0.0 and runs the patching process in the background without waiting for completion.

#### Example 28: Rolling Back Patch on a Subset of Nodes

```
dbascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --  
rollback --nodes node1, node2
```

Rolls back the last applied patch on node1 and node2 only for the specified Oracle home.

#### Example 29: Skipping Prerequisites and Patching Multiple Nodes

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.13.0.0.0 --skipPrereqs --nodes node3,node4
```

Patches Oracle home to version 19.13.0.0.0 on node3 and node4 without running prerequisite checks.

#### Example 30: Rolling Back Patch and Skipping Unreachable Nodes

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --rollback --
skipUnreachableNodes
```

Rolls back the last patch on Oracle home `DB_HOME_NAME` and skips unreachable nodes during the rollback process.

## dbaascli dbimage purge

The `dbimage purge` command removes the specified software image from your Oracle Exadata Database Service on Exascale 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
- `--dbName` 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 `dbascli dbaascli diag collect` command uses the syntax shown below to collect tooling log data:

## dbascli diag healthCheck

To run diagnostic health checks, use the `dbascli diag healthCheck` command.

**Prerequisite**

Run the command as the `root` user.

**Syntax**

```
dbascli 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*.

## Frequently Asked Questions

### Q: What is the purpose of the dbaascli diag healthCheck command?

A: The `dbaascli diag healthCheck` command is used to perform diagnostic health checks on an Oracle Database running in an Exadata Cloud@Customer environment.

### Q: What are the prerequisites for using the dbaascli diag healthCheck command?

A: The command must be run as the `root` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

### Q: How do I specify a custom directory for collecting the logs?

A: Use the `--destLocation` option to specify the directory where the health check logs will be collected. The default location is `/var/opt/oracle/dbaas_acfs`.

### Q: What is the default location for log collection if I don't specify --destLocation?

A: The default directory for log collection is `/var/opt/oracle/dbaas_acfs`.

### Q: Can I specify which nodes to run the health check on?

A: Yes, you can use the `--nodes` option to specify a comma-delimited list of nodes where the health check should be run.

### Q: How do I upload the health check logs to Object Storage?

A: Use the `--objectStoreBucketURI` option to provide a pre-authenticated request (PAR) URL from the Object Storage service. This will upload the collected logs to the specified bucket.

### Q: Can I collect logs from multiple nodes?

A: Yes, you can specify multiple nodes using the `--nodes` option in a comma-delimited format. For example: `--nodes node1,node2`.

### Q: What is an example command to run a health check on a specific node?

A: Here's an example command to run the health check on a specific node:

```
dbaascli diag healthCheck --nodes node1
```

### Q: How can I store the logs in Object Storage instead of on the local machine?

A: You can provide a pre-authenticated request (PAR) URL using the `--objectStoreBucketURI` option to store the logs in Object Storage.

### Q: Can I specify both --destLocation and --objectStoreBucketURI at the same time?

A: Yes, you can specify both `--destLocation` for local storage and `--objectStoreBucketURI` to upload logs to Object Storage.

**Q: What should I do if I encounter an error while running the dbaascli diag healthCheck command?**

A: Ensure that you are running the command as the root user and have provided valid options for `--destLocation`, `--nodes`, or `--objectStoreBucketURI`. Verify that the node names are correct if specified.

**Q: Can I run the health check in the background?**

A: The `dbaascli diag healthCheck` command does not have an explicit background mode, but you can run it in the background by appending `&` at the end of the command.

**Q: What happens if I don't provide the --nodes option?**

A: If the `--nodes` option is not provided, the health check will be performed on all nodes in the cluster by default.

**Q: Can I resume a previous health check session using this command?**

A: No, the `dbaascli diag healthCheck` command does not support resuming previous sessions. You must initiate a new health check each time.

**Q: What is an example command to run a health check and upload logs to Object Storage?**

A: Here's an example command:

```
dbaascli diag healthCheck --objectStoreBucketURI https://  
objectstorage.example.com/n/namespace-string/b/bucket-name/o/PAR-URL
```

**Q: What is the default behavior if I don't specify --destLocation or --objectStoreBucketURI?**

A: If neither `--destLocation` nor `--objectStoreBucketURI` is specified, the health check logs will be collected in the default directory `/var/opt/oracle/dbaas_acfs` on the local machine.

**Q: Can I limit the health check to specific components or logs?**

A: No, the `dbaascli diag healthCheck` command does not allow you to specify individual components or logs. It performs a general diagnostic health check for the system.

**Q: What should I verify before running the dbaascli diag healthCheck command?**

A: Ensure that you are connected to an Exadata Cloud@Customer virtual machine and running the command as the `root` user.

**Related Topics**

- [Using Pre-Authenticated Requests](#)

## 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

### Note:

By default, TCPS is enabled for databases on Oracle Exadata Database Service on Dedicated Infrastructure systems.

### 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
[--pkcs12WalletFilePath]
[--caCertChain]
[--precheckOnly]
[--serverCert]
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

### Where:

- `--pkcs12WalletFilePath` specifies the absolute path of the certificate file, which is in the `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`.

## Frequently Asked Questions

### Q: What is the purpose of the `dbaascli grid configureTCPS` command?

A: The `dbaascli grid configureTCPS` command is used to configure Transport Layer Security (TCPS) for the existing cluster in an Oracle Exadata environment.

**Q: What is the prerequisite for running the dbascli grid configureTCPS command?**

A: The command must be run as the `root` user.

**Q: Is TCPS enabled by default on Exadata Database Service on Dedicated Infrastructure systems?**

A: Yes, TCPS is enabled by default for databases on Oracle Exadata Database Service on Dedicated Infrastructure systems.

**Q: Is TCPS enabled by default on Exadata Database Service on Cloud@Customer systems?**

A: No, TCPS is not enabled by default on Exadata Database Service on Cloud@Customer systems. To enable TCPS, you need to update the `sqlnet.ora` file for the specific database and restart the database.

**Q: What is the consequence of enabling TCPS on Exadata Cloud@Customer systems?**

A: Enabling TCPS on Exadata Cloud@Customer systems will cause Zero Data Loss Recovery Appliance (ZDLRA) connections to fail. You can only enable either ZDLRA or TCPS configuration, but not both simultaneously.

**Q: What does the --pkcs12WalletFilePath option specify?**

A: The `--pkcs12WalletFilePath` option specifies the absolute path to the certificate file in the PKCS12 wallet format, which is used for TCPS configuration.

**Q: What is the --caCertChain option used for?**

A: The `--caCertChain` option specifies a concatenated list of certificates containing intermediate CA certificates and the root CA certificate.

**Q: What does the --precheckOnly option do?**

A: The `--precheckOnly` option specifies whether to run only the prechecks for the TCPS configuration operation. Accepted values are `yes` or `no`.

**Q: What does the --serverCert option specify?**

A: The `--serverCert` option specifies the path to the PEM certificate that will be used or rotated for the TCPS configuration.

**Q: How do I specify the private key for TCPS configuration?**

A: Use the `--privateKey` option to specify the path to the private key file associated with the server certificate.

**Q: What are the accepted values for the --certType option?**

A: The accepted values for the `--certType` option are:

`SELF_SIGNED_CERT`

`CA_SIGNED_CERT`

`PKCS12_CERT`

The default value is `SELF_SIGNED_CERT`.

**Q: Is the private key password protected by default?**

A: Yes, the `--privateKeyPasswordProtected` option is set to `true` by default, indicating that the private key is password protected. You can set it to `false` if the private key is not password protected.

**Q: Can I run a precheck before configuring TCPS?**

A: Yes, you can run only the prechecks for the operation by setting the `--precheckOnly` option to `yes`. This helps to validate the environment before making changes.

**Q: What happens if I provide an incorrect path for the PKCS12 wallet file?**

A: If the `--pkcs12WalletFilePath` contains an incorrect path, the command will fail, and the TCPS configuration will not proceed.

**Q: What should I do if the private key is password protected?**

A: If the private key is password protected, ensure that the `--privateKeyPasswordProtected` option is set to `true` (which is the default).

**Q: Can I specify my own CA-signed certificate for TCPS configuration?**

A: Yes, you can specify your own CA-signed certificate by using the `--serverCert` and `--privateKey` options, and by setting the `--certType` to `CA_SIGNED_CERT`.

**Q: What is an example command to configure TCPS using a self-signed certificate?**

A: Here's an example:

```
dbascli grid configureTCPS --serverCert /path/to/self_signed_cert.pem --privateKey /path/to/private_key.pem --certType SELF_SIGNED_CERT
```

**Q: Can I use a PKCS12 certificate for TCPS configuration?**

A: Yes, you can use a PKCS12 certificate by specifying the `--pkcs12WalletFilePath` option and setting the `--certType` to `PKCS12_CERT`.

**Q: What should I verify before running the dbascli grid configureTCPS command?**

A: Verify that you have the correct certificate files, private key files, and that you are logged in as the root user. Also, ensure you understand the implications if you're using ZDLRA as it cannot run simultaneously with TCPS.

**Example 6-27 dbascli grid configureTCPS**

To configure grid using self-signed certificate:

```
dbascli grid configureTCPS
```

To configure grid using CA-signed certificate:

```
dbascli 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`

### Frequently Asked Questions

**Q: What does the dbaascli grid patch command do?**

A: The `dbaascli grid patch` command is used to patch Oracle Grid Infrastructure to a specified minor version.

**Q: Do I need special permissions to run the dbaascli grid patch command?**

A: Yes, you need to run the `dbaascli grid patch` command as the `root` user.

**Q: Can I specify a target version when patching Oracle Grid Infrastructure?**

A: Yes, you can specify the target version using the `--targetVersion` option.

**Q: How do I specify the target version for the patch?**

A: Use the `--targetVersion` option followed by the version number in the format `19.12.0.0.0`.

**Q: What does the --containerURL option do in the dbaascli grid patch command?**

A: The `--containerURL` option allows you to specify a custom URL for fetching the Grid Infrastructure image.

**Q: What is the purpose of the --executePrereqs option?**

A: The `--executePrereqs` option is used to run prerequisite checks before applying the patch.

**Q: How can I patch a subset of nodes using the dbaascli grid patch command?**

A: Use the `--nodeList` option followed by a comma-delimited list of node names to patch only a subset of nodes.

**Q: What happens if I use the --rollback option?**

A: The `--rollback` option will roll back the patched Oracle home to its previous state.

**Q: Can I resume a previous patching session?**

A: Yes, you can use the `--resume` option to resume the last patching session. If you have a specific session ID, you can specify it with the `--sessionID` option.

**Q: What is the --continueWithDbDowntime option used for?**

A: The `--continueWithDbDowntime` option allows you to continue patching even if there is database downtime, typically used in environments where there is only one active instance.

**Q: How do I create an image from a patched Grid home?**

A: Use the `--createImage` option to create an image. You can specify the directory where the image should be created using the `--createImageDir` option.

**Q: What is the purpose of the --imageFile option?**

A: The `--imageFile` option allows you to specify the fully qualified path of the image file to be used for patching.

**Q: How can I run the dbaascli grid patch command in the background?**

A: You can use the `--waitForCompletion` option set to `false` to run the operation in the background.

**Q: Can I use a custom URL to fetch the patch image?**

A: Yes, you can use the `--containerURL` option to specify a custom URL for fetching the Grid Infrastructure image.

**Q: How do I specify which nodes to patch if I don't want to patch all of them?**

A: You can specify the nodes you want to patch using the `--nodeList` option with a comma-separated list of node names.

**Q: What should I do if I need to roll back a patch?**

A: Use the `--rollback` option in the `dbaascli grid patch` command to roll back the patch.

**Q: How do I handle a patching operation if my environment only has one active instance and I need to continue with downtime?**

A: Use the `--continueWithDbDowntime` option to continue patching even with database downtime.

**Q: Can I create an image of the patched Grid home?**

A: Yes, you can use the `--createImage` option to create an image of the patched Grid home. If needed, specify the directory where the image should be saved using `--createImageDir`.

**Q: What should I do if I want to resume a patching session after an interruption?**

A: Use the `--resume` option to resume the patching session. If you know the session ID, you can specify it with `--sessionID`.

**Q: What if the patching process fails midway?**

A: If the patching process fails, you can use the `--resume` option to restart the process. You can also use the `--rollback` option to revert to the previous state.

**Q: How can I ensure all prerequisites are met before patching?**

A: Use the `--executePrereqs` option to run all prerequisite checks before applying the patch.

**Q: Can I perform patching in the background to avoid holding up the terminal?**

A: Yes, you can use the `--waitForCompletion false` option to run the patching process in the background.

**Q: How can I create a Grid home image after patching?**

A: Use the `--createImage` option to create a new image from the patched Grid home. Specify the directory using `--createImageDir` if needed.

**Q: How do I use an existing image file for patching?**

A: You can use the `--imageFile` option to specify the fully qualified path to the image file you want to use for patching.

**Q: What should I do if I want to avoid database downtime during patching?**



A: Ensure that your environment has more than one active instance running. You can avoid using the `--continueWithDbDowntime` option, which is meant for environments with only one active instance.

**Q: How do I know the progress of a patch running in the background?**

A: If you run the patch with `--waitForCompletion false`, you can check the status of the background job using operating system commands like `ps` or check the logs located in the Grid home.

**Q: Is it possible to patch to a higher major version using dbaascli grid patch?**

A: No, `dbaascli grid patch` only allows patching to a minor version of the current major version. For major upgrades, you will need to follow a different upgrade process.

**Q: Can I skip specific prerequisite checks during patching?**

A: No, when you use `--executePrereqs`, all prerequisite checks will be executed. However, you can review the results of the prerequisite checks and manually handle any issues before proceeding.

**Q: What should I do if the patching process is stuck or hanging?**

A: If the patching process is unresponsive, you can stop it using operating system commands and then resume using the `--resume` option. If that doesn't work, try using the `--rollback` option to revert the patch.

**Q: Can I automate the patching process across multiple clusters?**

A: Yes, using scripts that include the `dbaascli grid patch` command with appropriate options, you can automate patching across different clusters.

**Q: Where can I find logs for the patching process?**

A: Logs are typically located in the Oracle Grid home logs directory or the default location specified during setup. You can monitor these logs for details about the patching process.

**Q: Is it possible to create a silent patch process with no user interaction?**

A: Yes, by specifying all necessary options in the command and running it in the background (`--waitForCompletion false`), you can create a non-interactive patching process.

**Q: Can I check for available patch updates before applying a patch?**

A: The `dbaascli grid patch` command itself does not have an option to list available patches. However, you can use Oracle's standard methods such as Oracle Support to identify the latest patches.

**Q: Can I use dbaascli to patch multiple Oracle Homes?**

A: No, the `dbaascli grid patch` command is designed to patch a specific Oracle Grid Infrastructure home at a time. You would need to run the command separately for each home.

**Q: Is there a way to prevent downtime completely when patching Grid Infrastructure?**

A: To minimize downtime, ensure that your environment has multiple active database instances (RAC configuration) so that patching can be done node by node. The `--continueWithDbDowntime` option should not be used in this case.

**Q: How do I handle patching for RAC One Node environments?**

A: In RAC One Node environments, you need to be cautious with the `--continueWithDbDowntime` option, as there may be only one active instance. Review the Oracle documentation for specific patching guidelines for RAC One Node.

**Q: Can I view the session history of previous patches?**

A: The dbaascli utility does not provide a direct way to view session history. However, logs of previous patching sessions can be found in the Grid home logs directory.

**Example Use Cases**

**Example 1: Basic Grid Patching**

```
dbaascli grid patch --targetVersion 19.12.0.0.0
```

Patches the Oracle Grid Infrastructure to version 19.12.0.0.0.

**Example 2: Patching with a Custom Container URL**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --containerURL https://  
example.com/custom/url
```

Patches Grid Infrastructure to version 19.12.0.0.0, using a custom container URL to fetch the Grid Infrastructure image.

**Example 3: Patching with Prerequisite Checks**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --executePrereqs
```

Patches Grid Infrastructure to version 19.12.0.0.0 after running the prerequisite checks.

**Example 4: Patching on a Subset of Nodes**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --nodeList node1,node2,node3
```

Patches Grid Infrastructure to version 19.12.0.0.0 on the specified nodes (node1, node2, and node3).

**Example 5: Rolling Back the Patch**

```
dbaascli grid patch --rollback
```

Rolls back the last applied patch on the Oracle Grid Infrastructure.

**Example 6: Resuming a Previous Patch Operation**

```
dbaascli grid patch --resume
```

Resumes the previous patching operation from where it was stopped.

**Example 7: Resuming a Patch Operation with a Specific Session ID**

```
dbaascli grid patch --resume --sessionID 12345
```

Resumes the patching operation using session ID 12345.

**Example 8: Patching with Database Downtime Allowed**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --continueWithDbDowntime
```

Patches the Grid Infrastructure to version 19.12.0.0.0 while allowing downtime of the database if needed.

**Example 9: Creating a Patched Image**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --createImage --createImageDir /  
path/to/dir
```

Creates an image of the patched Grid home (version 19.12.0.0.0) and stores it in the specified directory.

#### **Example 10: Using an Existing Image File**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --imageFile /path/to/image/  
file.zip
```

Patches Grid Infrastructure to version 19.12.0.0.0 using an existing image file located at /path/to/image/file.zip.

#### **Example 11: Running the Patching Operation in the Background**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --waitForCompletion false
```

Patches Grid Infrastructure to version 19.12.0.0.0 and runs the operation in the background.

#### **Example 12: Combination of Prerequisites, Custom URL, and Subset of Nodes**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --executePrereqs --containerURL  
https://example.com/custom/url --nodeList node1,node2
```

Patches Grid Infrastructure to version 19.12.0.0.0, runs prerequisite checks, uses a custom URL for the image, and applies the patch only on node1 and node2.

#### **Example 13: Creating a Patched Image with an Existing Image File**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --createImage --createImageDir /  
path/to/dir --imageFile /path/to/existing/image.zip
```

Creates a patched image and stores it in the specified directory while using an existing image file for the patch.

#### **Example 14: Verifying Prerequisites Without Patching**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --executePrereqs
```

Verifies whether all prerequisites are met for patching to version 19.12.0.0.0 without actually applying the patch.

#### **Example 15: Running Patch and Ignoring Prerequisite Failures**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --continueWithDbDowntime --  
executePrereqs
```

Runs the patch even if some prerequisite checks fail. This is useful in scenarios where downtime is allowed, and certain prerequisites can be ignored.

#### **Example 16: Checking Patch Logs for Issues**

```
tail -f /u01/app/grid/logs/grid_patch.log
```

Monitors the patch log in real time to diagnose any issues during the patching process.

#### **Example 17: Applying the Patch in a Parallel Environment**

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --nodeList node1,node2 --  
waitForCompletion false
```

Patches Grid Infrastructure on a subset of nodes (node1 and node2) and runs the process in the background.

### Example 18: Using a Specific Image File from an External Source

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --imageFile /mnt/images/  
grid_patch_19.12.zip
```

Patches Grid Infrastructure using a pre-downloaded image file located on an external storage device.

### Example 19: Running Patch with a Custom Session ID

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --resume --sessionID 67890
```

Resumes a patching operation that was interrupted, using session ID 67890.

### Example 20: Scheduling Patching to Run at a Later Time

```
echo "dbaascli grid patch --targetVersion 19.12.0.0.0" | at 02:00
```

Schedules the patching command to run at 2:00 AM using the at command in Linux.

### Example 21: Specifying Timeout for Completion

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --waitForCompletion true --  
continueWithDbDowntime --timeout 7200
```

Patches Grid Infrastructure while allowing downtime, but waits up to 7200 seconds (2 hours) for completion before timeout.

### Example 22: Creating a Custom Image for Another Environment

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --createImage --createImageDir /  
backups/images/grid_patch
```

Creates a custom image of the patched Grid Infrastructure to store in the `/backups/images/grid_patch` directory for use in other environments.

### Example 23: Patch Recovery After Interruption

```
dbaascli grid patch --resume --continueWithDbDowntime
```

Recovers and resumes the patching process if it was interrupted, with database downtime allowed.

### Example 24: Combining Prerequisites Check with Background Execution

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --executePrereqs --  
waitForCompletion false
```

Checks prerequisites and runs the patch in the background.

### Example 25: Skipping Image Creation for Faster Patching

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --patchInParallel --  
continueWithDbDowntime --waitForCompletion false
```

Patches Grid Infrastructure to version 19.12.0.0.0 in parallel across nodes, with database downtime allowed, and without creating an image to speed up the process.

### Example 26: Monitoring Patch Progress Through Logs

```
tail -f /u01/app/grid/logs/grid_patch_progress.log
```

Monitors the log file for patching progress in real-time, providing insights into each step of the patching process.

### Example 27: Patching with Custom Drain Timeout

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --drainTimeoutInSeconds 3600 --continueWithDbDowntime
```

Patches Grid Infrastructure and sets a custom timeout of 1 hour (3600 seconds) to allow graceful resource draining during database downtime.

### Example 28: Applying a Patch to Specific Nodes with Prerequisite Checks

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --nodeList node1,node4 --executePrereqs
```

Patches only nodes node1 and node4 to version 19.12.0.0.0 and runs the prerequisite checks beforehand.

### Example 29: Patching Without Waiting for Completion

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --waitForCompletion false
```

Begins patching the Grid Infrastructure to version 19.12.0.0.0 in the background, allowing other tasks to be performed without waiting for the process to complete.

### Example 30: Reapplying a Failed Patch After a Drain Timeout Issue

```
dbaascli grid patch --resume --drainTimeoutInSeconds 7200
```

Resumes the previous patching session and extends the resource draining timeout to 2 hours (7200 seconds) in case it failed due to insufficient time in the previous attempt.

### Example 31: Viewing Patch Logs in Real-Time with Specific Session ID

```
tail -f /u01/app/grid/logs/grid_patch_12345.log
```

Monitors the log file for the patching session with session ID 12345 in real-time.

### Example 32: Patching to a New Target Home

```
dbaascli grid patch --targetHome /u01/app/grid_home_19c --executePrereqs
```

Performs an out-of-place patch to a new Oracle Grid home located at /u01/app/grid\_home\_19c, with prerequisite checks.

### Example 33: Stopping a Background Patch Job

```
ps -ef | grep dbaascli | grep patch | awk '{print $2}' | xargs kill -9
```

Stops a background patch job by finding and killing the associated process ID (PID).

### Example 34: Verifying Patch Completion Without Logs

```
dbaascli grid status --targetVersion 19.12.0.0.0
```

Verifies if the patch to version 19.12.0.0.0 has been successfully applied by checking the current Grid Infrastructure version status.

## 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

## Frequently Asked Questions

### Q: What is the purpose of the `dbaascli grid removeTCPSCert` command?

A: The `dbaascli grid removeTCPSCert` command is used to remove existing TCPS certificates from the Grid Infrastructure wallet in an Oracle Exadata environment.

### Q: What is the prerequisite for running the `dbaascli grid removeTCPSCert` command?

A: The command must be run as the `root` user.

### Q: What does the `--subject` option specify in the `dbaascli grid removeTCPSCert` command?

A: The `--subject` option specifies the subject of the certificate to be removed from the Grid Infrastructure wallet.

### Q: What is the purpose of the `--userCert` flag?

A: The `--userCert` flag indicates that the certificate to be removed is a user certificate.

### Q: When should I use the `--trustedCert` flag?

A: Use the `--trustedCert` flag when removing a trusted certificate from the Grid Infrastructure wallet.

### Q: What does the `--requestedCert` flag do?

A: The `--requestedCert` flag indicates that the certificate being removed is a requested certificate.

### Q: What does the `--serialNumber` option specify?

A: The `--serialNumber` option specifies the serial number of the certificate to be removed. It is useful for uniquely identifying a certificate when there are multiple certificates with the same subject.

**Q: What is the purpose of the `--executePrereqs` option?**

A: The `--executePrereqs` option runs prerequisite checks before removing the certificate and reports the results, ensuring that the environment is properly prepared for the operation.

**Q: What does the `--resume` option do?**

A: The `--resume` option resumes the removal operation if it was previously interrupted.

**Q: How do I specify a session ID when resuming an interrupted operation?**

A: Use the `--sessionID` option to specify the session ID of the interrupted operation that you want to resume.

**Q: What does the `--bounceListeners` flag do?**

A: The `--bounceListeners` flag is used to restart the Grid Infrastructure listener and scan listener after the TCPS certificate is removed.

**Q: Can I remove a TCPS certificate without bouncing the listeners?**

A: Yes, the `--bounceListeners` flag is optional. If you don't specify it, the listeners will not be bounced automatically.

**Q: How can I ensure that the operation will run safely?**

A: You can use the `--executePrereqs` option to perform prerequisite checks before running the command, ensuring that everything is in order before the removal process.

**Q: What should I do if I need to remove a specific user certificate by serial number?**

A: Use the `--subject` option to specify the certificate's subject, the `--userCert` flag to indicate that it is a user certificate, and the `--serialNumber` option to specify the certificate's serial number.

**Q: Can I remove multiple certificates at once?**

A: No, the command is designed to remove a single certificate at a time based on the provided subject and other parameters.

**Q: What happens if the certificate removal process is interrupted?**

A: You can resume the operation using the `--resume` option along with the `--sessionID` of the interrupted process.

**Q: Do I need to run the command as the root user?**

A: Yes, the `dbaascli grid removeTCPSCert` command must be run as the `root` user to have the necessary privileges for removing TCPS certificates.

**Q: How can I identify the certificate I want to remove?**

A: You can identify the certificate by its subject, and optionally, by its serial number to ensure you're targeting the correct certificate for removal.

**Q: What is an example command to remove a trusted certificate?**

A: Here's an example:

```
dbaascli grid removeTCPSCert --subject "CN=example_cert" --trustedCert
```

## 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
[--pkcs12WalletFilePath]
[--caCertChain]
[--precheckOnly]
[--serverCert]
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

### Where:

- `--pkcs12WalletFilePath` specifies the absolute path of the certificate file, which is in the `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`.

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli grid rotateTCPSCert` command?

A: The `dbaascli grid rotateTCPSCert` command is used to rotate TCPS (Transport Layer Security Protocol) certificates in the Grid Infrastructure wallet in Oracle Exadata environments.

#### Q: What is the prerequisite for running the `dbaascli grid rotateTCPSCert` command?

A: The command must be run as the `root` user.

#### Q: What does the `--pkcs12WalletFilePath` option specify?

A: The `--pkcs12WalletFilePath` option specifies the absolute path to the certificate file in the PKCS12 wallet format for the TCPS configuration.

#### Q: What is the purpose of the `--caCertChain` option?

A: The `--caCertChain` option specifies a concatenated list of certificates, including intermediate CA and root CA certificates, for the TCPS configuration.



**Q: What does the --precheckOnly option do?**

A: The --precheckOnly option allows you to run prechecks without making any actual changes. The valid values are "yes" to run only the prechecks and "no" to proceed with the rotation.

**Q: How is the --serverCert option used?**

A: The --serverCert option specifies the path to the PEM (Privacy Enhanced Mail) server certificate that is used or rotated for the TCPS configuration.

**Q: What does the --privateKey option specify?**

A: The --privateKey option specifies the path to the private key file corresponding to the server certificate used for TCPS rotation.

**Q: What are the valid values for the --certType option?**

A: The --certType option accepts the following values for specifying the type of certificate to be added to the Grid Infrastructure wallet:

SELF\_SIGNED\_CERT (default)

CA\_SIGNED\_CERT

PKCS12\_CERT

**Q: What does the --privateKeyPasswordProtected option do?**

A: The --privateKeyPasswordProtected option indicates whether the private key is password-protected. Valid values are true (default) and false

**Q: Can I run the dbascli grid rotateTCPSCert command without rotating the certificates?**

A: Yes, by using the --precheckOnly yes option, you can run only the prechecks without rotating the certificates.

**Q: What is an example of a command to rotate a certificate using a PKCS12 wallet?**

A: Here's an example command:

```
dbascli grid rotateTCPSCert --pkcs12WalletFilePath /path/to/wallet.p12 --certType PKCS12_CERT
```

**Q: How do I rotate a server certificate with a CA-signed certificate chain?**

A: Use the --serverCert and --caCertChain options as shown below:

```
dbascli grid rotateTCPSCert --serverCert /path/to/serverCert.pem --caCertChain /path/to/caChain.pem
```

**Q: What happens if I don't specify --privateKeyPasswordProtected?**

A: If you don't specify the --privateKeyPasswordProtected option, the command assumes that the private key is password-protected (default: true).

**Q: Can I rotate a self-signed certificate?**

A: Yes, you can rotate a self-signed certificate by using the default --certType SELF\_SIGNED\_CERT option or specifying it explicitly.

**Q: How can I rotate a certificate without providing a private key?**

A: For certain certificate types, like PKCS12, you may not need to provide a separate private key file, as it is included in the wallet. However, if a private key is required, it must be provided using the `--privateKey` option.

**Q: What if I want to rotate a certificate in the background?**

A: The `dbaascli grid rotateTCPSCert` command doesn't provide an explicit option for background execution. You may run the command directly in a background session (e.g., using `nohup` or similar tools).

**Q: What is the default certificate type if not specified?**

A: The default certificate type is `SELF_SIGNED_CERT`.

**Example 6-28 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

## Frequently Asked Questions

### Q: What is the purpose of the `dbaascli grid upgrade` command?

A: The `dbaascli grid upgrade` command is used to upgrade Oracle Grid Infrastructure from one major version to another on an Exadata Cloud@Customer virtual machine.

### Q: What is the prerequisite for running the `dbaascli grid upgrade` command?

A: The command must be run as the `root` user, and you need to be connected to an Exadata Cloud@Customer virtual machine.

### Q: What does the `--version` option specify?

A: The `--version` option specifies the target version of Oracle Grid Infrastructure that you want to upgrade to.

### Q: What does the `--resume` option do?

A: The `--resume` option resumes a previously interrupted or failed Grid Infrastructure upgrade process.

### Q: How is the `--executePrereqs` option used?

A: The `--executePrereqs` option runs only the prerequisite checks for the Grid Infrastructure upgrade without performing the actual upgrade.

### Q: What is the purpose of the `--containerURL` option?

A: The `--containerURL` option specifies a custom URL to fetch the Grid Infrastructure software image for the upgrade.

### Q: What does the `--softwareOnly` option do?

A: The `--softwareOnly` option installs only the Grid Infrastructure software without configuring or upgrading the Grid environment.

### Q: When would you use the `--targetHome` option?

A: The `--targetHome` option specifies the path to the existing target Grid home where the upgrade will be performed.

### Q: What happens if the upgrade fails?

A: If the upgrade fails, you can use the `--revert` option to roll back the upgrade to its previous state.

### Q: Can I perform a Grid Infrastructure upgrade in stages?

A: Yes, by using the `--softwareOnly` option, you can first install the software and then later complete the full upgrade, allowing for staged upgrades.

### Q: How do I use the `dbaascli grid upgrade` command to upgrade only the software?

A: Use the following syntax to upgrade only the software:

```
dbaascli grid upgrade --version <target_version> --softwareOnly
```

### Q: Can I check for upgrade prerequisites without performing the upgrade?

A: Yes, you can run only the prerequisite checks using:

```
dbaascli grid upgrade --version <target_version> --executePrereqs
```

**Q: How can I upgrade Grid Infrastructure using a custom container URL?**

A: You can specify the URL for fetching the Grid Infrastructure image as follows:

```
dbaascli grid upgrade --version <target_version> --containerURL <custom_url>
```

**Q: How can I resume a previously interrupted upgrade process?**

A: To resume a previously interrupted or failed upgrade, use:

```
dbaascli grid upgrade --version <target_version> --resume
```

**Q: What does the --revert option do in the dbaascli grid upgrade command?**

A: The `--revert` option rolls back a failed or interrupted Grid Infrastructure upgrade to its original state.

**Q: Can I perform a full upgrade without configuring the Grid Infrastructure immediately?**

A: Yes, you can first install only the software using the `--softwareOnly` option and then configure it later.

**Q: What should I do if an upgrade fails and I want to undo the changes?**

A: Use the `--revert` option to roll back the failed upgrade:

```
dbaascli grid upgrade --version <target_version> --revert
```

**Example 6-29 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 6-30 dbaascli job getStatus**

```
dbaascli 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,
"updatedAt" : 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*

## dbascli patch db prereq

### 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*

## 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.]
```

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli pdb backup` command?

A: The `dbaascli pdb backup` command is used to create backups for a pluggable database (PDB), query backup status, generate backup reports, and delete PDB backups in an Exadata Cloud@Customer environment.

#### Q: What is the prerequisite for using the `dbaascli pdb backup` command?

A: The command must be run as the `root` user, and you need to be connected to an Exadata Cloud@Customer virtual machine.

**Q: How do I start a PDB backup using the dbaascli pdb backup command?**

A: You can start a PDB backup using the `--start` option. For example:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --start
```

**Q: What options can be used with the --start flag?**

A: With the `--start` flag, you can specify:

`--level1` for a Level-1 incremental backup

`--archival` for a full archival backup (which also requires a `--tag` to specify the backup tag)

**Q: How do I create a Level-1 incremental PDB backup?**

A: Use the `--level1` flag with the `--start` option to create a Level-1 incremental backup:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --start --level1
```

**Q: How do I create an archival PDB backup?**

A: Use the `--archival` flag with the `--start` option and specify a backup tag using `--tag`:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --start --archival --tag <backup_tag>
```

**Q: How do I delete a specific PDB backup?**

A: To delete a specific backup, use the `--delete` flag and specify the backup tag using `--backupTag`:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --delete --backupTag <backup_tag>
```

**Q: How can I check the status of a PDB backup?**

A: Use the `--status` flag along with the `backup --uuid` to check the status of a specific backup:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --status --uuid <backup_uuid>
```

**Q: How do I retrieve a PDB backup report in JSON format?**

A: To get a backup report in JSON format, use the `--getBackupReport` option, specify the file name with `--json`, and provide the backup tag with `--tag`:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --getBackupReport --json <file_name> --tag <backup_tag>
```

**Q: How can I list all PDB backups for a specific PDB?**

A: Use the `--list` option to get a list of all backups for a given PDB:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --list
```

Optionally, you can output the list in JSON format using the `--json` flag:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --list --json <file_name>
```

**Q: What does the --pdbName option do?**

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to back up, query, or delete backups for.

**Q: What is the purpose of the `--dbname` option?**

A: The `--dbname` option specifies the name of the Oracle Database to which the PDB belongs.

**Q: How do I specify a backup tag for a PDB backup?**

A: You specify a backup tag using the `--tag` option when starting an archival backup or when retrieving a backup report:

```
--tag <backup_tag>
```

**Q: Can I run PDB backups in JSON mode?**

A: Yes, both the backup report (`--getBackupReport`) and backup listing (`--list`) options support output in JSON format. You specify a JSON file name using the `--json` option.

**Example 6-31 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 Oracle Exadata Database Service on Exascale 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 --dbname --pdbName | --pdbUID  
[-openMode]
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbName` specifies the name of the PDB
- `--pdbUID` specifies the identifier of the PDB
- `--openMode` specifies the target `OPEN MODE` of PDB



## Frequently Asked Questions

### Q: What is the purpose of the dbaascli pdb bounce command?

A: The `dbaascli pdb bounce` command is used to bounce (restart) a pluggable database (PDB) in an Exadata Cloud@Customer environment.

### Q: What are the prerequisites for using the dbaascli pdb bounce command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

### Q: How do I bounce a PDB by specifying its name?

A: To bounce a PDB by specifying its name, use the following syntax:

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbName <PDB_Name>
```

### Q: How do I bounce a PDB by using its unique identifier (UID)?

A: To bounce a PDB using its unique identifier (UID), use the following syntax:

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbUID <PDB_UID>
```

### Q: What is the --dbname option used for?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the pluggable database (PDB) being bounced.

### Q: What is the --pdbName option used for?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to bounce.

### Q: What is the --pdbUID option used for?

A: The `--pdbUID` option specifies the unique identifier (UID) of the pluggable database (PDB) that you want to bounce.

### Q: How do I specify the target open mode for the PDB when bouncing it?

A: You can use the `--openMode` option to specify the desired open mode for the PDB after bouncing. The valid values are `READ_WRITE` and `READ_ONLY`. For example:

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbName <PDB_Name> --openMode  
READ_WRITE
```

### Q: Can I open the PDB in read-only mode after bouncing it?

A: Yes, you can use the `--openMode READ_ONLY` option to open the PDB in read-only mode after bouncing:

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbName <PDB_Name> --openMode READ_ONLY
```

### Q: What is the default open mode if --openMode is not specified?

A: If `--openMode` is not specified, the PDB will be opened in its default open mode, which is typically `READ_WRITE`.

### Q: Can I use both --pdbName and --pdbUID in the same command?

A: No, you should specify either `--pdbName` or `--pdbUID`, but not both in the same command.

**Q: How can I restart a PDB and ensure it opens in read-write mode?**

A: To restart a PDB and ensure it opens in read-write mode, use the `--openMode READ_WRITE` option:

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbName <PDB_Name> --openMode  
READ_WRITE
```

**Q: Is it mandatory to specify the open mode when using the dbaascli pdb bounce command?**

A: No, specifying the `--openMode` is optional. If not provided, the PDB will open in its default mode.

**Q: What happens if I don't specify the --openMode flag?**

A: If the `--openMode` flag is not specified, the PDB will be opened in its default mode, which is typically `READ_WRITE`.

**Example 6-32 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 --dbname --pdbName | --pdbUID
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB.
- `--pdbname` specifies the name of the PDB that you want to close.
- `--pdbUID` specifies the identifier of the PDB

Upon successful completion of running this command, the PDB is closed on all of the container database instances.

**Frequently Asked Questions**

**Q: What is the purpose of the dbaascli pdb close command?**

A: The `dbaascli pdb close` command is used to close a pluggable database (PDB) in an Exadata Cloud@Customer environment.

**Q: What are the prerequisites for using the dbaascli pdb close command?**

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

**Q: How do I close a PDB by specifying its name?**

A: To close a PDB by specifying its name, use the following syntax:

```
dbaascli pdb close --dbname <CDB_Name> --pdbName <PDB_Name>
```

**Q: How do I close a PDB by specifying its unique identifier (UID)?**

A: To close a PDB by using its unique identifier (UID), use the following syntax:

```
dbaascli pdb close --dbname <CDB_Name> --pdbUID <PDB_UID>
```

**Q: What does the --dbname option do in the dbaascli pdb close command?**

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the pluggable database (PDB) you want to close.

**Q: What does the --pdbName option do in the dbaascli pdb close command?**

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to close.

**Q: What is the purpose of the --pdbUID option in the dbaascli pdb close command?**

A: The `--pdbUID` option allows you to specify the unique identifier (UID) of the pluggable database (PDB) that you want to close.

**Q: Can I close the PDB on a specific instance of the CDB?**

A: No, upon successful completion, the PDB is closed on all instances of the container database (CDB).

**Q: Is it possible to specify both --pdbName and --pdbUID in the same command?**

A: No, you can specify either `--pdbName` or `--pdbUID`, but not both in the same command.

**Q: What happens when the dbaascli pdb close command completes successfully?**

A: When the command completes successfully, the pluggable database (PDB) is closed on all instances of the container database (CDB).

**Q: How do I close a specific PDB within a CDB using its UID?**

A: You can close a specific PDB using its UID by running:

```
dbaascli pdb close --dbname <CDB_Name> --pdbUID <PDB_UID>
```

**Q: What happens if I forget to specify either --pdbName or --pdbUID?**

A: You must specify either `--pdbName` or `--pdbUID` in the command. If neither is provided, the command will not run.

**Q: Can I use the dbaascli pdb close command for a CDB directly?**

A: No, the command is designed to close a pluggable database (PDB) within a container database (CDB), not the CDB itself.

### Example 6-33 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 getConnectionString

To display Oracle Net connect string information for a pluggable database (PDB) run the `dbaascli pdb getConnectionString` command.

### Prerequisite

Run the command as the `oracle` user.

### Syntax

```
dbaascli pdb getConnectionString --dbname --pdbName | --pdbUID
```

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

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli pdb getConnectionString` command?

A: The `dbaascli pdb getConnectionString` command is used to display the Oracle Net connect string information for a pluggable database (PDB) in an Exadata Cloud@Customer environment.

#### Q: What are the prerequisites for using the `dbaascli pdb getConnectionString` command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

#### Q: How do I retrieve the connect string of a PDB by specifying its name?

A: To retrieve the connect string by specifying the PDB name, use the following syntax:

```
dbaascli pdb getConnectionString --dbname <CDB_Name> --pdbName <PDB_Name>
```

#### Q: How do I retrieve the connect string of a PDB by specifying its unique identifier (UID)?

A: To retrieve the connect string using the PDB's unique identifier (UID), use the following syntax:

```
dbaascli pdb getConnectionString --dbname <CDB_Name> --pdbUID <PDB_UID>
```

#### Q: What does the `--dbname` option do in the `dbaascli pdb getConnectionString` command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the pluggable database (PDB) for which you want to display the connect string information.

**Q: What does the `--pdbName` option do in the `dbascli pdb getConnectString` command?**

A: The `--pdbName` option specifies the name of the pluggable database (PDB) for which you want to retrieve the Oracle Net connect string information.

**Q: What is the purpose of the `--pdbUID` option in the `dbascli pdb getConnectString` command?**

A: The `--pdbUID` option allows you to specify the unique identifier (UID) of the pluggable database (PDB) for which you want to display the connect string.

**Q: Can I use both `--pdbName` and `--pdbUID` in the same command?**

A: No, you can use either `--pdbName` or `--pdbUID`, but not both in the same command.

**Q: What type of information is returned by the `dbascli pdb getConnectString` command?**

A: The command returns the Oracle Net connect string information for the specified pluggable database (PDB).

**Q: Can I retrieve the connect string for a PDB on a specific container database instance?**

A: No, the connect string is for the PDB as a whole, not for a specific instance of the container database.

**Q: How can I get the connect string information if I only know the PDB's unique identifier (UID)?**

A: You can retrieve the connect string using the PDB's UID by running:

```
dbascli pdb getConnectString --dbname <CDB_Name> --pdbUID <PDB_UID>
```

**Q: What happens if I don't provide either `--pdbName` or `--pdbUID`?**

A: You must specify either `--pdbName` or `--pdbUID` to retrieve the connect string. The command will not run without one of these options.

**Q: Is the connect string information for the PDB always the same across all instances of the CDB?**

A: Yes, the connect string information is consistent for the PDB across all instances of the container database (CDB).

#### **Example 6-34 dbascli pdb getConnectString**

```
dbascli pdb getConnectString --dbname dbname --pdbName pdbName
```

## dbascli pdb create

To create a new pluggable database (PDB), use the `dbascli pdb create` command.

### **Prerequisite**

Run the command as the `oracle` user.

## Syntax

```
dbascli pdb create --pdbName <value> --dbName <value>
[--maxCPU <value>]
[--maxSize <value>]
[--pdbAdminUserName <value>]
[--lockPDBAdminAccount <value>]
[--resume [--sessionID <value>]]
[--executePrereqs <value>]
[--waitForCompletion <value>]
[--blobLocation |--standbyBlobFromPrimary <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.

During the PDB creation process, you are prompted to specify the administration password for the new PDB.

## Frequently Asked Questions

### Q: What is the purpose of the dbascli pdb create command?

A: The `dbascli pdb create` command is used to create a new pluggable database (PDB) in a container database (CDB) in an Exadata Cloud@Customer environment.

### Q: What are the prerequisites for using the dbascli pdb create command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

### Q: What does the --pdbName option do in the dbascli pdb create command?

A: The `--pdbName` option specifies the name of the new pluggable database (PDB) you want to create.

### Q: What does the --dbName option do in the dbascli pdb create command?

A: The `--dbName` option specifies the name of the container database (CDB) that will host the new pluggable database (PDB).

### Q: Can I limit the CPU resources for the new PDB?

A: Yes, you can use the `--maxCPU` option to specify the maximum number of CPUs that the PDB can use. This is equivalent to setting the `CPU_COUNT` parameter in the PDB.

### Q: How can I limit the storage size of a PDB?

A: You can use the `--maxSize` option to specify the maximum total size of data files and temporary files for the PDB. You can either set a size limit (in K, M, G, or T) or specify `UNLIMITED` for no limit.

### Q: What is the --pdbAdminUserName option used for?

A: The `--pdbAdminUserName` option specifies the name of the admin user for the new PDB. This user will have administrative privileges within the PDB.

### Q: Is it possible to lock the admin user account during PDB creation?

A: Yes, you can use the `--lockPDBAdminAccount` option to specify whether the PDB admin account should be locked. The default value is `true` (locked).

### Q: What does the --resume option do in the dbascli pdb create command?

A: The `--resume` option allows you to resume a previously failed PDB creation process.

### Q: How do I specify a session ID for resuming a previous run?

A: You can specify a session ID using the `--sessionID` option to resume a specific session of the PDB creation process.

### Q: What is the purpose of the --executePrereqs option?

A: The `--executePrereqs` option specifies whether to run only the prerequisite checks for PDB creation. You can set this option to `yes` or `no`.

### Q: Can I run the PDB creation process in the background?

A: Yes, you can use the `--waitForCompletion` option and set it to `false` to run the operation in the background.

**Q: What is the `--standbyBlobFromPrimary` option used for?**

A: The `--standbyBlobFromPrimary` option specifies the location of the standby blob file, which is prepared from the primary database. This is required for standby database PDB operations.

**Q: How will I be prompted for the PDB admin password during the creation process?**

A: During the PDB creation process, you will be prompted to specify the administration password for the new PDB.

**Q: Can I create a standby PDB using the `dbaascli pdb create` command?**

A: Yes, if you are creating a standby PDB, you can use the `--standbyBlobFromPrimary` option to specify the location of the standby blob file from the primary database.

**Q: What happens if I don't use the `--maxSize` option?**

A: If you do not specify the `--maxSize` option, the PDB will not have a storage size limit unless otherwise defined by the CDB policies.

**Q: What happens if I do not provide the `--pdbAdminUserName` option?**

A: If you do not provide the `--pdbAdminUserName` option, the PDB will be created without a specified admin user, and you will need to manually configure the admin user after creation.

**Q: Can I resume a failed PDB creation at any point in the process?**

A: Yes, as long as the session has not been terminated, you can resume a failed PDB creation using the `--resume` and `--sessionID` options.

**Example 6-35 dbaascli pdb create**

To create a PDB from seed in a standard database in a non-Data Guard environment:

```
dbaascli 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 createSnapshot

To create a snapshot of a given pluggable database (PDB), use the `dbaascli pdb createSnapshot` command.

### Prerequisite

Run the command as the `oracle` user.

### Syntax

```
dbaascli pdb createSnapshot  
{  
  --pdbName <value> | --pdbUID <value>  
}
```



```
--dbName <value>
--snapshotName <value>
[--pdbAdminUserName <value>]
[--executePrereqs] [--resume [--sessionID <value>]]
[--waitForCompletion true|false]
```

**Where:**

- `--pdbName` specifies the PDB name from which to create the snapshot.
- `--pdbUID` specifies the user ID (UID) of the PDB from which to create the snapshot.
- `--dbName` Oracle Database name.
- `--snapshotName` specifies the PDB snapshot name.
- `--pdbAdminUserName` specifies the PDB administrator user name.
- `--executePrereqs` runs the prerequisite checks and reports the results.
- `--resume [--sessionID <value>]` resumes the previous operation. It can take the flag `--sessionID <value>` to specify to resume from a specific session ID (<value>).
- `--waitForCompletion true|false` specifies whether to run the operation in foreground (true) or background (false). Valid values: true, false.

**Example 6-36 dbascli pdb createSnapshot**

In the following example, a snapshot is created from the database named `db721`, in the PDB name `pdb1`. The snapshot name that is given is `snap1`.

```
dbascli pdb createSnapshot --dbName db721 --pdbName pdb1 --snapshotName snap1
```

## dbascli pdb configureSnapshot

To configure automatic snapshots for a given pluggable database (PDB), use the `dbascli pdb configureSnapshot` command.

**Prerequisite**

Run the command as the `oracle` user.

**Syntax**

```
dbascli pdb configureSnapshot
{
  --pdbName <value> | --pdbUID <value>
}
--dbName <value>
--snapshotIntervalInMins <value>
[--executePrereqs]
[--maxPDBSnapshots <value>]
[--waitForCompletion <value>]
```

**Where:**

- `--pdbName <value>` specifies the name of the PDB for which automatic snapshot configuration will be set.

- `--pdbUID <value>` specifies the user ID (UID) of the PDB for which automatic snapshot configuration will be set.
- `--dbName` Oracle Database name.
- `--snapshotIntervalInMins <value>` specifies the interval, in minutes, for when automatic PDB snapshots will be taken.
- `--executePrereqs` runs the prerequisite checks and reports the results.
- `--maxPDBSnapshots <value>]` specifies the maximum number of snapshots to create for the given PDB. .
- `--waitForCompletion true|false` specifies whether to run the operation in foreground (true) or background (false). Valid values: true, false.

### Example 6-37 dbascli pdb configureSnapshot

In the following example, an automatic snapshot plan is configured on the database named db721, in the PDB name pdb1. The snapshot interval is set to run automatic snapshot creation every 60 minutes.

```
dbascli pdb configureSnapshot --dbName db721 --pdbName pdb1 --
snapshotIntervalInMins 60
```

## dbascli pdb delete

To delete a pluggable database (PDB) run the `dbascli pdb delete` command.

### Prerequisite

Run the command as the `oracle` user.

### Syntax

```
dbascli pdb delete --dbName value
{ --pdbName value | --pdbUID value }
[--executePrereqs value]
[--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` - option to cleanup source database after a PDB has been relocated.

### Frequently Asked Questions

#### **Q: What is the purpose of the dbascli pdb delete command?**

A: The `dbascli pdb delete` command is used to delete a pluggable database (PDB) from a container database (CDB) in an Exadata Cloud@Customer environment.

#### **Q: What are the prerequisites for running the dbascli pdb delete command?**

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

#### **Q: What does the --dbName option specify in the dbascli pdb delete command?**

A: The `--dbName` option specifies the name of the container database (CDB) that hosts the PDB you want to delete.

#### **Q: How can I specify which PDB to delete using the dbascli pdb delete command?**

A: You can specify the PDB to delete using either the `--pdbName` option (specifies the PDB name) or the `--pdbUID` option (specifies the PDB UID).

#### **Q: Can I run the prerequisite checks without actually deleting the PDB?**

A: Yes, you can use the `--executePrereqs` option and set it to `yes` to run only the prerequisite checks for the PDB deletion operation.

#### **Q: How can I run the PDB deletion process in the background?**

A: Use the `--waitForCompletion` option and set it to `false` to run the deletion process in the background.

#### **Q: What does the --resume option do in the dbascli pdb delete command?**

A: The `--resume` option allows you to resume a previously failed PDB deletion process.

#### **Q: How do I resume a specific session for a PDB deletion?**

A: You can specify a session ID using the `--sessionID` option to resume a specific session for the PDB deletion process.

#### **Q: What does the --allStandbyPrepared option do?**

A: The `--allStandbyPrepared` option is used to confirm that the deletion operation has been successfully run on all standby databases before proceeding with the primary PDB deletion.

#### **Q: What is the purpose of the --cleanupRelocatedPDB option?**

A: The `--cleanupRelocatedPDB` option cleans up the source database after a PDB has been relocated, ensuring no residuals are left after the relocation.

#### **Q: Can I delete a PDB that has already been relocated?**

A: Yes, you can use the `--cleanupRelocatedPDB` option to delete a PDB that has already been relocated to a new CDB.

#### **Q: How do I ensure that the delete operation runs successfully on standby databases?**

A: Use the `--allStandbyPrepared` option to confirm that the operation has run successfully on all standby databases before proceeding.

**Q: What happens if the delete process fails and needs to be resumed?**

A: You can resume the delete process by using the `--resume` option, and if needed, specify the session ID with `--sessionID`.

**Q: What does setting `--waitForCompletion` to false do?**

A: Setting `--waitForCompletion` to `false` allows the delete process to run in the background, letting you continue working without waiting for the operation to finish.

**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 deleteSnapshot

To delete a snapshot of a given pluggable database (PDB), use the `dbaascli pdb deleteSnapshot` command.

**Prerequisite**

Run the command as the `oracle` user.

**Syntax**

```
dbaascli pdb deleteSnapshot
{
  --pdbName <value> | --pdbUID <value>
}
{
  --snapshotName <value> | --snapshotUID <value>
}
--dbName <value>
[--executePrereqs]
[--waitForCompletion <value>]
[--resume [--sessionID <value>]]
]
```

Where:

- `--pdbName <value>` specifies the name of the PDB for which automatic snapshot configuration will be set.
- `--pdbUID <value>` specifies the user ID (UID) of the PDB for which automatic snapshot configuration will be set.
- `--snapshotName <value>` specifies the name of the PDB snapshot that you want to delete.

- `--snapshotUID <value>` specifies the user ID (UID) of the PDB snapshot that you want to delete.
- `--dbName` specifies the Oracle Database name.
- `--executePrereqs` runs the prerequisite checks and reports the results.
- `--waitForCompletion true|false` specifies whether to run the operation in foreground (`true`) or background (`false`). Valid values: `true`, `false`.
- `--resume [sessionID <value>]` specifies to resume the previous operation. To specify resuming from a particular session ID, add the flag `sessionID`, and provide the session ID number.

### Example 6-38 dbaascli pdb configureSnapshot

In the following example, the PDB snapshot `snap1` is specified for deletion in the PDB named `pdb1`, for the database named `db721`:

```
dbaascli pdb deleteSnapshot --dbName db721 --pdbName pdb1 --snapshotName snap1
```

## 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 --pdbName | --pdbUID
```

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

### Frequently Asked Questions

#### Q: What is the purpose of the dbaascli pdb getDetails command?

A: The `dbaascli pdb getDetails` command is used to view details of a pluggable database (PDB) hosted in a container database (CDB) in an Exadata Cloud@Customer environment.

#### Q: What are the prerequisites for running the dbaascli pdb getDetails command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

#### Q: What does the --dbname option specify in the dbaascli pdb getDetails command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the PDB for which you want to view details.

#### Q: How do you specify the PDB for which you want to view details?

A: You can specify the PDB using either the `--pdbName` option (to provide the PDB name) or the `--pdbUID` option (to provide the PDB UID).

**Q: What is the difference between --pdbName and --pdbUID?**

A: The --pdbName option uses the name of the PDB to fetch details, whereas the --pdbUID option uses the unique identifier (UID) of the PDB to fetch its details.

**Q: Can I use both --pdbName and --pdbUID together in the dbaascli pdb getDetails command?**

A: No, you can specify either the --pdbName or the --pdbUID option to get details of the PDB, but not both at the same time.

**Q: What are some use cases for the dbaascli pdb getDetails command?**

A: You can use the dbaascli pdb getDetails command to:

- Retrieve details about a specific PDB in a CDB.
- Verify the configuration of a PDB.
- Check the status of a PDB within a CDB.

**Q: How can I view details of a PDB based on its name?**

A: To view details of a PDB based on its name, use the following syntax:

```
dbaascli pdb getDetails --dbname <CDB_Name> --pdbName <PDB_Name>
```

**Q: How can I view details of a PDB based on its UID?**

A: To view details of a PDB based on its UID, use the following syntax:

```
dbaascli pdb getDetails --dbname <CDB_Name> --pdbUID <PDB_UID>
```

**Q: Can this command be used for multiple PDBs in one execution?**

A: No, the command can be used to fetch details of one PDB at a time by specifying either its name or UID.

**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 getSnapshot

To obtain details of a given pluggable database (PDB) snapshot, use the dbaascli pdb getSnapshot command.

### Prerequisite

Run the command as the oracle user.

### Syntax

```
dbaascli pdb getSnapshot  
{  
  --pdbName <value>| --pdbUID <value>  
}
```

```
{
  --snapshotName <value> | --snapshotUID <value>
}
--dbName <value>
```

Where:

- `--pdbName <value>` specifies the name of the PDB for which you want to obtain details.
- `--pdbUID <value>` specifies the user ID (UID) of the PDB for the snapshot for which you want to obtain details.
- `--snapshotName <value>` specifies the name of the snapshot for which you want to obtain details
- `--snapshotUID <value>` specifies the user ID (UID) of the snapshot for which you want to obtain details.
- `--dbName` specifies the Oracle Database name.

#### Example 6-40 dbaascli pdb configureSnapshot

In the following example, the details are obtained for the snapshot named `snap1` in the database named `db721`, in the PDB name `pdb1`:

```
dbaascli pdb getSnapshot --dbName db721 --pdbName pdb1 --snapshotName snap1
```

## 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

### Frequently Asked Questions

#### Q: What is the purpose of the dbaascli pdb list command?

A: The `dbaascli pdb list` command is used to view the list of pluggable databases (PDBs) in a specified container database (CDB) in an Exadata Cloud@Customer environment.

#### Q: What are the prerequisites for running the dbaascli pdb list command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

#### Q: What does the --dbname option specify in the dbaascli pdb list command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the pluggable databases (PDBs) for which you want to view the list.

**Q: Can I view the list of PDBs from multiple container databases at once?**

A: No, the `dbaascli pdb list` command allows you to list the PDBs from only one container database (CDB) at a time, specified by the `--dbname` option.

**Q: How do I list the PDBs in a specific container database (CDB)?**

A: You can list the PDBs in a specific CDB by using the following syntax:

```
dbaascli pdb list --dbname <CDB_Name>
```

**Q: What information is displayed when using the dbaascli pdb list command?**

A: The command returns a list of all pluggable databases (PDBs) within the specified container database (CDB). The list typically includes the names of the PDBs and possibly other details like their status.

**Q: Can I filter the PDB list using additional options?**

A: No, the `dbaascli pdb list` command does not support additional filtering options. It simply returns the complete list of PDBs within the specified CDB.

**Q: What happens if the specified --dbname does not exist or is incorrect?**

A: If the specified `--dbname` is incorrect or does not exist, the command will return an error, and no PDB list will be displayed.

**Q: Can the dbaascli pdb list command be used for any Oracle database environment?**

A: No, the `dbaascli pdb list` command is specifically designed for use in Exadata Cloud@Customer environments.

**Example 6-41 dbaascli pdb list**

```
dbaascli pdb list --dbname cdb name
```

## dbaascli pdb listSnapshots

To list the snapshots of a given pluggable database (PDB), use the `dbaascli pdb listSnapshots` command..

### Prerequisite

Run the command as the `oracle` user.

### Syntax

```
dbaascli pdb listSnapshots  
{  
  --pdbName <value> | --pdbUID <value>  
}  
--dbName <value>
```

Where:

- `--pdbName <value>` specifies the PDB name for which the snapshots will be listed.
- `--pdbUID <value>` specifies the UID of the PDB for which the snapshots will be listed.
- `--dbName <value>` specifies the Oracle Database name.



### Example 6-42 dbascli pdb listSnapshots

In the following example, the command lists the snapshots for database db721, and the pdb name pdb1:

```
dbascli pdb listSnapshots --dbName db721 --pdbName pdb1
```

## 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.

### Prerequisite

Run the command as the `oracle` user.

### Syntax

```
dbascli 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>]
```

### 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 comma separated list of source to target service names which need to be converted. Syntax is `source_srv1:new_srv1,source_srv2:new_srv2`.

The newly cloned PDB inherits administration passwords from the source PDB.

**Frequently Asked Questions**

**Q: What is the purpose of the dbascli pdb localClone command?**

A: The `dbascli pdb localClone` command is used to create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB) in an Exadata Cloud@Customer environment.

**Q: What are the prerequisites for running the dbascli pdb localClone command?**

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine. Additionally, the source PDB must already exist in the specified CDB.

**Q: What does the --dbName option specify in the dbascli pdb localClone command?**

A: The `--dbName` option specifies the name of the container database (CDB) that hosts the source PDB from which the new PDB will be cloned.

**Q: What does the --pdbName option specify in the dbascli pdb localClone command?**

A: The `--pdbName` option specifies the name of the new PDB that you want to create as a clone of the existing PDB in the same CDB.

**Q: Can I clone a PDB with a different name using the dbascli pdb localClone command?**

A: Yes, you can specify a different name for the cloned PDB using the `--targetPDBName` option. If this option is not provided, the cloned PDB will inherit the name of the source PDB.

**Q: What does the --resume option do in the dbascli pdb localClone command?**

A: The `--resume` option allows you to resume a previously interrupted PDB cloning operation.

**Q: How do I limit the CPU resources available to the cloned PDB?**

A: You can limit the CPU resources for the cloned PDB using the `--maxCPU` option, which specifies the maximum number of CPUs that will be allocated to the new PDB.

**Q: Can I run the PDB cloning operation in the background?**

A: Yes, you can run the operation in the background by setting the `--waitForCompletion` option to `false`. If you set it to `true`, the operation will run in the foreground and wait for completion.

**Q: What is the purpose of the `--maxSize` option in the `dbascli pdb localClone` command?**

A: The `--maxSize` option specifies the maximum storage size (in GB) for the newly cloned PDB. If no size is specified, the cloned PDB inherits the same storage limits as the source PDB.

**Q: Can I control the parallelism of the PDB clone operation?**

A: Yes, you can control the degree of parallelism for the cloning operation using the `--powerLimit` option. This option accepts values between 1 and 128 to define the degree of parallelism.

**Q: What is the `--primaryDBWalletTar` option used for?**

A: The `--primaryDBWalletTar` option specifies the location of the primary database wallet tar file. This option is only required if the cloning operation involves standby database PDB operations.

**Q: Can I run only the prerequisite checks for the cloning operation?**

A: Yes, you can run only the prerequisite checks by using the `--executePrereqs` option and setting it to `yes`. The valid values are `yes` and `no`.

**Q: What happens if the PDB cloning operation fails or is interrupted?**

A: If the cloning operation fails or is interrupted, you can resume it by using the `--resume` option to continue from where the operation stopped.

**Example 6-43 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>]]
```

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

Upon successful completion, the PDB is opened on all of the container database instances.

### Frequently Asked Questions

**Q: What is the purpose of the dbascli pdb open command?**

A: The `dbascli pdb open` command is used to open a pluggable database (PDB) in an Oracle container database (CDB) in an Exadata Cloud@Customer environment.

**Q: Who can run the dbascli pdb open command?**

A: The command can be run as either the `root` or `oracle` user.

**Q: What does the --pdbName option specify in the dbascli pdb open command?**

A: The `--pdbName` option specifies the name of the PDB that you want to open.

**Q: What does the --pdbUID option specify in the dbascli pdb open command?**

A: The `--pdbUID` option specifies the unique identifier (UID) of the PDB that you want to open.

**Q: What does the --dbname option specify in the dbascli pdb open command?**

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the PDB.

**Q: What is the purpose of the --openMode option?**

A: The `--openMode` option specifies the mode in which the PDB will be opened. Valid values are `READ_WRITE` and `READ_ONLY`.

**Q: Can I start services when opening the PDB?**

A: Yes, you can use the `--startServices` option to either start all services associated with the PDB by specifying `all` or provide a comma-delimited list of specific services to start.

**Q: What happens if I set the --waitForCompletion option to false?**

A: If `--waitForCompletion` is set to `false`, the command will run in the background, and the user does not need to wait for the operation to complete. If set to `true`, the command will wait for completion before exiting.

**Q: What does the --setPDBRefreshModeNone option do?**

A: The `--setPDBRefreshModeNone` option converts a refreshable PDB (one that is regularly updated from a primary database) into a non-refreshable PDB.

**Q: What is the function of the `--skipPDBRefresh` option?**

A: The `--skipPDBRefresh` option allows you to skip the refresh operation when opening a refreshable PDB, preventing the PDB from syncing with the primary database at that time.

**Q: What does the `--pdbAdminUserName` option do in the `dbascli pdb open` command?**

A: The `--pdbAdminUserName` option allows you to specify a new PDB admin username when opening the PDB.

**Q: What happens if the `dbascli pdb open` command is successful?**

A: Upon successful completion, the specified PDB will be opened on all instances of the container database (CDB).

**Q: Is it possible to run the `dbascli pdb open` command for a refreshable PDB?**

A: Yes, the command can be used for refreshable PDBs. The `--setPDBRefreshModeNone` option converts the PDB to non-refreshable, and the `--skipPDBRefresh` option skips the refresh operation during the opening process.

**Q: What is the default open mode for a PDB if no `--openMode` is specified?**

A: If no `--openMode` is specified, the PDB is typically opened in `READ_WRITE` mode by default.

**Example 6-44 dbascli pdb open**

```
dbascli pdb open --dbname cdb name --pdbName pdb name associated with the CDB
```

```
dbascli pdb open --dbname cdb name --pdbUID con_uid of that pdb
```

**Optional:** `--openMode READ_WRITE/READ_ONLY`

## dbascli pdb recover

To recover a pluggable database (PDB), use the `dbascli 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

```
dbascli pdb recover --pdbName <value> --dbname <value>
{
    --start
    {
        --untilTime <value>
        | --untilSCN <value>
        | --latest
        | --tag <value>
    }
}
```

```
    | --status --uuid <value>
  }
```

#### Where:

```
--pdbName: PDB name.
--dbname: Oracle Database name.
--start | --status
--start
    --untilTime | --untilSCN | --latest | --tag
    --untilTime: Recovers PDB until time. Input format: DD-MON-YYYY HH24:MI:SS.
    --untilSCN: Recovers PDB until SCN.
    --latest: Recovers PDB to last known state.
    --tag: Recovers PDB to archival tag.
--status
    --uuid <value>
```

### Frequently Asked Questions

#### Q: What is the purpose of the dbascli pdb recover command?

A: The `dbascli pdb recover` command is used to recover a pluggable database (PDB) to a previous state using backups stored in a configured backup storage destination.

#### Q: Who can run the dbascli pdb recover command?

A: The command must be run as the `root` user.

#### Q: What is required before running the dbascli pdb recover command?

A: Before running the command, the database must be configured with backup storage destination details where backups are stored.

#### Q: What does the --pdbName option specify in the dbascli pdb recover command?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to recover.

#### Q: What does the --dbname option specify in the dbascli pdb recover command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the PDB.

#### Q: What are the possible options for starting a PDB recovery using the --start option?

A: You can recover the PDB using one of the following options:

- `--untilTime <value>`: Recovers the PDB until a specified time (format: DD-MON-YYYY HH24:MI).
- `--untilSCN <value>`: Recovers the PDB until a specified System Change Number (SCN).
- `--latest`: Recovers the PDB to the latest known state.
- `--tag <value>`: Recovers the PDB to a specific archival tag.

#### Q: What is the format required for specifying the time in the --untilTime option?

A: The time must be in the format DD-MON-YYYY HH24:MI:SS.

#### Q: How can I recover a PDB to the latest state using dbascli pdb recover?

A: To recover the PDB to the latest known state, use the `--latest` option:

```
dbaascli pdb recover --pdbName <value> --dbname <value> --start --latest
```

**Q: How do I recover a PDB to a specific archival tag?**

A: You can recover the PDB to a specific tag using the `--tag` option:

```
dbaascli pdb recover --pdbName <value> --dbname <value> --start --tag <tag_value>
```

**Q: Can I recover a PDB using a specific SCN?**

A: Yes, you can recover the PDB to a specific SCN using the `--untilSCN` option:

```
dbaascli pdb recover --pdbName <value> --dbname <value> --start --untilSCN  
<SCN_value>
```

**Q: What does the `--status` option do in the `dbaascli pdb recover` command?**

A: The `--status` option is used to check the status of a recovery operation. You must provide the `--uuid` to specify the recovery session.

**Q: How can I check the status of a PDB recovery?**

A: To check the status of a recovery operation, use the `--status` option with the `--uuid` of the recovery session:

```
dbaascli pdb recover --pdbName <value> --dbname <value> --status --uuid  
<uuid_value>
```

**Q: What happens if I specify the `--latest` option in the recovery command?**

A: If you specify the `--latest` option, the PDB will be recovered to the most recent state available in the backup.

**Example 6-45 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 Oracle Exadata Database Service on Exascale 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`

## Frequently Asked Questions

### Q: What is the purpose of the dbaascli pdb refresh command?

A: The `dbaascli pdb refresh` command is used to refresh a specified pluggable database (PDB) in a container database (CDB).

### Q: Who can run the dbaascli pdb refresh command?

A: The command can be run by either the `root` or `oracle` user.

### Q: What does the --dbname option specify in the dbaascli pdb refresh command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the pluggable database (PDB) to be refreshed.

### Q: What does the --pdbName option specify in the dbaascli pdb refresh command?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to refresh.

### Q: What does the --pdbUID option specify in the dbaascli pdb refresh command?

A: The `--pdbUID` option specifies the unique identifier (UID) of the pluggable database (PDB) that you want to refresh.

### Q: What does the --waitForCompletion option do in the dbaascli pdb refresh command?

A: The `--waitForCompletion` option specifies whether the operation should be run in the foreground or background. If set to `true`, the operation will run in the foreground and wait for completion. If set to `false`, the operation will run in the background.

### Q: How can I refresh a PDB and run the operation in the background?

A: To refresh a PDB and run the operation in the background, use the `--waitForCompletion false` option:

```
dbaascli pdb refresh --dbname <value> --pdbName <value> --waitForCompletion false
```

### Q: How do I refresh a PDB using its unique identifier (UID)?

A: You can refresh the PDB using the `--pdbUID` option:



```
dbascli pdb refresh --dbname <value> --pdbUID <value>
```

**Q: Can I specify both --pdbName and --pdbUID together in the dbascli pdb refresh command?**

A: No, you must specify either --pdbName or --pdbUID, but not both, when refreshing a PDB.

**Q: What happens if I don't include the --waitForCompletion option in the command?**

A: If you don't specify the --waitForCompletion option, the default behavior will be to wait for the operation to complete before returning control to the user.

**Q: Can I refresh a PDB while the database is running?**

A: Yes, you can refresh a PDB while the database is running, as long as the command is executed by a user with appropriate privileges.

**Related Topics**

- [Connecting to a Virtual Machine with SSH](#)  
You can connect to the virtual machines in an Oracle Exadata Database Service on Exascale Infrastructure system by using a Secure Shell (SSH) connection.

## dbascli pdb relocate

To relocate the specified PDB from the remote database into local database, use the `dbascli 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**

```
dbascli pdb relocate --pdbName <value> --dbName <value> --
sourceDBConnectionString <value>
[--targetPDBName <value>]
[--powerLimit <value>]
[--maxCpu <value>]
[--maxSize <value>]
[--resume [--sessionID <value>]]
[--executePrereqs <value>]
[--sourcePDBServices <value>]
[--sourcePDBReadOnlyServices <value>]
[--waitForCompletion <value>]
{
    [--blobLocation <value>] | [--standbyBlobFromPrimary <value>]
}
[--upgradePDB <value>]
[--updateDBBlockCacheSize]
{
    [skipOpenPDB] | [--completePDBRelocate]
}
```

**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
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`
- `--blobLocation` specifies the location of a custom directory where the standby BLOB file will be generated in a Data Guard 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` specifies `true` to upgrade the PDB as part of this operation. Valid values : `true | false`.
- `--updateDBBlockCachesize` option to enable application to set db block cache size initialization parameters to support data copy with different block size.
- `--skipOpenPDB` - indicates that the PDB should not be opened at the end of the current operation.
- `--completePDBRelocate` - completes the PDB relocation if done as a two-step operation.

### Frequently Asked Questions

**Q: What is the dbaascli pdb relocate command used for?**

A: The `dbaascli pdb relocate` command is used to relocate a Pluggable Database (PDB) from a remote database to a local database.

**Q: What user should run the dbaascli pdb relocate command?**

A: The command should be run as the `Oracle` user.

**Q: What is required when prompted during the dbaascli pdb relocate operation?**

A: You must supply the SYS user password for the source database.

**Q: What does the --pdbName parameter specify?**

A: The --pdbName parameter specifies the name of the source PDB to be relocated.

**Q: What is the purpose of the --dbName parameter?**

A: The --dbName parameter specifies the target database name where the PDB will be relocated.

**Q: How should the --sourceDBConnectionString be formatted?**

A: The --sourceDBConnectionString should be formatted as `<scan_name>:<scan_port>/<database_service_name>`.

**Q: What does the --targetPDBName parameter do?**

A: The --targetPDBName parameter specifies a new name for the relocated PDB.

**Q: What is the use of --powerLimit?**

A: The --powerLimit parameter specifies the degree of parallelism to be used during the relocate operation.

**Q: How does --maxCpu affect the relocation process?**

A: The --maxCpu parameter specifies the maximum number of CPUs to be allocated for the PDB relocation process.

**Q: What does the --maxSize parameter define?**

A: The --maxSize parameter defines the maximum storage size in GB for the new PDB.

**Q: What is the function of --resume?**

A: The --resume parameter indicates that the relocation operation should resume from where it left off.

**Q: What should I provide with the --resume option?**

A: You can specify a --sessionId to resume a specific session if you are resuming a previous operation.

**Q: What does the --executePrereqs parameter do?**

A: The --executePrereqs parameter determines if only the prerequisites for the operation should be run. Valid values are yes or no.

**Q: What is specified by the --sourcePDBServices parameter?**

A: The --sourcePDBServices parameter specifies a list of comma-delimited source PDB services.

**Q: What does the --sourcePDBReadOnlyServices parameter list?**

A: The --sourcePDBReadOnlyServices parameter lists a comma-delimited list of source PDB read-only services.

**Q: What is the effect of --waitForCompletion?**

A: The --waitForCompletion parameter specifies whether to run the operation in the background. Valid values are true or false.

**Q: What does the --blobLocation parameter specify?**

A: The `--blobLocation` parameter specifies the location of a custom directory where the standby BLOB file will be generated in a Data Guard environment.

**Q: When should I use `--standbyBlobFromPrimary`?**

A: Use `--standbyBlobFromPrimary` to specify the location of the standby BLOB file, which is prepared from the primary database. This is required only for standby operations.

**Q: Can I use `--blobLocation` and `--standbyBlobFromPrimary` together?**

A: No, the `--blobLocation` and `--standbyBlobFromPrimary` parameters are mutually exclusive and cannot be used together.

**Q: What does `--upgradePDB` do?**

A: The `--upgradePDB` parameter specifies whether to upgrade the PDB as part of the relocation operation. Valid values are `true` or `false`.

**Q: What is the purpose of `--updateDBBlockCacheSize`?**

A: The `--updateDBBlockCacheSize` option allows the application to set the DB block cache size initialization parameters to support data copy with a different block size.

**Q: What does the `--skipOpenPDB` option do?**

A: The `--skipOpenPDB` option indicates that the PDB should not be opened at the end of the relocation operation.

**Q: When should I use `--completePDBRelocate`?**

A: Use `--completePDBRelocate` to complete the PDB relocation if it is done as a two-step operation.

**Q: What should I do if I encounter an error while using the `dbascli pdb relocate` command?**

A: Check the error message for details, ensure all parameters are correctly specified, and verify that you have the necessary permissions and credentials. You might also need to review the prerequisites and configurations.

**Q: What if I forget the SYS user password for the source database?**

A: You will need to reset or recover the SYS user password for the source database. Without it, you cannot complete the relocation operation.

**Example 6-46 dbascli pdb relocate**

```
dbascli pdb relocate --sourceDBConnectionString test-  
scan.dbaastools1rgsu.dbaastools1rgvc.oraclevcn.com:1521/  
source_cdb_service_name --pdbName source_pdb --dbName target_db
```

## 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.

Run the command as the `root` or `oracle` user.

## Syntax

```
dbascli 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>]
[--refreshablePDB --refreshMode <value> [--refreshIntervalInMinutes <value>]
--dblinkUsername <value> [--honorCaseSensitiveUserName]
[--updateDBBlockCacheSize]
```

## 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`.
- `--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
- `--updateDBBlockCacheSize`: specifies to enable application to set db block cache size initialization parameters to support data copy with a different block size

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.

### Frequently Asked Questions

**Q: What is the dbascli pdb remoteClone command used for?**

A: The `dbascli pdb remoteClone` command is used to create a new Pluggable Database (PDB) as a clone of an existing PDB in another container database (CDB).

**Q: What user should execute the dbascli pdb remoteClone command?**

A: The command should be executed as either the `root` or `oracle` user.

**Q: What is required when prompted during the dbascli pdb remoteClone operation?**

A: You must supply the SYS user password for the source PDB.

**Q: What does the --pdbName parameter specify?**

A: The `--pdbName` parameter specifies the name of the source PDB that you want to clone.

**Q: What does the --dbName parameter represent?**

A: The `--dbName` parameter represents the name (`DB_NAME`) of the CDB that will host the newly cloned PDB.

**Q: How should the `--sourceDBConnectionString` be formatted?**

A: The `--sourceDBConnectionString` should be formatted as `<scan_name>:<scan_port>/<database_service_name>`.

**Q: What is the purpose of the `--targetPDBName` parameter?**

A: The `--targetPDBName` parameter specifies the name for the newly cloned PDB.

**Q: What does `--powerLimit` control?**

A: The `--powerLimit` parameter controls the degree of parallelism used for the cloning operation. The valid value is between 1 and 128.

**Q: What does the `--maxCPU` parameter define?**

A: The `--maxCPU` parameter defines the maximum number of CPUs to be allocated for the PDB cloning process.

**Q: What is the function of `--maxSize`?**

A: The `--maxSize` parameter specifies the maximum storage size in GB for the new PDB.

**Q: What does the `--resume` parameter do?**

A: The `--resume` parameter resumes the previous cloning operation.

**Q: What should you provide with the `--resume` option?**

A: You can specify a `--sessionID` to resume a specific session if you are resuming a previous operation.

**Q: What does `--executePrereqs` control?**

A: The `--executePrereqs` parameter determines if only the prerequisites for the cloning operation should be run. Valid values are `yes` or `no`.

**Q: How does `--waitForCompletion` affect the operation?**

A: The `--waitForCompletion` parameter specifies whether to wait for the operation to complete or run it in the background. Valid values are `true` or `false`.

**Q: What is specified by the `--sourcePDBExportedTDEKeyFile` parameter?**

A: The `--sourcePDBExportedTDEKeyFile` parameter specifies the exported key file from the source PDB. This parameter is applicable only for Oracle Database version 12.1.

**Q: What does the `--blobLocation` parameter define?**

A: The `--blobLocation` parameter specifies the custom path where the standby BLOB file will be generated in a Data Guard environment.

**Q: When is `--standbyBlobFromPrimary` used?**

A: The `--standbyBlobFromPrimary` parameter specifies the location of the standby BLOB file prepared from the primary database. This is required only for standby database PDB operations.

**Q: Can `--blobLocation` and `--standbyBlobFromPrimary` be used together?**

A: No, `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive and cannot be used together.

**Q: What does the `--excludeUserTablespaces` option do?**

A: The `--excludeUserTablespaces` option allows you to skip specific user tablespaces from being cloned. For example, `t1,t2,t3`.

**Q: What is the effect of `--excludePDBData`?**

A: The `--excludePDBData` option specifies whether to skip user data from the source PDB during cloning. Valid values are `true` or `yes`.

**Q: What is specified by `--pdbAdminUserName`?**

A: The `--pdbAdminUserName` parameter specifies the new admin user name for the cloned PDB.

**Q: What does the `--lockPDBAdminAccount` option control?**

A: The `--lockPDBAdminAccount` option specifies whether to lock the PDB admin user account. The default value is `true`.

**Q: What does `--sourcePDBServiceConvertList` specify?**

A: The `--sourcePDBServiceConvertList` parameter specifies a comma-delimited list of source to target service name conversions. For example, `source_srv1:new_srv1,source_srv2:new_srv2`.

**Q: What is the purpose of `--refreshablePDB`?**

A: The `--refreshablePDB` parameter specifies whether to create a refreshable PDB.

**Q: What does `--refreshMode` control?**

A: The `--refreshMode` parameter controls the refresh mode for a refreshable PDB. Valid values are `AUTO` or `MANUAL`.

**Q: How does `--refreshIntervalInMinutes` work?**

A: The `--refreshIntervalInMinutes` parameter specifies the interval in minutes for refreshing the refreshable PDB.

**Q: What is `--dblinkUsername` used for?**

A: The `--dblinkUsername` parameter specifies a common user of a remote database used for the database link to connect to the remote database.

**Q: What does the `--honorCaseSensitiveUserName` option indicate?**

A: The `--honorCaseSensitiveUserName` option indicates that the specified username is case sensitive.

**Q: What is the effect of `--updateDBBlockCacheSize`?**

A: The `--updateDBBlockCacheSize` option enables the application to set the DB block cache size initialization parameters to support data copy with a different block size.

**Q: What should I do if I encounter an error with the `dbascli pdb remoteClone` command?**

A: Review the error message for details, ensure all parameters are correctly specified, and verify that you have the necessary permissions and credentials. Additionally, check that the source and target databases meet all the requirements.

**Q: What if I forget the SYS user password for the source PDB?**



A: You will need to reset or recover the SYS user password for the source PDB. Without it, the cloning operation cannot be completed.

#### Example 6-47 dbaascli pdb remoteClone

```
dbaascli pdb remoteClone --sourceDBConnectionString test-
can.dbaastoolslrsgsu.dbaastoolslrsvc.oraclevcn.com:1521 --pdbName source_pdb1
--dbName db9944 --targetPDBName new_pdb1 --maxsize 5 --maxcpu 2
```

```
dbaascli pdb remoteClone --sourceDBConnectionString
orcla.dbaastoolslrsgsu.dbaastoolslrsvc.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
```

### Frequently Asked Questions

#### Q: What is the dbaascli system getDBHomes command used for?

A: The `dbaascli system getDBHomes` command is used to view information about all the Oracle homes on a system.

#### Q: What user should execute the dbaascli system getDBHomes command?

A: The command should be executed as either the `root` or `oracle` user.

#### Q: Are there any parameters for the dbaascli system getDBHomes command?

A: No, the `dbaascli system getDBHomes` command does not have any parameters.

#### Q: What kind of information does the dbaascli system getDBHomes command provide?

A: The command provides details about all Oracle homes on the system, including their paths and other relevant information.

#### Q: How can I interpret the output from the dbaascli system getDBHomes command?

A: The output will list all Oracle homes with information such as the location of each Oracle home. This information can help in managing and configuring Oracle environments.

#### Q: What should I do if the dbaascli system getDBHomes command does not return any output?

A: Ensure that you are running the command as the `root` or `oracle` user and verify that Oracle homes are properly installed on the system. You may also want to check system permissions and configurations.

**Q: What if I receive an error message while executing the dbaascli system getDBHomes command?**

A: Check the error message for specific details, verify that you have the appropriate permissions, and ensure that the dbaascli tool is correctly installed and configured.

**Q: Can I run dbaascli system getDBHomes on a non-Oracle system?**

A: No, the dbaascli system getDBHomes command is specific to Oracle systems and requires Oracle software to be installed.

**Example 6-48 dbaascli system getDBHomes**

```
dbaascli system getDBHomes
```

## 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

```
dbaascli 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.

### Frequently Asked Questions

**Q: What does the dbaascli tde changePassword command do?**

A: The dbaascli tde changePassword command changes the Transparent Data Encryption (TDE) keystore password as well as the database wallet password for the alias tde\_ks\_passwd.

**Q: Who should run the dbaascli tde changePassword command?**

A: The command must be run as the root user.

**Q: When should I use the dbascli tde changePassword command?**

A: Use this command when you need to change the TDE keystore password or the DB wallet password for an Exadata Cloud@Customer database.

**Q: What does the --dbname option do?**

A: The --dbname option specifies the name of the database for which you want to change the TDE keystore password.

**Q: What does the --pdbName option do?**

A: The --pdbName option specifies the name of the pluggable database (PDB) for which the TDE keystore password needs to be changed. This option is used for multitenant databases.

**Q: Can you give an example of how to run this command for a specific database?**

A: Here's an example to change the TDE keystore password for a specific database:

```
dbascli tde changePassword --dbname mydatabase
```

**Q: How do I run the command for a specific PDB within a multitenant database?**

A: You can specify the PDB name using this syntax:

```
dbascli tde changePassword --dbname mydatabase --pdbName mypdb
```

**Q: What are the prerequisites for running the dbascli tde changePassword command?**

A: You must run the command as the root user and have access to the Exadata Cloud@Customer virtual machine where the database is running.

**Q: Do I need to stop the database to change the TDE keystore password?**

A: No, the database does not need to be stopped to change the TDE keystore password.

**Q: What should I do if the command fails?**

A: Ensure that you are running the command as the root user and that the database name (--dbname) and PDB name (--pdbName, if applicable) are correct.

**Q: What if I get an "invalid password" error when changing the TDE keystore password?**

A: Make sure the new password meets your system's password complexity requirements, and that you are entering the correct old password if prompted.

**Q: How do I check if the TDE keystore password has been changed successfully?**

A: You can check the database logs or use the Oracle Database Vault and Key Management views to verify that the TDE keystore password change was successful.

**Q: Can I change the TDE keystore password for a multitenant database and all PDBs at once?**

A: No, the dbascli tde changePassword command needs to be run for each PDB individually if you need to change the password for multiple PDBs.

**Q: What happens if I forget the new TDE keystore password?**

A: If the new password is forgotten, you may need to restore the keystore from a backup or follow Oracle's recovery process to reset it, depending on your setup.

**Q: Can I automate the process of changing the TDE keystore password?**

A: While the `dbaascli tde changePassword` command itself is not designed for automation, you can script it as part of your regular database maintenance procedures if needed.

**Q: How often should I change the TDE keystore password?**

A: Oracle recommends periodically changing your TDE keystore password based on your organization's security policies. Best practices typically involve rotating encryption keys and keystore passwords regularly.

```
dbaascli tde changepassword --dbname
    <dbname>
```

1. Change the TDE password in primary database.

```
dbaascli 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

```
dbaascli tde changepassword --dbname
    <dbname> --standbyBlobFromPrimary <Location of blob generated from
    primary>
```

## dbaascli tde addSecondaryHsmKey

To add a secondary HSM (KMS) key to the existing HSM (KMS) configuration, use the `dbaascli tde addSecondaryHsmKey` command.

### Prerequisite

Run the command as the `root` user.

### Syntax

```
dbaascli 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` sexecute the prerequisites checks and report the results.

### Frequently Asked Questions

**Q: What does the `dbaascli tde addSecondaryHsmKey` command do?**

A: The `dbaascli tde addSecondaryHsmKey` command adds a secondary HSM (KMS) key to the existing HSM (KMS) configuration for an Exadata Cloud@Customer database.

**Q: Who should run the `dbaascli tde addSecondaryHsmKey` command?**

A: The command must be run as the `root` user.

**Q: On which machine should I run the `dbaascli tde addSecondaryHsmKey` command?**

A: You need to connect to an Exadata Cloud@Customer virtual machine using SSH to run this command.

**Q: Where can I find more details about connecting to a virtual machine to run this command?**

A: You can refer to the guide "Connecting to a Virtual Machine with SSH" for instructions on how to connect.

**Q: What does the `--secondaryKmsKeyOCID` option specify?**

A: The `--secondaryKmsKeyOCID` option specifies the OCID (Oracle Cloud Identifier) of the secondary KMS key to be added to the existing HSM (KMS) configuration.

**Q: What does the `--dbname` option do?**

A: The `--dbname` option allows you to specify the name of the database for which the secondary KMS key should be added. It is optional.

**Q: What does the `--precheckOnly` option do?**

A: The `--precheckOnly` option, when set to `yes`, runs a precheck of the operation without making any actual changes. The valid values are `yes` or `no`.

**Q: Can I run the precheck only without making changes?**

A: Yes, you can use the `--precheckOnly yes` option to run just the precheck without making changes.

**Q: Can you give an example of how to run this command to add a secondary HSM key?**

A: Here's an example:

```
dbaascli tde addSecondaryHsmKey --secondaryKmsKeyOCID ocid1.kms.key.oc1..example
```

**Q: How do I run the command for a specific database?**

A: You can specify the database name like this:

```
dbaascli tde addSecondaryHsmKey --secondaryKmsKeyOCID ocid1.kms.key.oc1..example  
--dbname mydatabase
```

**Q: How do I run the command with a precheck only?**

A: To run the precheck, use the following syntax:

```
dbaascli tde addSecondaryHsmKey --secondaryKmsKeyOCID ocid1.kms.key.oc1..example  
--precheckOnly yes
```

**Q: What should I do if the command fails?**

A: Ensure that you are running the command as the `root` user and that you have connected to the correct Exadata Cloud@Customer virtual machine. Also, verify the OCID of the KMS key and check if the required permissions are granted.

**Q: How can I check if I have the correct OCID for the secondary KMS key?**

A: You can retrieve the OCID of the KMS key from the Oracle Cloud Infrastructure console, under the Key Management Service (KMS) section.

**Q: What permissions are required to add a secondary KMS key?**

A: You need appropriate permissions in Oracle Cloud Infrastructure for KMS operations, including the ability to manage KMS keys for the relevant compartment.

**Q: Can I use the dbaascli tde addSecondaryHsmKey command without specifying the --dbname option?**

A: Yes, the --dbname option is optional. If omitted, the command applies to all databases using the existing HSM (KMS) configuration.

**Q: What happens if I add a secondary KMS key?**

A: The secondary KMS key will be added to the existing configuration, providing an additional layer of encryption key management redundancy.

**Q: Can I remove a secondary KMS key once it is added?**

A: No, once a secondary KMS key is added, it cannot be removed. You can only rotate or update keys in the future.

**Example 6-49 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` spfile 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.

### Frequently Asked Questions

#### Q: What does the `dbaascli tde enableWalletRoot` command do?

A: The `dbaascli tde enableWalletRoot` command enables the `wallet_root` parameter in the `spfile` for an existing Oracle database on Exadata Cloud@Customer.

#### Q: Who should run the `dbaascli tde enableWalletRoot` command?

A: The command must be run as the `root` user.

#### Q: On which machine should I run the `dbaascli tde enableWalletRoot` command?

A: You must connect to an Exadata Cloud@Customer virtual machine using SSH to run this command.

#### Q: Where can I find instructions to connect to the virtual machine?

A: You can refer to the guide "Connecting to a Virtual Machine with SSH" for instructions on connecting.

#### Q: What does the `--dbRestart` option do?

A: The `--dbRestart` option specifies how the database should be restarted after enabling `wallet_root`. The valid values are:

- `rolling`: Restarts the database in a rolling manner (default behavior).
- `full`: Performs a full database restart.

#### Q: What does the `--dbname` option do?

A: The `--dbname` option allows you to specify the name of the Oracle Database for which the `wallet_root` parameter should be enabled.

#### Q: What does the `--precheckOnly` option do?

A: The `--precheckOnly` option runs a precheck of the operation without making actual changes. The valid values are `yes` or `no`.

#### Q: What happens if I do not specify the `--dbRestart` option?

A: If you do not specify the `--dbRestart` option, the database will restart in a rolling manner by default.

#### Q: Can you give an example of how to enable `wallet_root` for a specific database?

A: Here's an example to enable `wallet_root` for a database named `mydatabase`:

```
dbaascli tde enableWalletRoot --dbname mydatabase
```

#### Q: How do I enable `wallet_root` and specify a full database restart?

A: You can enable `wallet_root` with a full database restart using the following command:

```
dbaascli tde enableWalletRoot --dbname mydatabase --dbRestart full
```

#### Q: How do I run the command with a precheck only?

A: To perform a precheck without making changes, use the following syntax:

```
dbaascli tde enableWalletRoot --dbname mydatabase --precheckOnly yes
```

**Q: What are the prerequisites for running the dbaascli tde enableWalletRoot command?**

A: You must run the command as the `root` user and be connected to the correct Exadata Cloud@Customer virtual machine.

**Q: Do I need to restart the database to enable wallet\_root?**

A: Yes, the database will need to restart either in a rolling manner (default) or fully, depending on the option you choose.

**Q: What should I do if the command fails?**

A: Ensure that you are running the command as the `root` user, and verify that the database name (`--dbname`) is correct. Check for any precheck errors if you are running with `--precheckOnly`.

**Q: What if the database fails to restart after running the command?**

A: Verify that the correct restart option was used (`rolling` or `full`) and check the database logs for any errors. You may need to manually restart the database if the automatic restart fails.

**Q: How can I check if wallet\_root was enabled successfully?**

A: You can verify the change by checking the database's `spfile` or using Oracle SQL queries to confirm that the `wallet_root` parameter is enabled.

**Q: Can I enable wallet\_root without restarting the database?**

A: No, the database needs to restart for the change to take effect. You can choose between a rolling restart or a full restart.

**Q: What is the difference between a rolling and full database restart?**

A: A rolling restart restarts the database one instance at a time, allowing the database to remain partially available during the operation. A full restart shuts down and restarts the entire database, causing a complete downtime.

**Q: Can I run this command for multiple databases simultaneously?**

A: You need to run the `dbaascli tde enableWalletRoot` command separately for each database you wish to enable `wallet_root` on.

**Q: How does enabling wallet\_root affect the existing TDE keystore configuration?**

A: Enabling `wallet_root` updates the TDE keystore location to the new wallet root directory, making it easier to manage multiple keystores and wallets in Oracle databases.

**Example 6-50 dbaascli tde enableWalletRoot**

```
dbaascli tde enableWalletRoot --dbname db name --dbrestart rolling|full
```

```
dbaascli tde enableWalletRoot --dbname orcl
```

```
dbaascli tde enableWalletRoot --dbname orcl--dbrestart full
```



## dbaascli tde encryptTablespacesInPDB

To encrypt all the tablespaces in the specified PDB, use the `dbaascli tde encryptTablespacesInPDB` command.

### Prerequisite

Run the command as the `root` user.

### Syntax

```
dbaascli tde encryptTablespacesInPDB --dbname <value> --pdbName <value>
[--executePrereqs]
```

Where:

- `--pdbName` specifies the name of the PDB to encrypt all the tablespaces.
- `--dbname` specifies the name of the Oracle Database.
- `--executePrereqs` execute the prerequisites checks and report the results.

### Frequently Asked Questions

#### Q: What does the `dbaascli tde encryptTablespacesInPDB` command do?

A: The `dbaascli tde encryptTablespacesInPDB` command encrypts all the tablespaces in the specified pluggable database (PDB) for an Oracle Database on Exadata Cloud@Customer.

#### Q: Who should run the `dbaascli tde encryptTablespacesInPDB` command?

A: The command must be run as the `root` user.

#### Q: On which machine should I run the `dbaascli tde encryptTablespacesInPDB` command?

A: You need to connect to an Exadata Cloud@Customer virtual machine using SSH to run this command.

#### Q: Where can I find instructions for connecting to the virtual machine?

A: Refer to the guide "Connecting to a Virtual Machine with SSH" for connection instructions.

#### Q: What does the `--pdbName` option specify?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) whose tablespaces need to be encrypted.

#### Q: What does the `--dbname` option do?

A: The `--dbname` option allows you to specify the name of the Oracle Database to which the PDB belongs.

#### Q: What does the `--precheckOnly` option do?

A: The `--precheckOnly` option runs a precheck of the encryption operation without making any actual changes. Valid values are `yes` or `no`.

#### Q: What does the `--useSysdbaCredential` option do?

A: The `--useSysdbaCredential` option specifies whether SYSDBA credentials should be used for the operation. Valid values are true or false.

**Q: Can you give an example of how to encrypt tablespaces in a specific PDB?**

A: Here's an example to encrypt all tablespaces in a PDB named `mypdb`:

```
dbascli tde encryptTablespacesInPDB --pdbName mypdb
```

**Q: How do I encrypt tablespaces in a specific PDB within a database?**

A: Use the following command to specify both the PDB and the database:

```
dbascli tde encryptTablespacesInPDB --pdbName mypdb --dbname mydatabase
```

**Q: How do I run a precheck without performing the encryption?**

A: You can run a precheck only using this syntax:

```
dbascli tde encryptTablespacesInPDB --pdbName mypdb --precheckOnly yes
```

**Q: How do I use SYSDBA credentials to encrypt the tablespaces?**

A: You can use the SYSDBA credentials by adding the `--useSysdbaCredential true` option:

```
dbascli tde encryptTablespacesInPDB --pdbName mypdb --useSysdbaCredential true
```

**Q: What are the prerequisites for running the dbascli tde encryptTablespacesInPDB command?**

A: You must run the command as the `root` user and have access to the Exadata Cloud@Customer virtual machine.

**Q: Do I need to restart the database to encrypt the tablespaces?**

A: No, the command does not require a database restart. The encryption is performed while the database is online.

**Q: Do I need SYSDBA credentials to encrypt tablespaces?**

A: You may need SYSDBA credentials for this operation if specified using the `--useSysdbaCredential` option.

**Q: What should I do if the command fails?**

A: Ensure you are running the command as the `root` user, and verify that the PDB name (`--pdbName`) and database name (`--dbname`) are correct. You can also run the command with `--precheckOnly yes` to check for issues before running the full encryption.

**Q: What should I do if encryption of the tablespaces fails?**

A: Check the database logs and ensure that you have the necessary privileges and resources to perform the encryption. You may also need to verify that there is enough space to handle the encryption process.

**Q: How can I check if the tablespaces in a PDB are encrypted?**

A: You can query the database views related to encryption, such as `V$ENCRYPTED_TABLESPACES`, to verify if the tablespaces have been successfully encrypted.

**Q: How do I verify if the precheck was successful?**

A: If you ran the command with `--precheckOnly yes`, you can check the output for any warnings or errors indicating potential issues with the encryption process.

**Q: Can I encrypt the tablespaces for multiple PDBs simultaneously?**

A: No, you need to run the `dbaascli tde encryptTablespacesInPDB` command separately for each PDB.

**Q: Can I partially encrypt some tablespaces in a PDB?**

A: No, this command encrypts all tablespaces within the specified PDB. For partial encryption, you would need to use different database management commands.

**Q: Does encrypting tablespaces impact database performance?**

A: Encrypting tablespaces can have a temporary performance impact during the encryption process. However, the impact should be minimal once the encryption is complete.

**Q: Can I undo the encryption of tablespaces?**

A: No, once the tablespaces are encrypted, the encryption cannot be undone. You can only rotate or re-encrypt the keys as needed.

**Q: What happens if the operation is interrupted during the encryption process?**

A: If the operation is interrupted, you may need to rerun the command. The system will resume encryption from where it left off, and you can verify the status using database views.

**Example 6-51 dbaascli tde encryptTablespacesInPDB**

```
dbaascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb
```

```
dbaascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb --
executePrereqs
```

## 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`.

### Frequently Asked Questions

#### Q: What is the purpose of the dbascli tde fileToHsm command?

A: The `dbascli tde fileToHsm` command is used to convert a FILE-based Transparent Data Encryption (TDE) to Hardware Security Module (HSM)-based TDE, such as KMS or OKV, in an Oracle Database Cloud Service environment.

#### Q: Who can run the dbascli tde fileToHsm command?

A: The command must be run as the `root` user.

#### Q: What is the purpose of the --kmsKeyOCID parameter?

A: The `--kmsKeyOCID` parameter specifies the KMS key OCID that will be used for TDE encryption when transitioning from file-based to HSM-based TDE.

#### Q: What does the --dbname parameter do?

A: The `--dbname` parameter specifies the name of the database for which you are converting the TDE from file-based to HSM-based.

#### Q: Can I skip the patch validation check while converting TDE?

A: Yes, by using the `--skipPatchCheck` parameter with the value `true`, you can skip the validation check for required patches.

#### Q: What is the --executePrereqs parameter used for?

A: The `--executePrereqs` parameter allows you to run only the prechecks for the TDE conversion process without performing the actual conversion. Valid values are `yes` or `no`.

#### Q: What does the --primarySuc parameter do in a Data Guard setup?

A: The `--primarySuc` parameter is used in a Data Guard environment to indicate that the command has been successfully run on the primary database. It should be specified on the standby database after the primary conversion is complete.

#### Q: How do I resume a previous TDE conversion?

A: You can resume a previously incomplete TDE conversion by using the `--resume` parameter. Optionally, you can specify a specific session ID with `--sessionID`.

#### Q: How do I revert a TDE conversion?

A: To revert a previous TDE conversion, use the `--revert` parameter. You can also provide the specific session ID you want to revert using `--sessionID`.

**Q: How do I specify a session ID when resuming or reverting a TDE conversion?**

A: You can use the `--sessionID` parameter to specify the ID of the session you want to resume or revert. Example: `--resume --sessionID <ID>` or `--revert --sessionID <ID>`.

**Q: What happens if I set `--waitForCompletion` to false?**

A: If you set `--waitForCompletion` to `false`, the TDE conversion process will run in the background, and the command prompt will return immediately. If set to `true`, the command will wait for the process to finish before returning control to the user.

**Q: What are the valid values for the `--waitForCompletion` parameter?**

A: Valid values are `true` or `false`. Setting it to `true` makes the command wait until the process is complete; setting it to `false` runs the process in the background.

**Q: Can I run `dbascli tde fileToHsm` without converting the TDE immediately?**

A: Yes, you can use the `--executePrereqs yes` parameter to perform only the prechecks for the conversion, without making any changes to the TDE.

**Q: In a Data Guard environment, how do I handle the standby database after converting TDE on the primary?**

A: After successfully running the conversion on the primary database, you need to specify `--primarySuc` when running the command on the standby database.

**Q: What should I do if the TDE conversion process fails?**

A: If the process fails, you can use the `--resume` parameter to try resuming from where it left off. If necessary, you can use the `--revert` parameter to roll back the changes made during the failed session.

**Example 6-52 dbascli tde fileToHsm --kmsKeyOCID**

```
dbascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5uxa
invvxza
```

```
dbascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5uxa
invvxza --executePrereqs
```

```
dbascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5uxa
invvxza --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

```
dbaascli 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

## Frequently Asked Questions

### Q: What does the `dbaascli tde getHsmKeys` command do?

A: The `dbaascli tde getHsmKeys` command retrieves details of active Transparent Data Encryption (TDE) keys from the Hardware Security Module (HSM) for a specified database.

### Q: Who should run the `dbaascli tde getHsmKeys` command?

A: The command must be run as the `root` user.

### Q: On which machine should I run the `dbaascli tde getHsmKeys` command?

A: You must connect to an Exadata Cloud@Customer virtual machine using SSH to run this command.

### Q: Where can I find instructions for connecting to the virtual machine?

A: Refer to the guide "Connecting to a Virtual Machine with SSH" for instructions on connecting.

### Q: What does the `--dbname` option do?

A: The `--dbname` option allows you to specify the name of the Oracle Database for which you want to retrieve TDE key details.

### Q: What does the `--infoFile` option do?

A: The `--infoFile` option specifies the file path where the list of key OCIDs (Oracle Cloud Identifiers) will be saved. The output is in JSON format.

### Q: Can you give an example of how to retrieve TDE key details for a specific database?

A: Here's an example to get the TDE key details for a database named `mydatabase`:

```
dbaascli tde getHsmKeys --dbname mydatabase
```

### Q: How do I save the TDE key details to a file?

A: You can specify a file path using the `--infoFile` option to save the output in JSON format:

```
dbaascli tde getHsmKeys --dbname mydatabase --infoFile /path/to/output.json
```

### Q: What are the prerequisites for running the `dbaascli tde getHsmKeys` command?

A: You must run the command as the `root` user and be connected to the Exadata Cloud@Customer virtual machine.

### Q: Do I need SYSDBA credentials to retrieve TDE key details?

A: No, SYSDBA credentials are not required to run the `dbaascli tde getHsmKeys` command.

**Q: In what format is the TDE key information saved when using the `--infoFile` option?**

A: The output is saved in JSON format.

**Q: What information is included in the TDE key details?**

A: The details include key OCIDs and other metadata about the active encryption keys stored in the HSM for the specified database.

**Q: What should I do if the command fails to retrieve the key details?**

A: Ensure that you are running the command as the root user and that the database name (`--dbname`) is correct. Check your connection to the Exadata Cloud@Customer virtual machine.

**Q: How do I check if the output file was created successfully?**

A: You can check the specified file path for the output JSON file. If the file is missing, verify that the file path is correct and that you have write permissions to the directory.

**Q: What should I do if the output file is empty?**

A: Ensure that the database specified contains active TDE keys and that the `--dbname` parameter is correct. You may also need to check if there are any errors in the database logs.

**Q: Can I retrieve TDE key details for multiple databases at once?**

A: No, you must run the `dbaascli tde getHsmKeys` command separately for each database.

**Q: How can I use the output file from the `--infoFile` option in other operations?**

A: Since the output is in JSON format, you can parse the file programmatically or use it as input for other database or encryption management tasks.

**Q: Can I get historical TDE key details using this command?**

A: No, the command only retrieves details about the currently active keys in the HSM.

**Q: How do I verify that the keys retrieved are correct?**

A: You can verify the keys by cross-referencing them with the Oracle Cloud Infrastructure (OCI) console or using database views related to encryption management.

### **Example 6-53 dbaascli tde getHsmKeys**

```
dbaascli tde getHsmkeys --dbname dbname
```

```
dbaascli tde getHsmkeys --dbname dbname --infoFile infoFilePath
```

## dbaascli tde getMkidForKeyVersionOCID

To get Master Key ID associated with the KMS key version OCID, use the `dbaascli 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`.

## Frequently Asked Questions

### Q: What is the purpose of the dbaascli tde getMkidForKeyVersionOCID command?

A: The `dbaascli tde getMkidForKeyVersionOCID` command retrieves the Master Key ID (MKID) associated with a specific KMS key version OCID in Oracle Database Cloud Service environments.

### Q: What are the prerequisites for running the dbaascli tde getMkidForKeyVersionOCID command?

A: You must:

- Run the command as the `root` user.
- Be connected to an Exadata Cloud@Customer virtual machine via SSH.

### Q: Who can run the dbaascli tde getMkidForKeyVersionOCID command?

A: Only the `root` user can run this command.

### Q: What does the --kmsKeyVersionOCID parameter specify?

A: The `--kmsKeyVersionOCID` parameter specifies the KMS key version OCID for which you want to retrieve the associated Master Key ID (MKID).

### Q: What does the --dbname parameter specify?

A: The `--dbname` parameter specifies the name of the database for which the KMS key version OCID is being queried.

### Q: Is the --dbname parameter mandatory?

A: No, the `--dbname` parameter is optional. If you don't specify a database name, the command will retrieve the MKID for the default database on the system.

### Q: What should I do if I don't know the KMS key version OCID?

A: You must retrieve the KMS key version OCID from your KMS management console or service provider before using this command. Without it, the command cannot retrieve the Master Key ID (MKID).

### Q: Can I run this command on a non-Exadata Cloud@Customer environment?

A: No, this command is specifically for use in an Exadata Cloud@Customer environment, and you need to connect to a virtual machine using SSH to execute it.



**Q: What happens if I run the command without specifying a database name using --dbname?**

A: If the --dbname parameter is not provided, the command will attempt to retrieve the MKID for the default database configured on the system.

**Q: What should I do if I encounter an error while retrieving the MKID?**

A: Ensure that:

- You are running the command as the root user.
- You are correctly connected to the Exadata Cloud@Customer virtual machine.
- The KMS key version OCID you provided is valid. If the error persists, check the system logs for more details.

**Q: How do I connect to the Exadata Cloud@Customer virtual machine?**

A: You can connect to the virtual machine via SSH. Refer to the Exadata Cloud@Customer documentation for steps on how to securely connect.

**Example 6-54 dbaascli tde getMkidForKeyVersionOCID**

```
dbaascli tde getMkidForKeyVersionOCID --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgi23nha
```

## 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

### Frequently Asked Questions

**Q: What is the purpose of the dbaascli tde getPrimaryHsmKey command?**

A: The dbaascli tde getPrimaryHsmKey command retrieves the primary 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 getPrimaryHsmKey command?**

A: You must:

- Run the command as the `root` user.
- Be connected to an Exadata Cloud@Customer virtual machine via SSH.

**Q: Who can execute the `dbascli tde getPrimaryHsmKey` command?**

A: Only the `root` user can execute this command.

**Q: What does the `--dbname` parameter specify in this command?**

A: The `--dbname` parameter specifies the name of the database for which you want to retrieve the primary HSM key.

**Q: Is the `--dbname` parameter mandatory?**

A: No, the `--dbname` parameter is optional. If not provided, the command will retrieve the primary HSM key for the default database on the system.

**Q: What should I do if I don't specify a database name with `--dbname`?**

A: If the `--dbname` parameter is not specified, the command will attempt to retrieve the primary HSM key for the default database configured on the system.

**Q: Can I run this command on a non-Exadata Cloud@Customer environment?**

A: No, this command is designed specifically for use in an Exadata Cloud@Customer environment, and you must be connected to the virtual machine using SSH to run it.

**Q: How do I connect to the Exadata Cloud@Customer virtual machine to run the command?**

A: You can connect to the virtual machine via SSH. Refer to the Exadata Cloud@Customer documentation for instructions on how to securely connect.

**Q: What should I check if I encounter an error while retrieving the primary HSM key?**

A: If you encounter an error, ensure that:

- You are running the command as the `root` user.
- You are correctly connected to the Exadata Cloud@Customer virtual machine.
- The database name (if specified) is valid. If the issue persists, consult system logs or error messages for more details.

**Q: Do I need to stop the database to run the `dbascli tde getPrimaryHsmKey` command?**

A: No, the database does not need to be stopped to run this command. You can execute it while the database is running.

**Q: What is the purpose of retrieving the primary HSM key?**

A: Retrieving the primary HSM key allows you to identify the current HSM key that is being used for encryption in the database's existing HSM (KMS) configuration.

**Example 6-55 `dbascli tde getPrimaryHsmKey`**

```
dbascli 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`

### Frequently Asked Questions

**Q: What is the purpose of the `dbaascli tde hsmToFile` command?**

A: The `dbascli tde hsmToFile` command is used to convert a Hardware Security Module (HSM)-based Transparent Data Encryption (TDE) to a file-based TDE in Oracle Database Cloud Service environments.

**Q: What are the prerequisites for running the `dbascli tde hsmToFile` command?**

A: You must:

- Run the command as the `root` user.
- Ensure you have the necessary permissions and configurations set in the database environment.

**Q: What does the `--dbname` parameter specify?**

A: The `--dbname` parameter specifies the name of the database for which you are converting TDE from HSM-based to file-based.

**Q: When is the `--primaryDBWalletTar` parameter required?**

A: The `--primaryDBWalletTar` parameter is required only when performing the `hsmToFile` conversion on a standby database. It specifies the tar file of the primary database's wallet.

**Q: What is the purpose of the `--skipPatchCheck` parameter?**

A: The `--skipPatchCheck` parameter allows you to skip the validation check for required patches. Set this to `true` to skip the check or `false` to enforce it.

**Q: How do I run only prechecks for the conversion process without performing the actual conversion?**

A: You can use the `--executePrereqs` parameter and set it to `yes` to run only the prechecks. Set it to `no` to perform the full conversion.

**Q: What does the `--primarySuc` parameter do in a Data Guard environment?**

A: The `--primarySuc` parameter is used in a Data Guard setup to indicate that the conversion has successfully run on the primary database. It should be used when running the conversion on the standby database.

**Q: How can I resume a previous `hsmToFile` conversion?**

A: You can resume a previous conversion by using the `--resume` parameter. Optionally, you can specify the session ID of the previous run with `--sessionID`.

**Q: What is the purpose of the `--revert` parameter?**

A: The `--revert` parameter is used to roll back a previously initiated conversion process in case of failure or if you need to undo the operation.

**Q: What happens if I set `--waitForCompletion` to `false`?**

A: If you set `--waitForCompletion` to `false`, the operation will run in the background, allowing you to continue other tasks. If set to `true`, the command will wait for the process to complete before returning control to the user.

**Q: What should I do if I need to convert the TDE in a standby database in a Data Guard setup?**

A: In a Data Guard setup, after converting TDE on the primary database, you must run the command on the standby database using the `--primaryDBWalletTar` parameter, specifying the wallet tar file from the primary database, and include `--primarySuc`.

**Q: What should I do if I want to skip checking for required patches during the conversion?**

A: You can skip the patch check by using the `--skipPatchCheck` parameter and setting it to `true`.

**Q: How do I check if the system is ready for the hsmToFile conversion without making changes?**

A: You can perform only the prechecks by using the `--executePrereqs` parameter and setting it to `yes`.

**Q: What should I do if the conversion process is interrupted?**

A: You can use the `--resume` parameter to restart the process from where it left off. Optionally, you can specify a particular session ID with `--sessionID`.

**Q: What should I do if the conversion process fails?**

A: If the conversion fails, you can roll back the process using the `--revert` parameter. Additionally, review any error messages and check system logs for more details.

**Q: Can I run the dbascli tde hsmToFile command on a non-Exadata environment?**

A: This command is designed for use in Exadata Cloud@Customer environments. If you are not using Exadata, ensure that you are in a supported environment for the command to work properly.

**Example 6-56 dbascli tde hsmToFile**

```
dbascli tde hsmToFile --dbname dbname
```

```
dbascli tde hsmToFile --dbname dbname --executePrereqs
```

```
dbascli tde hsmToFile --dbname dbname --resume
```

## dbascli tde listKeys

To list TDE master keys, use the `dbascli tde listKeys` command.

### Prerequisite

Run the command as the `root` user.

### Syntax

```
dbascli 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.

## Frequently Asked Questions

### Q: What is the purpose of the dbaascli tde listKeys command?

A: The `dbaascli tde listKeys` command is used to list all the Transparent Data Encryption (TDE) master keys for a specified database in an Oracle Database environment.

### Q: What are the prerequisites for running the dbaascli tde listKeys command?

A: You must:

- Run the command as the `root` user.
- Be connected to an Exadata Cloud@Customer virtual machine using SSH.

### Q: What does the --file parameter do in the dbaascli tde listKeys command?

A: The `--file` parameter specifies the file path where the list of TDE master keys should be saved. If this parameter is not provided, the results will be displayed directly in the terminal.

### Q: What does the --dbname parameter specify?

A: The `--dbname` parameter specifies the name of the database for which you want to list the TDE master keys.

### Q: Is the --file parameter mandatory?

A: No, the `--file` parameter is optional. If not provided, the list of TDE keys will be shown in the terminal output instead of being saved to a file.

### Q: Is the --dbname parameter mandatory?

A: No, the `--dbname` parameter is optional. If not specified, the command will list the TDE master keys for the default database configured on the system.

### Q: What should I do if I want to save the list of keys to a file?

A: You should provide the `--file` parameter along with the desired file path. For example:

```
dbaascli tde listKeys --file /path/to/output.txt
```

### Q: What happens if I don't provide a database name with --dbname?

A: If the `--dbname` parameter is not provided, the command will list the TDE master keys for the default database on the system.

### Q: Can I use this command in environments other than Exadata Cloud@Customer?

A: This command is designed specifically for Exadata Cloud@Customer environments. Ensure you are connected to the appropriate virtual machine to run it.

### Q: What should I do if the command fails to list keys?

A: Ensure that:

- You are running the command as the `root` user.
- You are connected to the Exadata Cloud@Customer virtual machine.
- The database name (if specified) is correct. Check the error messages and logs for more details on the failure.

### Q: Can I run the dbaascli tde listKeys command while the database is running?

A: Yes, the command can be executed while the database is running. It simply lists the TDE master keys and does not alter the state of the database.

**Q: Do I need special permissions to run this command?**

A: You must run this command as the `root` user. Without root permissions, you will not be able to execute the command.

**Q: What is the purpose of listing TDE master keys?**

A: Listing TDE master keys allows you to review the encryption keys being used for protecting your database's data. It is essential for monitoring and managing encryption settings.

**Q: How do I connect to the Exadata Cloud@Customer virtual machine to run the command?**

A: You can connect to the virtual machine using SSH. Refer to the Exadata Cloud@Customer documentation for instructions on how to establish a secure connection.

**Example 6-57 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@Customer 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@Customer 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@Customer virtual machine to run the command?**

A: You can connect to the virtual machine using SSH. Refer to the Exadata Cloud@Customer documentation for instructions on establishing a secure connection.

**Example 6-58 dbaascli tde removeSecondaryHsmKey**

```
dbaascli tde removeSecondaryHsmKey --dbname dbname
```

```
dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsghfa2xe5prvlzdxtygoiqm2pu2afgta54krxwllk5ux
ainvvxza
```

```
dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsghfa2xe5prvlzdxtygoiqm2pu2afgta54krxwllk5ux
ainvvxza --executePrereqs
```

## dbaascli tde rotateMasterKey

To rotate the master key for database encryption, use the `dbaascli tde rotateMasterKey` command.

### Prerequisites:

Run the command as the `root` user.

### Syntax

```
dbaascli tde rotateMasterKey --dbname <value>
[--rotateMasterKeyOnAllPDBs]
[--pdbName <value>]
[--executePrereqs]
[--resume [--sessionID <value>]]
{
    [--prepareStandbyBlob <value> [--blobLocation <value>]]
    | [--standbyBlobFromPrimary <value>]
}
```

### Where:

- `--dbname` specifies the name of the Oracle Database
- `--rotateMasterKeyOnAllPDBs` specifies `true` to rotate master key of all PDBs in CDB.  
Valid values: `true|false`
- `--pdbName` specifies the name of the PDB
- `--executePrereqs` runs the prerequisites checks and report the results
- `--resume` specifies to resume the previous execution
- `--sessionID` specifies to resume a specific session ID
- `--prepareStandbyBlob` specifies `true` to generate a BLOB file containing the artifacts needed to perform the operation in a Data Guard environment
- `--blobLocation` specifies the location of the custom directory where the standby BLOB file will be generated in a Data Guard environment
- `--standbyBlobFromPrimary` specifies the location of the standby BLOB file, which is prepared from the primary database. This is required only for standby operations.

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli tde rotateMasterKey` command?

A: The `dbaascli tde rotateMasterKey` command is used to rotate the master key used for Transparent Data Encryption (TDE) in an Oracle Database. This process ensures the encryption keys are updated for better security.

#### Q: What are the prerequisites for running the `dbaascli tde rotateMasterKey` command?

A: You must:

- Run the command as the `root` user.
- Ensure that the database is configured correctly for TDE.

**Q: What does the --dbname parameter specify?**

A: The --dbname parameter specifies the name of the Oracle Database for which you want to rotate the master encryption key.

**Q: What is the purpose of the --rotateMasterKeyOnAllPDBs parameter?**

A: The --rotateMasterKeyOnAllPDBs parameter specifies whether to rotate the master key for all Pluggable Databases (PDBs) in a Container Database (CDB). Valid values are true or false.

**Q: What does the --pdbName parameter do?**

A: The --pdbName parameter specifies the name of a particular Pluggable Database (PDB) if you want to rotate the master key for a specific PDB rather than all PDBs.

**Q: What does the --executePrereqs parameter do?**

A: The --executePrereqs parameter runs prerequisite checks to validate whether the environment is ready for the master key rotation without performing the actual rotation.

**Q: What does the --resume parameter specify?**

A: The --resume parameter is used to resume a previously started operation. You can also provide a specific session ID using --sessionID to resume a particular session.

**Q: What is the purpose of the --prepareStandbyBlob parameter?**

A: The --prepareStandbyBlob parameter, if set to true, generates a BLOB file containing the necessary artifacts to perform master key rotation in a Data Guard environment.

**Q: What does the --blobLocation parameter do?**

A: The --blobLocation parameter specifies a custom directory path where the standby BLOB file will be generated. This is applicable when --prepareStandbyBlob is set to true.

**Q: What does the --standbyBlobFromPrimary parameter specify?**

A: The --standbyBlobFromPrimary parameter specifies the location of the standby BLOB file that was generated from the primary database. This parameter is used when performing the master key rotation on a standby database in a Data Guard environment.

**Q: Is the --rotateMasterKeyOnAllPDBs parameter mandatory?**

A: No, the --rotateMasterKeyOnAllPDBs parameter is optional. If it is not specified, the master key will only be rotated for the database (or specific PDB) provided in the --dbname or --pdbName parameters.

**Q: Is the --pdbName parameter required if I'm rotating keys for a CDB?**

A: No, the --pdbName parameter is only required if you want to rotate the master key for a specific Pluggable Database (PDB). It is optional when rotating the key for the entire CDB.

**Q: Do I need to use the --prepareStandbyBlob and --standbyBlobFromPrimary parameters for standalone databases?**

A: No, these parameters are only relevant in a Data Guard environment where a standby database is involved.

**Q: How can I rotate the master key for all PDBs in a CDB?**

A: You should use the --rotateMasterKeyOnAllPDBs parameter set to true to rotate the master key for all PDBs in the CDB. For example:

```
dbaascli tde rotateMasterKey --dbname CDB_NAME --rotateMasterKeyOnAllPDBs true
```

**Q: How do I run a check to validate that the system is ready for master key rotation without performing the actual operation?**

A: You can use the `--executePrereqs` parameter to run the prerequisite checks. This will report any issues that might prevent the master key rotation:

```
dbaascli tde rotateMasterKey --dbname DB_NAME --executePrereqs
```

**Q: What should I do if the operation was interrupted, and I want to resume it?**

A: You can use the `--resume` parameter to resume the previously interrupted operation. If you have a session ID, provide it with the `--sessionID` parameter:

```
dbaascli tde rotateMasterKey --dbname DB_NAME --resume --sessionID <value>
```

**Q: How can I prepare for key rotation in a Data Guard environment?**

A: You should use the `--prepareStandbyBlob` parameter to generate a BLOB file that contains the required artifacts for rotating the master key in a standby environment:

```
dbaascli tde rotateMasterKey --dbname DB_NAME --prepareStandbyBlob true --blobLocation /path/to/blob
```

**Q: How do I apply the standby BLOB file from the primary database when rotating keys on a standby database?**

A: Use the `--standbyBlobFromPrimary` parameter to specify the location of the BLOB file that was prepared on the primary database:

```
dbaascli tde rotateMasterKey --dbname DB_NAME --standbyBlobFromPrimary /path/to/blob
```

**Q: What should I check if the master key rotation fails?**

A: Ensure that:

- You are running the command as the `root` user.
- The database name (`--dbname`) is correct.
- Any prerequisite checks were run using `--executePrereqs` to ensure readiness. Review the error logs for more detailed information on the failure.

**Q: What should I do if the operation does not complete successfully in a Data Guard environment?**

A: Ensure that the BLOB file from the primary database was prepared correctly using `--prepareStandbyBlob`, and then use `--standbyBlobFromPrimary` to apply it on the standby database.

**Q: Can I run the `dbaascli tde rotateMasterKey` command while the database is running?**

A: Yes, the command can be executed while the database is running. However, it is recommended to run prerequisite checks beforehand using the `--executePrereqs` option.

**Q: Why is rotating the master key important?**

A: Rotating the master key improves database security by ensuring that the encryption keys used for data protection are periodically updated, reducing the risk of key compromise.

**Q: Do I need to restart the database after rotating the master key?**

A: No, restarting the database is not required after rotating the master key. The key rotation will take effect immediately without any service disruption.

## 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

```
dbaascli 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`

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli tde setKeyVersion` command?

A: The `dbaascli tde setKeyVersion` command is used to set the version of the primary encryption key that should be used for Transparent Data Encryption (TDE) in a database or Pluggable Database (PDB). This allows for the specific version of a KMS key to be assigned to the database.

#### Q: What are the prerequisites for using the `dbaascli tde setKeyVersion` command?

A: You must run the command as the `root` user and ensure that you are connected to an Exadata Cloud@Customer virtual machine.

#### Q: What does the `--kmsKeyVersionOCID` parameter specify?

A: The `--kmsKeyVersionOCID` parameter specifies the KMS key version OCID (Oracle Cloud Identifier) that you want to set for the database or PDB.

#### Q: What does the `--dbname` parameter specify?

A: The `--dbname` parameter specifies the name of the Oracle Database for which the key version will be set.

**Q: What is the purpose of the --pdbName parameter?**

A: The --pdbName parameter specifies the name of the Pluggable Database (PDB) within a Container Database (CDB) where you want to set the specific KMS key version.

**Q: What is the --masterKeyID parameter used for?**

A: The --masterKeyID parameter specifies the master key ID that is associated with the given KMS key version OCID. This is particularly important in a Data Guard environment.

**Q: What is the role of the --standbySuc parameter?**

A: The --standbySuc parameter is used in a Data Guard environment. It specifies that this property should be set on the primary database after successfully running the command on the standby database.

**Q: What does the --executePrereqs parameter do?**

A: The --executePrereqs parameter specifies whether to run prerequisite checks before executing the operation. Valid values are `yes` or `no`.

**Q: What does the --waitForCompletion parameter control?**

A: The --waitForCompletion parameter determines whether the operation will run synchronously (waiting for completion) or asynchronously (in the background). Valid values are `true` or `false`.

**Q: Is the --pdbName parameter required if setting the key version for a CDB?**

A: No, the --pdbName parameter is only required if you are setting the key version for a specific Pluggable Database (PDB). It is optional if you are setting the key version for the entire Container Database (CDB).

**Q: Is the --masterKeyID parameter necessary for non-Data Guard environments?**

A: No, the --masterKeyID parameter is typically only used in Data Guard environments. For standalone databases, this parameter is not required.

**Q: How do I set the key version for a database?**

A: You can set the key version for a database by running:

```
dbaascli tde setKeyVersion --kmsKeyVersionOCID <value> --dbname <DB_NAME>
```

**Q: How do I set the key version for a specific PDB?**

A: To set the key version for a specific Pluggable Database (PDB), use the --pdbName parameter along with the database name:

```
dbaascli tde setKeyVersion --kmsKeyVersionOCID <value> --dbname <DB_NAME> --  
pdbName <PDB_NAME>
```

**Q: How can I ensure that all prerequisites are met before setting the key version?**

A: You can run the prerequisite checks by using the --executePrereqs parameter:

```
dbaascli tde setKeyVersion --kmsKeyVersionOCID <value> --executePrereqs yes
```

**Q: How do I set the key version in a Data Guard environment?**

A: In a Data Guard environment, you should:

- Run the command on the standby database:

```
dbascli tde setKeyVersion --kmsKeyVersionOCID <value> --masterKeyID <keyID>
--dbname <DB_NAME>
```

- After successfully running the command on the standby database, run the command on the primary database using the `--standbySuc` parameter:

```
dbascli tde setKeyVersion --kmsKeyVersionOCID <value> --dbname <DB_NAME> --
standbySuc yes
```

**Q: How can I run the operation in the background without waiting for it to complete?**

A: You can run the operation asynchronously by setting `--waitForCompletion` to `false`:

```
dbascli tde setKeyVersion --kmsKeyVersionOCID <value> --waitForCompletion false
```

**Q: What should I do if the key version fails to set?**

A: Ensure that:

- You are running the command as the `root` user.
- The KMS key version OCID is correct.
- Any prerequisite checks were run using `--executePrereqs` to ensure readiness. Review error logs for specific details and rerun the operation if needed.

**Q: What should I check if the operation doesn't complete successfully in a Data Guard environment?**

A: Ensure that the `--masterKeyID` parameter is correctly specified when running the command on the standby database. Once completed on the standby, the `--standbySuc` parameter should be used when running the command on the primary database.

**Q: Can I run the dbascli tde setKeyVersion command while the database is running?**

A: Yes, the command can be executed while the database is running. However, running the prerequisite checks beforehand using `--executePrereqs` is recommended.

**Q: Why is it important to set the correct KMS key version for a database?**

A: Setting the correct KMS key version ensures that the database is using the appropriate encryption key version for TDE, which helps maintain data security and compliance with organizational policies.

**Q: What happens if I use the wrong KMS key version OCID?**

A: If an incorrect KMS key version OCID is used, the encryption may fail, and the database will not be able to use the incorrect key for encryption operations. You should ensure that the correct key version OCID is provided.

**Q: Do I need to restart the database after setting the key version?**

A: No, restarting the database is not necessary after setting the key version. The new key version will take effect immediately without requiring a restart.

**Example 6-59 dbascli tde setKeyVersion**

```
dbascli tde setKeyVersion --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
```

```
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha
```

```
dbaascli tde setKeyVersion --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha --executePrereqs
```

```
dbaascli tde setKeyVersion --dbname dbname --pdbName pdb --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha
```

## dbaascli tde setPrimaryHsmKey

To change the primary HSM (KMS) key for the existing HSM (KMS) configuration, use the `dbaascli 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.

### Frequently Asked Questions

#### Q: What is the purpose of the `dbaascli tde setPrimaryHsmKey` command?

A: The `dbaascli tde setPrimaryHsmKey` command is used to change the primary HSM (Hardware Security Module) or KMS (Key Management Service) key in an existing HSM/KMS configuration for Transparent Data Encryption (TDE).

#### Q: What are the prerequisites for running the `dbaascli tde setPrimaryHsmKey` command?



A: The command must be executed as the `root` user, and the environment should be an Exadata Cloud@Customer virtual machine.

**Q: What does the `--primaryKmsKeyOCID` parameter specify?**

A: The `--primaryKmsKeyOCID` parameter specifies the OCID (Oracle Cloud Identifier) of the primary KMS key to be set for the TDE environment.

**Q: What is the function of the `--dbname` parameter?**

A: The `--dbname` parameter specifies the name of the Oracle Database for which the primary HSM/KMS key will be set.

**Q: What does the `--standbySuc` parameter do?**

A: The `--standbySuc` parameter is used in a Data Guard environment. It specifies that the command should be run on the primary database after successfully executing it on the standby database.

**Q: What is the purpose of the `--precheckOnly` parameter?**

A: The `--precheckOnly` parameter allows you to run only the prechecks for this operation. It validates the environment without making any actual changes. Valid values are `yes` or `no`.

**Q: What does the `--bounceDatabase` parameter control?**

A: The `--bounceDatabase` parameter specifies whether the database should be bounced (restarted) in a rolling manner as part of the operation. This ensures minimal downtime by restarting parts of the database one by one.

**Q: How do I set the primary KMS key for my database?**

A: To set the primary KMS key, run the following command:

```
dbascli tde setPrimaryHsmKey --primaryKmsKeyOCID <key_OCID> --dbname <DB_NAME>
```

**Q: How do I ensure that the operation can be executed without any issues?**

A: Run the operation with the `--precheckOnly` parameter to verify that all prerequisites are met:

```
dbascli tde setPrimaryHsmKey --primaryKmsKeyOCID <key_OCID> --precheckOnly yes
```

**Q: How do I set the primary KMS key in a Data Guard environment?**

A: First, run the command on the standby database:

```
dbascli tde setPrimaryHsmKey --primaryKmsKeyOCID <key_OCID> --dbname <DB_NAME>
```

Then, run the command on the primary database with the `--standbySuc` parameter:

```
dbascli tde setPrimaryHsmKey --primaryKmsKeyOCID <key_OCID> --dbname <DB_NAME>  
--standbySuc yes
```

**Q: How do I minimize downtime while changing the primary KMS key?**

A: You can use the `--bounceDatabase` parameter to perform a rolling restart, minimizing downtime:

```
dbascli tde setPrimaryHsmKey --primaryKmsKeyOCID <key_OCID> --bounceDatabase
```

**Q: Is the `--dbname` parameter required for all databases?**

A: Yes, you should specify the `--dbname` parameter to indicate the target database for which the primary KMS key should be set.

**Q: Is it mandatory to use the `--standbySuc` parameter in a Data Guard environment?**

A: Yes, the `--standbySuc` parameter must be used when running the command on the primary database after successfully executing it on the standby database.

**Q: Can I skip the bounce operation for the database?**

A: Yes, if you do not specify the `--bounceDatabase` parameter, the database will not be bounced (restarted) as part of the operation.

**Q: What should I do if the command fails during execution?**

A: If the command fails, ensure:

- You are running it as the `root` user.
- The correct `--primaryKmsKeyOCID` and `--dbname` values are provided.
- The environment passes all prerequisite checks (run with `--precheckOnly`).

**Q: What if the operation fails in a Data Guard environment?**

A: Ensure that the command has been successfully executed on the standby database before running it on the primary. Check for errors in the logs and rerun the operation with correct parameters.

**Q: Can I run the `dbaascli tde setPrimaryHsmKey` command on a live database?**

A: Yes, the command can be run while the database is live. However, using the `--bounceDatabase` parameter will restart the database in a rolling fashion, which minimizes the impact.

**Q: How do I run the command in a rolling manner to avoid complete downtime?**

A: Use the `--bounceDatabase` parameter to perform a rolling restart of the database while changing the primary KMS key:

```
dbaascli tde setPrimaryHsmKey --primaryKmsKeyOCID <key_OCID> --bounceDatabase
```

**Q: What is the significance of changing the primary KMS key?**

A: Changing the primary KMS key ensures that the database uses an updated or different encryption key for Transparent Data Encryption (TDE). This may be required for security or compliance reasons.

**Q: How often should the primary KMS key be rotated or changed?**

A: While there is no strict rule, organizations may rotate the primary KMS key based on security policies, such as key rotation intervals or compliance requirements.

**Q: What happens if the primary KMS key is set incorrectly?**

A: If the incorrect key OCID is set, database encryption operations may fail, and you may need to revert to the correct key or rectify the configuration by setting the correct KMS key OCID.

**Q: Do I need to restart the database after changing the primary KMS key?**

A: No, you do not need to restart the database unless you choose to use the `--bounceDatabase` parameter, which will automatically restart the database to apply the change.

### Example 6-60 dbaascli tde setPrimaryHsmKey

```
dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocidl.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsghfa2xe5prvlzdxygoiqm2pu2afgta54krxwllk5ux
ainvvxza
```

```
dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocidl.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsghfa2xe5prvlzdxygoiqm2pu2afgta54krxwllk5ux
ainvvxza --executePrereqs
```

## dbaascli tde status

To display information about the keystore for the specified database, use the `dbaascli tde status` command.

### 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.

### Frequently Asked Questions

#### Q: What does the dbaascli tde status command do?

A: The `dbaascli tde status` command displays information about the keystore for a specified database. This includes details about the type of keystore and its status.

#### Q: Who should run the dbaascli tde status command?

A: The command must be executed as the `oracle` user.

#### Q: Where should the dbaascli tde status command be run?

A: The command should be run on an Exadata Cloud@Customer virtual machine. You must connect to the virtual machine via SSH to run the utility.

#### Q: What is the function of the --dbname parameter?

A: The `--dbname` parameter specifies the name of the database for which the TDE keystore status will be checked.

#### Q: What information does the dbaascli tde status command return?

A: The output of the command includes the type of keystore (e.g., HSM-based or file-based) and the current status of the keystore, such as whether it is open, closed, or in some other state.

**Q: How can I know if the keystore is open or closed using the dbascli tde status command?**

A: The status of the keystore, including whether it is open or closed, is part of the output returned by the `dbascli tde status` command.

**Q: How do I check the status of the TDE keystore for a specific database?**

A: To check the TDE keystore status for a specific database, run:

```
dbascli tde status --dbname <DB_NAME>
```

**Q: Can I check the status of the keystore for multiple databases?**

A: Yes, but you need to run the command separately for each database, specifying its name using the `--dbname` parameter.

**Q: Can the dbascli tde status command be run as the root user?**

A: No, the command should be executed as the `oracle` user, not the `root` user.

**Q: Do I need special permissions to run the dbascli tde status command?**

A: Yes, you need to have `oracle` user privileges and be connected to an Exadata Cloud@Customer virtual machine to run the command.

**Q: What should I do if I get an error when running the dbascli tde status command?**

A: Ensure you are running the command as the `oracle` user, and verify that you have the necessary permissions and are connected to the correct virtual machine.

**Q: How do I know what type of keystore my database is using?**

A: The type of keystore, such as whether it is file-based or HSM/KMS-based, is displayed in the output of the `dbascli tde status` command.

**Q: What should I do if the keystore is closed?**

A: If the keystore is closed, you may need to open it manually, depending on the operation you are trying to perform. The exact process will depend on the keystore type and your environment.

**Q: Can I view keystore status for a CDB (Container Database) or PDB (Pluggable Database)?**

A: Yes, by specifying the appropriate database name using the `--dbname` parameter, you can view the keystore status for both CDBs and PDBs.

**Q: What does it mean if the command returns an error about database connectivity?**

A: This could indicate an issue with the connection to the database or a problem with your environment. Ensure that the database is running and accessible, and verify your SSH connection to the Exadata Cloud@Customer virtual machine.

**Q: What happens if the database name is incorrect?**

A: If the `--dbname` parameter specifies an incorrect or non-existent database, the command will fail, and you'll receive an error message indicating the problem.

**Q: How do I troubleshoot if the keystore status indicates an unexpected state?**

A: If the keystore status indicates an unexpected state, review the database logs for more details and check the configuration of the keystore to ensure it is properly set up.

**Q: Can I automate the check of keystore status for monitoring purposes?**

A: Yes, you can script the `dbaascli tde status` command to check keystore status periodically or integrate it into your database monitoring tools.

**Q: How do I verify that Transparent Data Encryption (TDE) is correctly enabled?**

A: You can verify TDE is correctly enabled by checking the status of the keystore using the `dbaascli tde status` command. A valid and open keystore indicates that TDE is properly configured.

**Example 6-61 dbaascli tde status**

```
dbaascli tde status --dbname dbname
```

## 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.

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 Oracle Exadata Database Service on Exascale 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
\u25cf syslens.service
Loaded: loaded (/etc/systemd/system/syslens.service; disabled; vendor
preset: disabled)
Active: active (running) since Wed 2022-03-16 18:08:59 UTC; 34s ago
Main PID: 358039 (python3)
Memory: 31.6M
CGroup: /system.slice/syslens.service
\u2514\u2500358039 /usr/bin/python3 /var/opt/oracle/syslens/bin/
syslens_main.py --archive /var/opt/oracle/log/...

Mar 16 18:08:59 node1 systemd[1]: Started syslens.service.
Mar 16 18:09:09 node1 su[360495]: (to oracle) root on none
Mar 16 18:09:09 node1 su[360539]: (to grid) root on none
Mar 16 18:09:10 node1 su[360611]: (to grid) root on none
Mar 16 18:09:11 node1 su[360653]: (to oracle) root on none
```

- 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 configured to notify events, run the `/usr/bin/syslens --config /var/opt/oracle/syslens/data/exacc.syslens.config --get-key enable_telemetry` command as root in the domU:

```
# /usr/bin/syslens --config /var/opt/oracle/syslens/data/
exacc.syslens.config --get-key enable_telemetry
syslens Collection 2.3.3
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

- [Using the Console 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](#)

## 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.

- [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](#)

## 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 Database**, then click **Exadata Database Service on Exascale 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 Database**, then click **Exadata on Oracle Public Cloud**.
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 Database**, then click **Exadata on Oracle Public Cloud**.
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 Exascale Infrastructure in the Monitoring Service

learn about 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.

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
CpuUtilization	CPU Utilization	percentage	Percent CPU utilization	1 minute	hostName deploymentType

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
FilesystemUtilization	<b>Filesystem Utilization</b>	percentage	Percent utilization of provisioned filesystem	1 minute	hostName deploymentType filesystemName
LoadAverage	<b>Load Average</b>	integer	System load average over 5 minutes	1 minute	hostName deploymentType
MemoryUtilization	<b>Memory Utilization</b>	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
NodeStatus	<b>Node Status</b>	integer	Indicates whether the host is reachable.	1 minute	hostName deploymentType
SwapUtilization	<b>Swap Utilization</b>	percentage	Percent utilization of total swap space	1 minute	hostName deploymentType

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	<b>CPU Utilization</b>	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.	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	<b>Storage Utilization</b>	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	<b>DB Block Changes</b>	Changes per second	The Average number of blocks changed per second.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database pdb} resourceName_{database pdb}
ExecuteCount	<b>Execute Count</b>	Count	The number of user and recursive calls that executed SQL statements during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType
ExecuteCount	<b>Execute Count</b>	Count	The number of user and recursive calls that executed SQL statements during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType
CurrentLogons	<b>Current Logons</b>	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	<b>Transaction Count</b>	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	<b>User Calls</b>	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	<b>Parse Count</b>	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	<b>Storage Space Used</b>	GB	Total amount of storage space used by the database at the collection time.	1 hour	deploymentType resourceId_{database pdb} resourceName_{database pdb}
StorageAllocated	<b>Storage Space Allocated</b>	GB	Total amount of storage space allocated to 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
StorageUsedByTablespace	<b>Storage Space Used By Tablespace</b>	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	<b>Allocated Storage Space By Tablespace</b>	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	<b>Storage Space Utilization By Tablespace</b>	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

## Oracle Exadata Database Service on Exascale Infrastructure Events

Oracle Exadata Database Service on Exascale Infrastructure resources emit events, which are structured messages that indicate changes in resources.

- [About Event Types on Oracle Exadata Database Service on Exascale Infrastructure](#)  
Learn about the event types available for Oracle Exadata Database Service on Exascale 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 Exascale Infrastructure Event Types](#)  
Learn about the event types available for Exadata Database Service on Exascale Infrastructure resources.

- [Oracle Exadata Database Service on Exascale 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.
- [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 Association Event Types](#)  
Review the list of event types that 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.

## About Event Types on Oracle Exadata Database Service on Exascale Infrastructure

Learn about the event types available for Oracle Exadata Database Service on Exascale Infrastructure resources.

Oracle Exadata Database Service on Exascale 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.
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 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**.

## Oracle Exadata Database Service on Exascale Infrastructure Event Types

Learn about the event types available for Exadata Database Service on Exascale Infrastructure resources.

Oracle Exadata Database Service on Exascale 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*.

### Resource Events and Operations for ExaDB-XS

**Table 6-1 Resource Events and Operations for ExaDB-XS**

Friendly Name	Begin Event Sample	End event Sample
Create Storage Vault	com.oraclecloud.DatabaseService.CreateExascaleDbStorageVault.begin	com.oraclecloud.DatabaseService.CreateExascaleDbStorageVault.end

**Table 6-1 (Cont.) Resource Events and Operations for ExaDB-XS**

Friendly Name	Begin Event Sample	End event Sample
Create VM Cluster	com.oraclecloud.DatabaseService.CreateExadbVmCluster.begin	com.oraclecloud.DatabaseService.CreateExadbVmCluster.end
Get ExaDB VM Cluster	com.oraclecloud.DatabaseService.GetExadbVmCluster	This is synchronous operation, so there is no end event.
List ExaDB VM Cluster	com.oraclecloud.databaseservice.ListExadbVmClusters	This is synchronous operation, so there is no end event.
Update ExaDB VM Cluster	com.oraclecloud.DatabaseService.UpdateExadbVmCluster.begin	com.oraclecloud.DatabaseService.UpdateExadbVmCluster.end
Delete ExaDB VM Cluster	com.oraclecloud.DatabaseService.DeleteExadbVmCluster.begin	com.oraclecloud.DatabaseService.DeleteExadbVmCluster.end
Change Compartment ExaDB VM Cluster	com.oraclecloud.DatabaseService.ChangeExadbVmClusterCompartment.begin	com.oraclecloud.DatabaseService.ChangeExadbVmClusterCompartment.end
Remove Virtual Machine ExaDB VM Cluster	com.oraclecloud.DatabaseService.ExadbVmClusterTerminateVirtualMachine.begin	com.oraclecloud.DatabaseService.ExadbVmClusterTerminateVirtualMachine.end
Get Exascale DB Storage Vault	com.oraclecloud.DatabaseService.GetExascaleDbStorageVault	This is synchronous operation, so there is no end event.
List Exascale DB Storage Vaults	com.oraclecloud.databaseservice.ListExascaleDbStorageVaults	This is synchronous operation, so there is no end event.
Update Exascale DB Storage Vault	com.oraclecloud.DatabaseService.UpdateExascaleDbStorageVault.begin	com.oraclecloud.DatabaseService.UpdateExascaleDbStorageVault.end
Delete Exascale DB Storage Vault	com.oraclecloud.DatabaseService.DeleteExascaleDbStorageVault.begin	com.oraclecloud.DatabaseService.DeleteExascaleDbStorageVault.end
ChangeCompartment Exascale DB Storage Vault	com.oraclecloud.DatabaseService.ChangeExascaleDbStorageVaultCompartment.begin	com.oraclecloud.DatabaseService.ChangeExascaleDbStorageVaultCompartment.end

This is a reference event for an Oracle Exadata Database Service on Exascale Infrastructure resource:

```
{
  "datetime": <date>,
  "logContent": {
    "data": {
      "additionalDetails": {
        "cpuCoreCount": 4,
        "dbNodeIds": "<DBNodeID>, <DBNodeID>",
        "exascaleDatabaseStorageVaultId": "<StorageVaultID>",
        "giVersion": "23.4.0.23.00",
```

```

    "licenseType": "LICENSE_INCLUDED",
    "lifecycleState": "TERMINATING",
    "localStorageInGbs": 586,
    "reservedCpuCoreCount": 4,
    "timeCreated": "2024-06-13T00:52:43Z",
    "timeUpdated": "2024-06-13T18:19:55Z",
    "timeZone": "UTC"
  },
  "availabilityDomain": "",
  "compartmentId": "ocidl.compartment.oc1<unique_ID>",
  "compartmentName": "<UniqueID>",
  "definedTags": {},
  "eventGroupingId": "/<ID>",
  "eventName": "GetExadbVmCluster",
  "freeformTags": {},
  "identity": {
    "authType": "natv",
    "callerId": null,
    "callerName": null,
    "consoleSessionId": null,
    "credentials": null,
    "ipAddress": "192.0.2.4",
    "principalId": "splat/<ID>",
    "principalName": "splat",
    "tenantId": "ocidl.tenancy.oc1<UniqueID>",
    "userAgent": "Jersey/2.38 (URLConnection 17.0.6)"
  },
  "message": "GetExadbVmCluster succeeded",
  "request": {
    "action": "GET",
    "headers": {},
    "id": "/<uniqueID>",
    "parameters": {},
    "path": "/20160918/exadbVmClusters/ocidl.<uniqueID>"
  },
  "resourceId": "ocidl.exadbvmcluster.oc1.<UniqueID>",
  "response": {
    "headers": {},
    "message": null,
    "payload": null,
    "responseTime": "2024-06-13T18:21:00.379Z",
    "status": "200"
  },
  "stateChange": {
    "current": {
      "cpuCoreCount": 4,
      "definedTags": {},
      "displayName": "audittest",
      "freeTags": {},
      "licenseType": "LICENSE_INCLUDED",
      "lifecycleState": "TERMINATING",
      "localStorageInGbs": 586,
      "reservedCpuCoreCount": 4,
      "sshPublicKeys": "..."
    },
    "previous": null
  }
}

```

```
    }  
  },  
  "dataschema": "2.0",  
  "id": "<uniqueID>",  
  "oracle": {  
    "compartmentid": "ocidl.compartment.oc1<UniqueID>",  
    "ingestedtime": "2024-06-13T18:21:06.462Z",  
    "loggroupid": "_Audit",  
    "tenantid": "ocidl.tenancy.oc1<UniqueID>"  
  },  
  "source": "audittest",  
  "specversion": "1.0",  
  "time": "2024-06-13T18:21:00.277Z",  
  "type": "com.oraclecloud.DatabaseService.GetExadbVmCluster"  
}  
}
```

### Related Topics

- [Overview of Events](#)
- [Notifications Overview](#)
- [ExadbVmClusterUpdate Reference](#)
- [ExascaleDbStorageVault Reference](#)
- [ExadbVmCluster Reference](#)

## Oracle Exadata Database Service on Exascale 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 will be desupported in a future release. The DB system event are not described.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure – Maintenance Scheduled	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancescheduled	<ul style="list-style-type: none"> <li data-bbox="1154 247 1469 709">• <b>Rolling:</b> 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 <i>&lt;infra-name&gt;</i>, <i>ocid &lt;infra-ocid&gt;</i> on <i>&lt;time-scheduled&gt;</i>. The maintenance method for this maintenance is <i>&lt;maintenance-method&gt;</i> as selected per the maintenance preferences.</li> <li data-bbox="1154 720 1469 1295">• <b>Non Rolling:</b> 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 <i>&lt;infra-name&gt;</i>, <i>ocid &lt;infra-ocid&gt;</i> on <i>&lt;time-scheduled&gt;</i>. The maintenance method for this maintenance is <i>&lt;maintenance-method&gt;</i> as selected per the maintenance preferences. Non-rolling maintenance minimizes maintenance time but will result in full system downtime.</li> </ul>



Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure – Maintenance Reminder	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancereminder	<ul style="list-style-type: none"> <li>• <b>Rolling:</b> This is an Oracle Cloud Operations reminder notice. Oracle has scheduled a quarterly maintenance update installation for Cloud Exadata Infrastructure &lt;infra-name&gt;, ocid &lt;ocid&gt; in approximately &lt;no-of-days&gt; days on &lt;time-scheduled&gt;. The maintenance method for this maintenance is &lt;maintenance-method&gt; as selected per the maintenance preferences.</li> <li>• <b>Non Rolling:</b> This is an Oracle Cloud Operations reminder notice. Oracle has scheduled a quarterly maintenance update installation for Cloud Exadata Infrastructure &lt;infra-name&gt;, ocid &lt;ocid&gt; in approximately &lt;no-of-days&gt; days on &lt;time-scheduled&gt;. The maintenance method for this maintenance is &lt;maintenance-method&gt; as selected per the maintenance preferences. Non-rolling maintenance minimizes maintenance time but will result in full system downtime.</li> </ul>
Cloud Exadata Infrastructure - Maintenance Begin	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.begin	<p>This is an Oracle Cloud Operations notice regarding the quarterly maintenance update installation for your Cloud Exadata Infrastructure instance &lt;infra-name&gt;, ocid &lt;infra-ocid&gt;. The update installation for the service started at &lt;time-scheduled&gt;.</p> <p>A follow-up notice will be sent when the maintenance update operation has completed.</p>
Cloud Exadata Infrastructure - Maintenance End Success	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.end.success	<p>This is an Oracle Cloud Operations notice that your Cloud Exadata Infrastructure quarterly maintenance update installation for service instance &lt;infra-name&gt;, ocid &lt;infra-ocid&gt; which started at &lt;maintenance-start-time&gt; is now successfully complete.</p>

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance End Failed	com.oraclecloud.databaservice.cloudexadatainfrastructuremaintenance.end.failed.	<p>This is an Oracle Cloud Operations notice that your Cloud Exadata Infrastructure quarterly maintenance update installation for service instance <i>&lt;infra-name&gt;</i>, ocid <i>&lt;infra-ocid&gt;</i> which started at <i>&lt;maintenance-start-time&gt;</i> has failed to complete due to technical reasons and operations team are currently looking into the issue.</p> <p>You will receive regular notifications to track progress of this maintenance.</p>
Cloud Exadata Infrastructure - Maintenance VM Begin	com.oraclecloud.databaservice.cloudexadatainfrastructuremaintenancevm.begin.	<p>This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of Virtual Machines component of your Cloud Exadata Infrastructure instance <i>&lt;infra-name&gt;</i>, ocid <i>&lt;infra-ocid&gt;</i>, Database Server <i>&lt;dbserver name&gt;</i>, ocid <i>&lt;dbserver ocid&gt;</i> has started.</p> <p>A follow-up notice will be sent when Virtual Machines maintenance operation has completed.</p>
Cloud Exadata Infrastructure - MaintenanceVM End	com.oraclecloud.databaservice.cloudexadatainfrastructuremaintenancevm.end	<p>This is an Oracle Cloud Operations notice that quarterly maintenance update of the Database Server component of your Cloud Exadata Infrastructure instance <i>&lt;infra-name&gt;</i>, ocid <i>&lt;infra-ocid&gt;</i>; Database Server <i>&lt;dbserver name&gt;</i> ocid <i>&lt;dbserver ocid&gt;</i> has completed.</p>
Cloud Exadata Infrastructure - Maintenance Storage Servers Start	com.oraclecloud.databaservice.cloudexadatainfrastructuremaintenancestorageservers.start	<p>This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of Storage servers component of your Cloud Exadata Infrastructure instance <i>&lt;infra-name&gt;</i>, ocid <i>&lt;infra-ocid&gt;</i> has started.</p> <p>A follow-up notice will be sent when storage servers maintenance operation has completed.</p>
Cloud Exadata Infrastructure - Maintenance Storage Servers End	com.oraclecloud.databaservice.cloudexadatainfrastructuremaintenancestorageservers.end	<p>This is an Oracle Cloud Operations notice that quarterly maintenance update of Storage servers component of your Cloud Exadata Infrastructure instance <i>&lt;infra-name&gt;</i>, ocid <i>&lt;infra-ocid&gt;</i> has completed.</p>

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance Network Switches Begin	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancenetworkswitches.begin	<p>This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of the network fabric switches component of your Cloud Exadata Infrastructure instance <i>&lt;infra-name&gt;</i>, <i>ocid &lt;infra-ocid&gt;</i> has started.</p> <p>A follow-up notice will be sent when the network fabric switches maintenance operation has completed.</p>
Cloud Exadata Infrastructure - Maintenance Network Switches End	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancenetworkswitches.end	<p>This is an Oracle Cloud Operations notice that quarterly maintenance update of the network fabric switches component of your Cloud Exadata Infrastructure instance <i>&lt;infra-name&gt;</i>, <i>ocid &lt;infra-ocid&gt;</i> has completed.</p>
Cloud Exadata Infrastructure - Maintenance Custom Action Time Begin	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancencustomactiontime.begin	<p>This is an Oracle Cloud Operations notice that the custom action timeout for your Cloud Exadata Infrastructure instance <i>&lt;infra-name&gt;</i>, <i>ocid &lt;infra-ocid&gt;</i>; Database Server <i>&lt;dbserver name&gt;</i>, <i>ocid &lt;dbserver ocid&gt;</i> has started.</p> <p>A follow-up notice will be sent when the custom action timeout has ended.</p>
Cloud Exadata Infrastructure - Maintenance Custom Action Time End	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancencustomactiontime.end	<p>This is an Oracle Cloud Operations notice that the custom action timeout for your Cloud Exadata Infrastructure instance <i>&lt;infra-name&gt;</i>, <i>ocid &lt;infra-ocid&gt;</i>; Database Server <i>&lt;dbserver name&gt;</i>, <i>ocid &lt;dbserver ocid&gt;</i> has ended.</p>
Cloud Exadata Infrastructure - Maintenance Rescheduled	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancerescheduled	<p>Oracle Cloud Operations is announcing reschedule of a quarterly maintenance update for Cloud Exadata Infrastructure.</p> <p>A maintenance run has been rescheduled on your service instance <i>&lt;infra-name&gt;</i>, <i>ocid &lt;infra-ocid&gt;</i> to <i>&lt;new-schedule-time&gt;</i>.</p>

Friendly Name	Event Type	Event Messages
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 <new-maintenance-method>.

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": "ocid1.compartment.oc1.<unique_ID>"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
    "compartmentName": "example_name",
    "resourceName": "my_exadata_infrastructure",
    "resourceId": "ocid1.dbsystem.oc1.eu-frankfurt-1.<unique_ID>",
    "availabilityDomain": "tXPJ:EU-FRANKFURT-1-AD-3",
    "freeFormTags": {
      "Department": "Finance"
    },
    "definedTags": {
      "Operations": {
        "CostCenter": "42"
      }
    },
    "additionalDetails" : {
      "subnetId" : "ocid1.subnet.oc1.eu-frankfurt-1.<unique_ID>",
      "lifecycleState" : "MAINTENANCE_IN_PROGRESS",
      "sshPublicKeys" : "...",
      "cpuCoreCount" : 32,
      "version" : "19.2.8.0.0.191119",
      "nsgIds" : "null",
      "backupSubnetId" : "ocid1.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>"
}
}

```

## 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"> <li>\"set serveroutput off\"</li> <li>\"select sql_id, sql_text from v\$sqltext where sql_id =&lt;sqlId&gt;\"</li> <li>\"set serveroutput on\"</li> </ul>

#### 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.exadatainfrastructureservice.critical",
  "cloudEventsVersion" : "0.1",
  "eventTypeVersion" : "2.0",
  "source" : "Exadata Storage",
  "eventTime" : "2021-07-30T04:53:18Z",
  "contentType" : "application/json",
  "data" : {
    "compartmentId" : "ocidl.tenancy.oc1.<unique_ID>",
    "compartmentName" : "example_name",
    "resourceName" : "my_exadata_resource",
    "resourceId" : "ocidl.dbssystem.oc1.phx.<unique_ID>",
    "availabilityDomain" : "phx-ad-2",
    "additionalDetails" : {
      "serviceType" : "exacs",
      "sqlID" : "gnwfm1jgqcfuu",
      "systemId" : "ocidl.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" : "ocidl.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..aaaaaaaa03lj36x6lwxyvc4wausjouca7pwyjfw5ebsq5emrpqlq12gj5iq",
    "compartmentName" : "intexadatateam",
    "resourceId" :
    "ocid1.dbssystem.oc1.phx.abyhqljt5y3taezn7ug445fzwlngjfszbedxlcctw45ykkaxyzc5isxoula",
    "availabilityDomain" : "phx-ad-2",
    "additionalDetails" : {
      "serviceType" : "exacs",
      "component" : "storage",
      "systemId" :
      "ocid1.dbssystem.oc1.phx.abyhqljt5y3taezn7ug445fzwlngjfszbedxlcctw45ykkaxyzc5isxoula",
      "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..aaaaaaaa03lj36x6lwxyvc4wausjouca7pwyjfw5ebsq5emrpqlq12gj5iq"
  }
}
```



## 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
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.updatecloudvmclusteriorconfig.begin
Cloud VM Cluster - Update IORM Configuration End	com.oraclecloud.databaseservice.updatecloudvmclusteriorconfig.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.updatecloudvmclusteriorconfig.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":
"ocid1.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": "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": "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.v87pkelz9k9u6xaqo51taf6bunf0gc2wyhrbmjzbh3h1pjawakav
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": {
    "eventGroupId": "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",
    "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.ekmz1phzp4b11k7m7tbygulbnakmjnr99eqjobs3zvp337pn
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": "ocidl.cloudvmcluster.oc1.....unique_id",
    "resourceVersion": null,
    "additionalDetails": {
      "cloudExadataInfrastructureId":
"ocidl.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.svwkildsx63clp1q6phba7d6lns1r192yc3uyc2ea5utjprqcu
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":
"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"
}

```

## VM Node Subsetting Event Types

Review the list of event types that VM Node Subsetting emits.

**Table 6-2 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 6-62 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.oc1..unique_ID"
    },
    "data": {
"compartmentId": "ocidl.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Vmcluster-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
" id": "ocidl.id.oc1..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.oc1..unique_ID"
  },
  "data": {
"compartmentId": "ocidl.compartment.oc1..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 - 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": "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 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": {
"eventId": "ocid1.id..oc1...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"
    }
}
}

```

## Data Guard Association Event Types

Review the list of event types that Data Guard associations emit.

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 Create Standby Database - Create Begin	com.oraclecloud.databaseservice.creates tandbydatabase.begin
Data Guard Create Standby Database - Create End	com.oraclecloud.databaseservice.creates tandbydatabase.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.updated ataguardconfig.begin
Data Guard Update Config - End	com.oraclecloud.databaseservice.updated ataguardconfig.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

Friendly Name	Event Type
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

Friendly Name	Event Type
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
Database - Critical	com.oraclecloud.databaseservice.databas e.critical
Database - Information	com.oraclecloud.databaseservice.databas e.information
Database - Delete Backup Begin	com.oraclecloud.databaseservice.deleteb ackup.begin
Database - Delete Backup End	com.oraclecloud.databaseservice.deleteb ackup.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.movedat abase.begin
Database - Move End	com.oraclecloud.databaseservice.movedat abase.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.rotated atabasekmskey.begin
Database - Rotate KMS Key End	com.oraclecloud.databaseservice.rotated atabasekmskey.end
Database - Terminate Begin	com.oraclecloud.databaseservice.databas e.terminate.begin
Database - Terminate End	com.oraclecloud.databaseservice.databas e.terminate.end
Database - Update Begin	com.oraclecloud.databaseservice.updated atabase.begin
Database - Update End	com.oraclecloud.databaseservice.updated atabase.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.remotecopypluggabledatabase.begin

Friendly Name	Event Type
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
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": "ocid1.compartment.oc1.<unique_ID>"
  },
  "eventType":
"com.oraclecloud.databaseservice.remoteclonepluggabledatabase.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",
  "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": "ocid1.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": "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",
      "isRefreshableClone": true,
      "timeCreated": "2021-03-13T21:15:59.000Z",
      "timeUpdated": "2021-03-13T21:15:59.000Z",
      "databaseId": "ocid1.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": "ocid1.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": "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",
        "isRefreshableClone": false,
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.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": "ocid1.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": "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": "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": "ocid1.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": "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 - 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": "ocid1.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": "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",
      "isRefreshableClone": true,
      "databaseId": "ocid1.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": "ocid1.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": "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": 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": "ocid1.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": "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",
      "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.

- [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.

## 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 Oracle Exadata Database Service on Exascale 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
\u25cf syslens.service
Loaded: loaded (/etc/systemd/system/syslens.service; disabled; vendor
preset: disabled)
Active: active (running) since Wed 2022-03-16 18:08:59 UTC; 34s ago
Main PID: 358039 (python3)
Memory: 31.6M
CGroup: /system.slice/syslens.service
\u2514\u2500358039 /usr/bin/python3 /var/opt/oracle/syslens/bin/
syslens_main.py --archive /var/opt/oracle/log/...

Mar 16 18:08:59 node1 systemd[1]: Started syslens.service.
Mar 16 18:09:09 node1 su[360495]: (to oracle) root on none
Mar 16 18:09:09 node1 su[360539]: (to grid) root on none
Mar 16 18:09:10 node1 su[360611]: (to grid) root on none
Mar 16 18:09:11 node1 su[360653]: (to oracle) root on none
```

- 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 configured to notify events, run the `/usr/bin/syslens --config /var/opt/oracle/syslens/data/exacc.syslens.config --get-key enable_telemetry` command as root in the domU:

```
# /usr/bin/syslens --config /var/opt/oracle/syslens/data/
exacc.syslens.config --get-key enable_telemetry
syslens Collection 2.3.3
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

- [Using the Console 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](#)

## 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
- DB System - Critical

### Related Topics

- [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	Event Type	Threshold
Resource Utilization - Disk Usage	HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE	com.oraclecloud.dat abaseservice.dbnode .critical	Critical Threshold: 90%
	<p>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:</p> <ul style="list-style-type: none"> <li>• /</li> <li>• /u01</li> <li>• /u02</li> <li>• /var (X8M and later only)</li> <li>• /tmp (X8M and later only)</li> </ul>		
CRS status Up/Down	AVAILABILITY.DB_GUE ST.CRS_INSTANCE.DOW N.	com.oraclecloud.dat abaseservice.dbnode .critical (if .DOWN and NOT "user_action")	N/A
	AVAILABILITY.DB_GUE ST.CRS_INSTANCE.DOW N_CLEARED	com.oraclecloud.dat abaseservice.dbnode .information (if .DOWN_CLEARED)	N/A
	<p>An event of type <b>CRITICAL</b> is created when the Cluster Ready Service (CRS) is detected to be down.</p> <p>An event of type <b>INFORMATION</b> is created once it is determined that the event for CRS down has cleared.</p>		



Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Event Type	Threshold
SCAN Listener Up/Down	AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.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]	com.oraclecloud.dat abaseservice.dbnode .critical (if .DOWN and NOT "user_action")	N/A
	AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN_CLEARED An event of type INFORMATION is created once it is determined that the event for SCAN Listener down has cleared.	com.oraclecloud.dat abaseservice.dbnode .information (if .DOWN_CLEARED)	N/A

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Event Type	Threshold
Net Listener Up/Down	AVAILABILITY.DB_GUE ST.CLIENT_LISTENER. DOWN	com.oraclecloud.dat abaseservice.databa se.critical (if .DOWN and NOT "user_action")	N/A
	<p>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.</p> <p>There is one client listener per node, each called LISTENER.</p>		
	AVAILABILITY.DB_GUE ST.CLIENT_LISTENER. DOWN_CLEARED	com.oraclecloud.dat abaseservice.databa se.information (if .DOWN_CLEARED)	N/A
	<p>An event of type INFORMATION is created once it is determined that the event for Client Listener down has cleared.</p>		

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Event Type	Threshold
CDB Up/Down	AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN	com.oraclecloud.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_INSTANCE.DOWN_CLEARED	com.oraclecloud.database.information (if .DOWN_CLEARED)	N/A
	An event of type INFORMATION is created once it is determined that the event for the CDB down has cleared.		
CRS Eviction	AVAILABILITY.DB_GUEST.CRS_INSTANCE.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

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Event Type	Threshold
Critical DB Errors	HEALTH.DB_CLUSTER.C DB.CORRUPTION	com.oraclecloud.dat abaseservice.databa se.critical	N/A
	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.		
Other DB Errors	HEALTH.DB_CLUSTER.C DB.ARCHIVER_HANG	com.oraclecloud.dat abaseservice.databa se.critical	N/A
	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_CLUSTER.C DB.DATABASE_HANG	N/A	N/A
	An event of type CRITICAL is created when a process/session hang is detected in the CDB.		
Backup Failures	HEALTH.DB_CLUSTER.C DB.BACKUP_FAILURE	com.oraclecloud.dat abaseservice.databa se.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.D ISK_GROUP.FREE_SPAC E	com.oraclecloud.dat abaseservice.dbsyst em.critical	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%.	com.oraclecloud.dat abaseservice.dbsyst em.information (if < 90%)	

**Table 6-3 (Cont.) Database Service Events**

Friendly Name	Event Name	Event Type	Threshold
Memory Usage	CONFIGURATION.DB_GUEST.MEMORY.HUGE_PAGES_TOO_LARGE	com.oraclecloud.dat abaseservice.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.		
sshd Configuration	CONFIGURATION.DB_GUEST.SSHD.INVALID_CONFIG	com.oraclecloud.dat abaseservice.dbnode .critical	N/A
	An event of type CRITICAL is created if unexpected values are set in the /etc/ssh/sshd_config file.		
Disk Issues	HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION	com.oraclecloud.dat abaseservice.dbnode .critical	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"> <li>• /</li> <li>• /u01</li> <li>• /u02</li> </ul>		

**Table 6-3 (Cont.) Database Service Events**

Friendly Name	Event Name	Event Type	Threshold
Oracle EXAchk Reported Issues	HEALTH.DB_CLUSTER.E XACHK.CRITICAL_ALER T	com.oraclecloud.dat abaseservice.dbnode .critical	N/A
	<p>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.</p>		

**Example 6-63 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" : "SYSLENS/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"
```

```
}  
}
```

## 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|  
pdb_tablespace|pdb|listener|service|os  
-target all|name  
[-container name]  
[-pdb pdb_name]  
-event all|"event_str1,event_str2"|availability  
[-timeout nm|nh|nd|none]  
[-c|-local|-nodes "node1,node2"]  
[-reason "reason for blackout"]  
[-docollection]
```

### Parameters

**Table 6-4** `tfactl blackout` Command Parameters

Parameter	Description
add remove print	Adds, removes, or prints blackout conditions.

**Table 6-4 (Cont.) tfactl blackout Command Parameters**

Parameter	Description
targettype <i>type</i>	Limits blackout only to the specified target type.
<b>Target type:</b> host crs asm  asmdg database dbbackup  db_dataguard  db_tablespace   pdb_tablespace pdb  listener service os	<p>host: The whole node is under blackout. If there is host blackout, then every blackout element that's shown true in the Telemetry JSON will have the reason for the blackout.</p> <p>crs: Blackout the availability of the Oracle Clusterware resource or events in the Oracle Clusterware logs.</p> <p>asm: Blackout the availability of Oracle Automatic Storage Management (Oracle ASM) on this machine or events in the Oracle ASM alert logs.</p> <p>asmdg: Blackout an Oracle ASM disk group.</p> <p>database: Blackout the availability of an Oracle Database, Oracle Database backup, tablespace, and so on, or events in the Oracle Database alert logs.</p> <p>dbbackup: Blackout Oracle Database backup events (such as CDB or archive backups).</p> <p>db_dataguard: Blackout Oracle Data Guard events.</p> <p>db_tablespace: Blackout Oracle Database tablespace events (container database).</p> <p>pdb_tablespace: Blackout Oracle Pluggable Database tablespace events (Pluggable database).</p> <p>pdb: Blackout Oracle Pluggable Database events.</p> <p>listener: Blackout the availability of a listener.</p> <p>service: Blackout the availability of a service.</p> <p>os: Blackout one or more operating system records.</p>
target all  <i>name</i>	<p>Specify the target for blackout. You can specify a comma-delimited list of targets.</p> <p>By default, the target is set to all.</p>
container <i>name</i>	Specify the database container name ( <i>db_unique_name</i> ) where the blackout will take effect (for PDB, DB_TABLESPACE, and PDB_TABLESPACE).
pdb <i>pdb_name</i>	Specify the PDB where the blackout will take effect (for PDB_TABLESPACE only).
events all "str1,str2"	<p>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.</p> <p>all: Blackout everything for the target specified.</p> <p><i>string</i>: Blackout for incidents where any part of the line contains the strings specified.</p> <p>Specify a comma-delimited list of strings.</p>
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	<p>Specify if blackout should be set to cluster-wide or local.</p> <p>By default, blackout is set to local.</p>
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 6-64 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
```

## 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
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":
"ocidl.eventschema.oc1.phx.5ur5er8bddumnu9r84rtt2c3282s5no31vsthibyqvvisotnwp
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": "ocidl.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.m2gheil6f1nfbz9ggpkkv17wdomdks8zin9nntqlghui6bckh17
  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": "ocid1.applicationvip.oc1.....unique_id",
    "resourceName": "my_application_vip",
    "tagSlug": null,
    "compartmentId": "ocid1.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": "ocid1.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": "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 End:

```

{
  "id":
"ocid1.eventschema.oc1.phx.9d1tjgkavhn0rq4qdlmofrjro9npvugu73dp07uht0igxs9732x
6yar1m5l5",
  "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": "17619cal-07ae-4e2d-a818-5b5f1fcd4f70",
    "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": "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/
b3e19f76a81549e6b7bf1d8619f7c191/C683214FCB0BF3CEC1C8B23C2FEE983E",
  "eventName": "DeleteApplicationVip",
  "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": "TERMINATED",
    "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"
}

```

## 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

Friendly Name	Event Type
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":
  "ocid1.eventschema.oc1.phx.abyhqljrsl1p7rfneajgq2knxbqopwux24za7qzoe3mfj2bzfxtnwqcxpbcq",
  "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": "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 - 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":
  "ocid1.eventschema.oc1.phx.abyhqljrj4vuvph4qvj5eateeel6axblhkq3caqndgmjvw13sld
  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": "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 - 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":
  "ocid1.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": "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 - 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":
```

```

"ocid1.eventschema.oc1.phx.abyhqljrgwk2gvx5lmx6fiwotgdy32mdmrnkyzsnz37dpb4mmeh
gzt37vl7a",
  "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": "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 - 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":
"ocid1.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":
"ocidl.eventschema.oc1.phx.abyhqljrn2lruez55ah56kqksi5qfg6m7igvven7o2qkahlk5tk
wrj51l3oa",
  "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": "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 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.created bnodeconsoleconnection.begin
DB Node Console Connection - Create End	com.oraclecloud.databaseservice.created bnodeconsoleconnection.end
DB Node Console Connection - Delete Begin	com.oraclecloud.databaseservice.deleted bnodeconsoleconnection.begin

**Table 6-5 (Cont.) Serial Console Connection Events**

Friendly Name	Event Type
DB Node Console Connection - Delete End	com.oraclecloud.databaseservice.deleted bnodeconsoleconnection.end
DB Node Console Connection - Update	com.oraclecloud.databaseservice.updated bnodeconsoleconnection
DB Node - Update	com.oraclecloud.databaseservice.updated bnode

**Example 6-65 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"
    }
  }
}
}

```

- [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 6-66 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>aaaaaxxxxx",
    "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](#)

## 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 **Exadata on Oracle Public Cloud**.
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.
  - 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 (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*.
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 in the bare metal and virtual machine DB systems.
  - Private endpoint for Oracle RAC Databases in the virtual machine DB system.

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 **Exadata on Oracle Public Cloud**.
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 in the bare metal and virtual machine DB systems.
  - Private endpoint for Oracle RAC Databases in the virtual machine DB system.

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 **Exadata on Oracle Public Cloud**.
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.
  - 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 (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*.
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 in the bare metal and virtual machine DB systems.
  - Private endpoint for Oracle RAC Databases in the virtual machine DB system.

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 Disable Database Management for a Pluggable Database (PDB)

To disable Database Management for a pluggable database (PDB), use this procedure.

1. Open the navigation menu. Click **Oracle Database**, then click **Exadata on Oracle Public Cloud**.
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 **Exadata on Oracle Public Cloud**.
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 **Exadata on Oracle Public Cloud**.
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 **Exadata on Oracle Public Cloud**.

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 **Exadata on Oracle Public Cloud**.
  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.

## Policy Details for Oracle Exadata Database Service on Exascale Infrastructure

This topic covers details for writing policies to control access to Oracle Exadata Database Service on Exascale Infrastructure resources.

 **Note:**

For more information on Policies, see "How Policies Work".

For a sample policy, see "Let database admins manage Oracle Exadata Database Service on Exascale Infrastructure instances".

- [About Resource-Types](#)  
Learn about resource-types you can use in your policies.
- [Resource-Types for Exadata Cloud Service Instances](#)  
Instance resource types include aggregate resource types and individual resource types.
- [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 Oracle Exadata Database Service on Exascale 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

Instance resource types include aggregate resource types and individual resource types.

### Aggregate Resource-Type

`database-family`

### Individual Resource-Types

`db-nodes`

`db-homes`

`databases`

`pluggable-databases`

`db-backups`

`dbnode-console-connection`

## Supported Variables

Use variables when adding conditions to a policy.

Oracle Exadata Database Service on Exascale 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 DB Backups](#)
- [Permissions and API operation details for Databases \(CDBs\)](#)
- [Permissions and API operation details for Data Guard Association](#)
- [Permissions and API operation details for DB Nodes](#)
- [Permissions and API operation details for DB Homes](#)

- [Permissions and API operation details for Database Software Image](#)
- [exadb-vm-clusters](#)  
Review the list of permissions and API operations for the `exadb-vm-clusters` resource-type.
- [exascale-db-storage-vaults](#)  
Review the list of permissions and API operations for the `exascale-db-storage-vaults` resource-type.
- [Permissions and API operation details for Key Stores](#)
- [Permissions Required for Each API Operation](#)
- [Permissions and API operation details for Pluggable Databases \(PDBs\)](#)

#### 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 DB Backups

The table below lists permissions and API operations for `db-backups`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
<code>inspect</code>	<code>DB_BACKUP_INSPECT</code>	<code>GetBackup</code> <code>ListBackups</code>	<code>ChangeCloudVmClusterCompartments</code> (also <b>needs</b> <code>use cloud-vmclusters</code> , <code>use db-homes</code> , and <code>use databases</code> )
<code>read</code>	<i>INSPECT +</i> <code>DB_BACKUP_CONTENT_READ</code>	<i>none</i>	<code>RestoreDatabase</code> (also <b>needs</b> <code>use databases</code> )
<code>use</code>	<i>no extra</i>	<i>no extra</i>	<i>none</i>
<code>manage</code>	<i>USE +</i> <code>DB_BACKUP_CREATE</code> <code>DB_BACKUP_DELETE</code>	<code>DeleteBackup</code>	<code>CreateBackup</code> (also <b>needs</b> <code>read databases</code> )

## 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 ListDataGuardAssociations GetDataGuardAssociation	enableDatabaseManagement disableDatabaseManagement updateDatabaseManagement
read	<i>INSPECT+</i> DATABASE_CONTENT_READ	<i>no extra</i>	<i>no extra</i>
use	<i>READ +</i> DATABASE_CONTENT_WRITE DATABASE_UPDATE	UpdateDatabase SwitchoverDataGuardAssociation FailoverDataGuardAssociation ReinstateDataGuardAssociation	CreateDataGuardAssociation ChangeCloudVmClusterCompartment (also needs use cloud-vmclusters, use db-homes, and inspect db-backups) enableDatabaseManagement disableDatabaseManagement updateDatabaseManagement
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 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	<b>READ +</b> VM_CLUSTER_UPDATE + DB_HOME_UPDATE DATABASE_UPDATE	DeleteDatabase SwitchoverDataGuardAssociation, FailoverDataGuardAssociation, ReinstateDataGuardAssociation	CreateDataGuardAssociation
MANAGE	<b>USE +</b> DATABASE_DELETE	DeleteDatabase	<i>none</i>

## Permissions and API operation details for DB Nodes



### Note:

For Oracle Exadata Database Service on Exascale 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	<b>USE +</b> DB_NODE_POWER_ACTIONS	DbNodeAction	<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 ListDbHomePatchHistoryEntries GetDbHomePatch GetDbHomePatchHistoryEntry	<i>none</i>
read	<i>no extra</i>	<i>no extra</i>	<i>none</i>
use	DB_HOME_UPDATE	UpdateDBHome	ChangeCloudVmClusterCompartment ( <b>also needs</b> use cloud-vmclusters, use databases, and inspect backups)
manage	USE + DB_HOME_CREATE DB_HOME_DELETE	<i>no extra</i>	CreateCloudVmCluster, DeleteCloudVmCluster ( <b>both also need</b> manage cloud-vmclusters, manage databases, use vnics, and use subnets). <b>If automatic backups are enabled on the default database, also needs</b> manage backups  CreateDbHome, ( <b>also needs</b> use cloud-vmclusters and manage databases). <b>If creating the Database Home by restoring from a backup, also needs</b> read backups  DeleteDbHome, ( <b>also needs</b> use cloud-vmclusters and manage databases). <b>If automatic backups are enabled on the default database, also needs</b> manage backups. <b>If the performFinalBackup option is selected, also needs</b> manage backups and read databases.

## 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	<b>READ +</b> DB_SOFTWARE_IMG_UPDATE	UpdateDatabaseSoftwareImage ChangeDatabaseSoftwareImageCompartment	<i>none</i>
manage	<b>USE +</b> DB_SOFTWARE_IMG_CREATE DB_SOFTWARE_IMG_DELETE	CreateDatabaseSoftwareImage DeleteDatabaseSoftwareImage	<i>none</i>

## exadb-vm-clusters

Review the list of permissions and API operations for the `exadb-vm-clusters` resource-type.

**Table 6-6 INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
EXADB_VM_CLUSTER_INSPECT	ListExadbVmClusters GetExadbVmCluster ListExadbVmClusterUpdates GetExadbVmClusterUpdate ListExadbVmClusterUpdateHistoryEntries GetExadbVmClusterUpdateHistoryEntry	None

**Table 6-7 READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>No extra</i>	<i>No extra</i>	None

**Table 6-8 USE**

Permissions	APIs Fully Covered	APIs Partially Covered
inspect + EXADB_VM_CLUSTER_UPDATE	RemoveVirtualMachineFromExadbVmClusterDetails	ChangeExadbVmClusterComparison  (also needs use db-homes, use databases, and inspect db-backups)

**Table 6-9 MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
use + EXADB_VM_CLUSTER_CREATE, EXADB_VM_CLUSTER_DELETE	<i>No extra</i>	CreateExadbVmCluster (also needs manage db-homes, manage databases, use exascale-db-storage-vaults, use vnics, and use subnets)  DeleteExadbVmCluster (also needs manage db-homes, manage databases, use exascale-db-storage-vaults, use vnics, and use subnets)  UpdateExadbVmCluster (also needs use subnets, use vnics, and use private-ip)

## exascale-db-storage-vaults

Review the list of permissions and API operations for the `exascale-db-storage-vaults` resource-type.

**Table 6-10 INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
EXASCALE_DB_STORAGE_VAULT_INSPECT	ListExascaleDbStorageVaults  GetExascaleDbStorageVault	None

**Table 6-11 READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>No extra</i>	<i>No extra</i>	None

**Table 6-12 USE**

Permissions	APIs Fully Covered	APIs Partially Covered
inspect + EXASCALE_DB_STORAGE_VAULT_ UPDATE	ChangeExascaleDbStorageVau ltCompartment UpdateExascaleDbStorageVau lt	None

**Table 6-13 MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
use + EXASCALE_DB_STORAGE_VAULT_ CREATE EXASCALE_DB_STORAGE_VAULT_ DELETE	CreateExascaleDbStorageVau lt DeleteExascaleDbStorageVau lt	None

## 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 AUTONOMOUS_CONTAINE R_DATABASE_INSPECT AUTONOMOUS_DATABASE _INSPECT AUTONOMOUS_DB_BACKU P_INSPECT	GetKeyStore GetAutonomousContai nerDatabase GetAutonomousDataba se GetAutonomousDataba seBackup	ChangeKeyStoreCompa rtment RotateAutonomousCon tainerDatabaseKey
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>
USE	<b>READ +</b> KEY_STORE_UPDATE + AUTONOMOUS_VM_CLUST ER_UPDATE + AUTONOMOUS_CONTAINE R_DATABASE_UPDATE AUTONOMOUS_DATABASE _UPDATE	UpdateKeyStore <i>none</i> <i>none</i> RotateAutonomousDat abaseKey	ChangeKeyStoreCompa rtment CreateAutonomousCon tainerDatabase RotateAutonomousCon tainerDatabaseKey <i>none</i>
MANAGE	<b>USE +</b> KEY_STORE_CREATE + KEY_STORE_DELETE + AUTONOMOUS_CONTAINE R_DATABASE_CREATE	CreateKeyStore DeleteKeyStore CreateAutonomousCon tainerDatabase	<i>none</i>



## 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 6-14 Cloud Exadata Infrastructure Resource**

API Operation	Permissions Required to Use the Operation
ListCloudExadataInfrastructures	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
GetCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
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

**Table 6-15 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)

**Table 6-15 (Cont.) Cloud VM Cluster**

API Operation	Permissions Required to Use the Operation
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 6-16 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 6-17 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 6-18 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

**Table 6-18 (Cont.) Database Homes**

API Operation	Permissions Required to Use the Operation
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 6-19 Databases (CDB)**

API Operation	Permissions Required to Use the Operation
ListDatabases	DATABASE_INSPECT
GetDatabase	DATABASE_INSPECT
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 6-20 Pluggable Databases (PBDs)**

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

**Table 6-20 (Cont.) Pluggable Databases (PBDs)**

API Operation	Permissions Required to Use the Operation
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 6-21 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 6-22 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 6-23 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

**Table 6-23 (Cont.) Backups and Database Restore**

API Operation	Permissions Required to Use the Operation
RestoreDatabase	DB_BACKUP_INSPECT and DB_BACKUP_CONTENT_READ and DATABASE_CONTENT_WRITE

**Table 6-24 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 6-25 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

## 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	UpdatePluggableDatabase
		GetPluggableDatabase	StartPluggableDatabase StopPluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase RefreshPluggableDatabase ConvertRefreshablePluggableDatabase
	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

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
	PLUGGABLE_DATABASE_ UPDATE	<i>no extra</i>	UpdatePluggableData base StartPluggableDatab ase StopPluggableDataba se LocalClonePluggable Database RemoteClonePluggabl eDatabase RefreshPluggableDat abase ConvertRefreshableP luggableDatabase
	DATABASE_UPDATE	<i>no extra</i>	CreatePluggableData base DeletePluggableData base LocalClonePluggable Database RemoteClonePluggabl eDatabase
manage	<b>USE +</b> PLUGGABLE_DATABASE_ CREATE	<i>no extra</i>	CreatePluggableData base LocalClonePluggable Database RemoteClonePluggabl eDatabase
	PLUGGABLE_DATABASE_ DELETE	<i>no extra</i>	DeletePluggableData base

## 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 Oracle Cloud (Bare Metal, Virtual Machine VM, and 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](#)

## 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 Exascale Infrastructure databases and databases running on virtual machine and bare metal systems.

### 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 Exascale Infrastructure databases and databases running on virtual machine and bare metal systems.

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 Oracle Exadata Database Service on Exascale Infrastructure, Virtual Machine, and Bare Metal databases, 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. For information about using Database Metrics and Performance Hub with Virtual Machine, Bare Metal, Oracle Exadata Database Service on Exascale Infrastructure and external databases, see [Enable Database 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, databases running on virtual machine and bare metal systems, Oracle Database Cloud Service, Oracle Exadata Database Service on Exascale Infrastructure, and external databases. For information about IAM policies and ExaDB-XS databases, see *Required IAM Policy*. For information about policies and how to use them, see [How Policies Work](#).

**Related Topics**

- [Required IAM Policy for Oracle Exadata Database Service on Exascale Infrastructure](#)  
Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Exascale Infrastructure systems.

## Troubleshooting Oracle Exadata Database Service on Exascale Infrastructure Systems

These topics cover some common issues you might run into and how to address them.

- [Known Issues for Exadata Database Service on Exascale Infrastructure](#)  
General known issues.
- [Troubleshooting Oracle Data Guard](#)  
Learn to identify and resolve Oracle Data Guard issues.
- [Obtaining Further Assistance](#)

### Known Issues for Exadata Database Service on Exascale Infrastructure

General known issues.

### 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 reinstate. 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](#)  
 Review errors that can 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

Review errors that can 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 stage of enablement. In this stage, configurations are performed on the standby database, including the rman duplicate of the primary database. To resolve this issue:
  1. Verify the connectivity status for the primary and standby sites.
  2. Ensure that the host can communicate from port 1521 to all ports. Check the network setup, including Network Security Groups (NSGs), Network Security Lists, and the remote VCN peering setup (if applicable). The best way to test communication between the host and other nodes is to access the databases using SQL\*PLUS from the primary to standby and from the standby to the primary.
- The SCAN VIPs or listeners may not be running. Use the test above to help identify the issue.

### NOT\_SUPPORTED

Possible causes:

- SCAN VIPs or listeners may not be running. You can confirm this issue by using the following commands on any cluster node.

```
- [grid@exal-***** ~]$ srvctl status  
scan
```

```
- [grid@exal-***** ~]$ srvctl status  
scan_listener
```

- Databases may not be reachable. You can confirm this issue by attempting to connect using an existing Oracle Net alias.

Troubleshooting procedure:

1. As the oracle OS user, check for the existence of an Oracle Net alias for the container database (CDB). Look for an alias in \$ORACLE\_HOME/network/admin/<dbname>/tnsnames.ora.

The following example shows an entry for a container database named db12c:

```
cat $ORACLE_HOME/network/admin/db12c/tnsnames.ora  
DB12C = (DESCRIPTION = (ADDRESS = (PROTOCOL = TCP) (HOST = exal-*****-  
scan.*****.*****.*****.com) (PORT = 1521))  
(CONNECT_DATA = (SERVER = DEDICATED) (SERVICE_NAME =
```

```
db12c.*****.*****.*****.com)
(FAILOVER_MODE = (TYPE = select) (METHOD = basic)))
```

2. Verify that you can use the alias to connect to the database. For example, as sysdba, enter the following command:

```
sqlplus sys@db12c
```

### NOT\_SUPPORTED

A possible cause for this error is that the Oracle Database sys or system user passwords for the database and the TDE wallet may not be the same. To compare the passwords:

1. Connect to the database as the **sys user** and check the TDE status in

```
V$ENCRYPTION_WALLET
.
.
```

2. Connect to the database as the **system user** and check the TDE status in

```
V$ENCRYPTION_WALLET
.
.
```

3. Update the applicable passwords to match. Log on to the system host as **opc** 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

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>
DGMGR> VALIDATE CONFIGURATION VERBOSE
DGMGR> VALIDATE DATABASE VERBOSE <PRIMARY>
DGMGR> 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)

## 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](#)