Oracle® Cloud

Exadata Database Service on Exascale Infrastructure

F92371-68
November 2026

Oracle Cloud Exadata Database Service on Exascale Infrastructure,

F92371-68

Copyright © 2023, 2026, Oracle and/or its affiliates.

Primary Authors: Nirmal Kumar, Doug Williams, Tejas N K

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, Bhuvaneswari Thiagarajan, Sujeet Vasudevan, Nayana Vishwa, Ravi Wijayaratne

Contents

1 Oracle Exadata Database Service on Exascale Infrastructure Overview

About Oracle Exadata Database Service on Exascale Infrastructure 1
Accessing the Oracle Exadata Database Service on Exascale Infrastructure Using the OCI Console 6
Licensing Considerations for Oracle Exadata Database Service on Exascale Infrastructure 6
Supported Database Edition and Versions for Oracle Exadata Database Service on Exascale Infrastructure 6
Subscription Types 7
Service Limits for Exadata Database Service on Exascale Infrastructure 7
Metering Frequency and Per-Second Billing 8
Support Requirements 8
Exadata Cloud Management Interfaces 9
- Introduction to Exadata Cloud Management Interfaces 9
- OCI Control Plane Interfaces for Oracle Exadata Database Service on Exascale Infrastructure 10
- Local VM Command-Line Interfaces 11

2 What’s New in Oracle Exadata Database Service on Exascale Infrastructure

New Regions and Realms for Exadata Database Service on Exascale Infrastructure 3
Change in Backup Destination for New Tenancies in Select OCI Regions (Effective November 04, 2026) 6
Container Database (CDB) Snapshot and Clone 7
Flexible Software Release Selection for Exadata Database 7
Managed MCP Server Support for OCI Database Services 9
Key Pools for Oracle AI Database@AWS 9
Automatic Data Safe Registration for New Databases 10
AWS Key Management Service Integration for Exadata Database Service on Exascale Infrastructure@AWS 10
Data Guard Automatic Failover for Exadata Database Service on Exascale Infrastructure 11
Data Guard Support to Convert Physical Standby to Snapshot Standby and Vice Versa 12
Google Cloud Key Management Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud 13

Copyright © 2023, 2026, Oracle and/or its affiliates.

Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure 13
Change in Backup Destination for New Tenancies in Select OCI Regions (Effective June 23, 2026) 13
Enhanced Controls for Managing Unresponsive Virtual Machines 14
Oracle Managed Database Software Updates 15
Delegate Access Control for Exadata Database Service on Exascale Infrastructure 16
Change in Backup Destination for New Tenancies in Select OCI Regions (Effective May 15, 2026) 16
Cross-Service Data Guard Operations for Oracle Database 19c 17
Protect Sensitive Data With Oracle Cloud Infrastructure Zero Trust Packet Routing 17
Managing Certificates using DBCS Certificate Management 18
Manage Serial Console Access in Oracle Exadata Database Service on Exascale Infrastructure 18
Enhanced Out-of-Place Restore 19
Enhanced Data Guard Health Insights 19
Virtual Machine (VM) Snapshots 19
Auto scaling Exascale Storage Vault 19
Oracle Database 19c Support 20
Pluggable Database (PDB) Snapshots 20
Cross-Service Data Guard Operations 21
Change in Backup Destination for New Tenancies in Select OCI Regions (Effective August 06, 2025) 21
Expanded Tagging Support Across Database Workflows 22
Data Guard Setup with Pre-Check Validation 22
Auto Setup (Inline) of Network During Provisioning of VM Clusters 22
Long-Term Retention Backup (LTR) 22
Performance Hub & metrics available for databases running in Oracle Exadata Database Service on Exascale Infrastructure 23
Multiple Standby Databases 23
Deploy Single Node VM Clusters 24
Scale ECPUs to Zero 24

3 Preparing for Oracle Exadata Database Service on Exascale Infrastructure

Oracle Cloud Infrastructure (OCI) Requirements for Oracle Exadata Database Service on Exascale Infrastructure 1
Required IAM Policy for Oracle Exadata Database Service on Exascale Infrastructure 2
Network Setup for Oracle Exadata Database Service on Exascale Infrastructure Instances 3
VCN and Subnets 3
Option 1: Public Client Subnet with Internet Gateway 4
Option 2: Private Subnets 6
Requirements for IP Address Space 7

Copyright © 2023, 2026, Oracle and/or its affiliates.

Configuring a Static Route for Accessing the Object Store 8
Setting Up DNS for an Oracle Exadata Database Service on Exascale Infrastructure Instance 8
DNS: Short Names for the VCN, Subnets, and Oracle Exadata Database Service on Exascale Infrastructure instance 8
Configure Private DNS 9
Node Access to Object Storage: Static Route 10
Object Storage IP allocations 10
To configure a static route for Object Storage access 11
Service Gateway for the VCN 12
Option 1: Service Gateway Access to OCI Services 12
Option 2: Service Gateway Access to Both Object Storage and YUM Repos 13
Security Rules for the Oracle Exadata Database Service on Exascale Infrastructure 14
Rules Required for Both the Client Network and Backup Network 18
Rules Required Specifically for the Client Network 20
Rule Required Specifically for the Backup Network 21
Rules Required for Events Service 22
Rules Required for Monitoring Service 22
Ways to Implement the Security Rules 22
If you use network security groups 23
If you use security lists 23
Network Requirements for Oracle Database Autonomous Recovery Service 24
Create a Service Gateway to Object Storage 24
Cross-Service Interoperability Between ExaDB-D and ExaDB-XS: Data Guard, Backup, and Restore 25

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

Regional Availability for Exadata Database Service on Exascale Infrastructure 1
Tagging Oracle Exadata Database Service on Exascale Infrastructure Resources 3
Restarting a VM for Planned Maintenance 7
Connecting to an Oracle Exadata Database Service on Exascale Infrastructure VM 8
Prerequisites for Accessing Oracle Exadata Database Service on Exascale Infrastructure 8
SCAN Listener Port Setting 9
Connecting to a Virtual Machine with SSH 9
Connecting from a Unix-Style System 10
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY 10
Accessing a Database After You Connect to the Virtual Machine 11
Using Oracle Net Services to Connect to a Database 13
Prerequisites for Connecting to a Database with Oracle Net Services 13
Connecting to a Database with SQL Developer 14

Copyright © 2023, 2026, Oracle and/or its affiliates.

Connecting to a Database Using SCAN 15
Connecting to a Database Using a Node Listener 17
Connect to the Oracle Exadata Database Service on Exascale Infrastructure Service 18
Connecting to a Database with SQL Developer 18
Connecting to a Database with Oracle Net Services 19
Capacity Limits for Exadata Database Service on Exascale Infrastructure 23
Best Practices for Oracle Exadata Database Service on Exascale Infrastructure VMs 24

5 How-to Guides

Managed MCP Servers for OCI Database Services 2
Create MCP Server 6
View the List of MCP Servers 6
View the Details of an MCP Server 7
Enable or Disable an MCP Server 7
Move an MCP Server to Another Compartment 8
Delete an MCP Server 8
Edit the Details of an MCP Server 9
Manage Database Security with Oracle Data Safe 9
About Oracle Data Safe 9
Get Started 11
Using Oracle Data Safe 11
Automatic Registration of Databases with Oracle Data Safe 13
Connecting to an Oracle Exadata Database Service on Exascale Infrastructure VM 14
Connection Prerequisites 15
About Connecting to a VM with SSH 15
Connecting from a Unix-Style System 15
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY 16
To access a database after you connect to the VM 17
Connect to the Oracle Exadata Database Service on Exascale Infrastructure Service 17
Connecting to a Database with SQL Developer 17
Connecting to a Database with Oracle Net Services 18
Manage Oracle Exadata Database Service on Exascale Infrastructure 22
Using the Console to Provision Oracle Exadata Database Service on Exascale Infrastructure 22
Lifecycle Management Operations 23
Network Management Operations 24
Management Tasks for the Oracle Cloud Infrastructure Platform 25
Oracle Database License Management Tasks 27
Using the API to Create Infrastructure Components 28
Using the API to Manage Oracle Exadata Database Service on Exascale Infrastructure Instance 29
Manage VM Clusters 29

Copyright © 2023, 2026, Oracle and/or its affiliates.

Using the Console to Manage VM Clusters on Oracle Exadata Database Service on Exascale Infrastructure 30
To create a cloud VM cluster 31
To add security attribute to an Exadata VM Cluster 38
To edit a security attribute 39
To remove a Security Attribute 39
To scale database storage capacity of a cluster database in the VM cluster 40
To add database server or storage server capacity to a cloud VM cluster 41
To Enable, Partially Enable, or Disable Diagnostics Collection 41
Using the Console to Update the License Type on a VM Cluster 43
To scale VM Clusters 43
To add SSH keys to a VM cluster 45
To Add or Delete SSH Keys After VM Cluster Creation 45
Using the Console to Stop, Start, or Reboot a Virtual Machine 46
To Check the Status of a VM Cluster Virtual Machine 48
To Create a Snapshot of Virtual Machines 48
To Delete a Snapshot of Virtual Machines 49
To Mount a Snapshot of Virtual Machines 49
To Unmount a Snapshot of Virtual Machines 49
To Create VM File System Snapshots for VM Cluster 50
To Move a VM Cluster to Another Compartment 50
Using the Console to Terminate a VM Cluster 51
To view details about private DNS configuration 51
Adding or Removing a VM From a VM Cluster 52
Add a VM to a VM Cluster 52
Terminate or Remove a VM from a VM Cluster 52
Overview of Automatic Diagnostic Collection 53
Incident Logs and Trace Files 54
Health Metrics 57
Using the API to Manage Oracle Exadata Database Service on Exascale Infrastructure Instance 61
Troubleshooting Virtual Machines Using Console Connections 62
Required IAM Policies 63
Prerequisites 63
Create the Virtual Machine Serial Console Connection 66
Make an SSH Connection to the Serial Console 67
Using Cloud Shell to Connect to the Serial Console 70
Displaying the Console History for a Virtual Machine 71
Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems 74
Exiting the Virtual Machine Serial Console Connection 76
Manage Exascale Database Vaults on Oracle Exadata Database Service on Exascale Infrastructure 77

Copyright © 2023, 2026, Oracle and/or its affiliates.

Manage Software Images 78
Using Software Images in Oracle Cloud Infrastructure 78
Creation and Storage of Software Images 79
Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Home 79
Using a Software Image with an Oracle Exadata Database Service on Exascale Infrastructure Instance 80
Using the Console for Software Images 80
To Create a Database Software Image 80
To create a database software image from a Database Home 81
To update database software using custom database software image 81
To Delete a Database Software Image 82
Using the API to manage database software images 82
Flexible Software Release Selection for Exadata Database 82
Acknowledging Older Release Updates and Maintenance Releases in API Requests 85
Create Oracle Database Homes on an Oracle Exadata Database Service on Exascale Infrastructure System 86
About Creating Oracle Database Homes on an Oracle Exadata Database Service on Exascale Infrastructure System 87
To create a new Database Home in an existing Oracle Exadata Database Service on Exascale Infrastructure instance 87
To Create a Database Software Image from a Database Home 88
Using the API to Create Oracle Database Home on Oracle Exadata Database Service on Exascale Infrastructure 89
Managing Oracle Database Homes on an Oracle Exadata Database Service on Exascale Infrastructure Instance 89
Manage Database Home Using the Console 90
To view information about a Database Home 90
To delete a database home 90
To manage tags for your Database Home 91
To Move a Database to Another Database Home 91
Using the API to Manage Oracle Database Home on Oracle Exadata Database Service on Exascale Infrastructure 92
Manage Databases on Oracle Exadata Database Service on Exascale Infrastructure 93
Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Exascale Infrastructure 93
Oracle Database Releases Supported by Oracle Exadata Database Service on Exascale Infrastructure 94
Provisioning and Managing Exadata Databases 94
About Oracle Managed Database Software Updates 95
Database Memory Initialization Parameters 96
Customer-Managed Keys in Oracle Exadata Database Service on Exascale Infrastructure 96
Using the Console to Manage Databases on Oracle Exadata Database Service on Exascale Infrastructure 98

Copyright © 2023, 2026, Oracle and/or its affiliates.

Using the API to manage Databases 135
Create and Manage Exadata Pluggable Databases 135
Limitations for Pluggable Database Management 137
Creating an Exadata Pluggable Database 137
Managing an Exadata Pluggable Database 142
Cloning an Exadata Pluggable Database 149
To download the PDB Connection Bundle 156
Restoring an Exadata Pluggable Database 157
Using the Console to Perform an In-Place Restore of a Pluggable Database (PDB) 157
Using the Console to Perform an Out-of-Place Restore of a Pluggable Database (PDB) 157
Changing the Database Passwords 158
Manage Database Backup and Recovery on Oracle Exadata Database Service on Exascale Infrastructure 159
Oracle Recommended Options to Perform Backup and Recovery Operations 160
Managing Exadata Database Backups 161
Managed Backup Types and Usage Information 162
Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console 164
Long-Term Retention Backup with Recovery Service 165
Default Backup Channel Allocation 166
Prerequisites for Backups on Oracle Exadata Database Service on Exascale Infrastructure 167
Using the Console to Manage Backups 168
To configure automatic backups for a database 168
To create an on-demand backup of a database 171
To view backup status 171
To cancel a backup 172
To delete full backups from Object Storage 172
To delete standalone backups from Object Storage 173
To change the retention period of an LTR backup with Recovery Service 173
To designate Autonomous Recovery Service as a Backup Destination for an Existing Database 174
Recovering an Exadata Database from Backup Destination 175
Using the Console to restore a database 175
Managing Exadata Database Backups by Using dbaascli 176
Default Backup Configuration 177
To create a backup configuration file 177
To create an on-demand backup 184
To remove the backup configuration 185
To delete a local backup 185
To delete a backup in Object Storage 186
Using the API to Manage Backup and Recovery 186

Copyright © 2023, 2026, Oracle and/or its affiliates.

Using the API to manage backups 187
Alternative Backup Methods 187
Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management 188
Recovering a Database Using Oracle Recovery Manager (RMAN) 190
Patch and Update an Oracle Exadata Database Service on Exascale Infrastructure System 190
User-Managed Maintenance Updates 191
Patching and Updating an Oracle Exadata Database Service on Exascale Infrastructure System 191
Patching and Updating VM Cluster's GI and Database Homes 191
Updating an Exadata Cloud VM Cluster Operating System 208
Applying Out-of-Cycle Kernel and Individual Package Updates 210
Interim Software Updates 210
Create an Interim Software Update 211
Download an Interim Software Update 211
Delete an Interim Software Update 212
Move an Interim Software Update Resource to Another Compartment 212
Using the API to Manage Interim Software Updates 212
Use Oracle Data Guard with Oracle Exadata Database Service on Exascale Infrastructure 213
About Using Oracle Data Guard with Oracle Exadata Database Service on Exascale Infrastructure 213
Prerequisites for Using Oracle Data Guard with Oracle Exadata Database Service on Exascale Infrastructure 214
Network Requirements for Data Guard 215
Password Requirements 217
Known Issues for Exadata Cloud Infrastructure and Data Guard 217
Adding a Node to a VM Cluster 217
Removing a Node from a VM Cluster 217
Working with Oracle Data Guard 217
Switchover 218
Failover 218
Reinstate 218
Enhanced Data Guard Health Status Reporting 219
Using the Console to Manage Oracle Data Guard Associations 220
To enable Data Guard on Oracle Exadata Database Service on Exascale Infrastructure 221
To view Data Guard associations of databases in a Cloud VM Cluster 230
To enable automatic backups on a standby database 230
To perform a database switchover 231
To edit the Oracle Data Guard association 232
To Disable Automatic Failover 232
To perform a database failover 233
To reinstate a database 233

Copyright © 2023, 2026, Oracle and/or its affiliates.

To terminate a Data Guard association on an Oracle Exadata Database Service on Exascale Infrastructure instance 234
Handling Standby Database Termination 234
Using the API to manage Data Guard associations 234
Configure Oracle Database Features for Oracle Exadata Database Service on Exascale Infrastructure 235
Using Oracle Multitenant on an Oracle Exadata Database Service on Exascale Infrastructure Instance 235
To determine if you need to create and activate an encryption key for the PDB 236
To create and activate the master encryption key in a PDB 237
To export and import a master encryption key 238
Managing Tablespace Encryption 239
Migrate to Oracle Exadata Database Service on Exascale Infrastructure 241
Connect Identity and Access Management (IAM) Users to Oracle Exadata Database Service on Exascale Infrastructure 241
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database 241
About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database 242
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication 243
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication 243
Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database 245
Prerequisites for IAM Authentication on Oracle Database 246
Disable External Authentication Scheme 246
Configure TLS to Use IAM Tokens 247
Enable, Disable, and Re-enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database 248
Enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database 248
Disable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database 249
Using Oracle Database Tools with Identity and Access Management (IAM) Authentication 249
Manage Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies, Users, Roles, and Database Passwords 250
Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies for IAM Users 250
Authorize Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users on Oracle Database 251
To Exclusively Map a Local IAM User to an Oracle Database Global User 252
Add Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Roles on Oracle Database 253

Copyright © 2023, 2026, Oracle and/or its affiliates.

Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password for IAM Users 253
Configuring Client Connection 254
Configure a Client Connection for SQL*Plus that Uses an IAM Database Password Verifier 254
Configure Client Connection for SQL*Plus that Uses an IAM Token 255
Client Connections That Use a Token Requested by an IAM User Name and Database Password 256
Use Instance Principal to Access Database with IAM Authentication 257
Configure Proxy Authentication 257
Use Database Link with IAM Authenticated Users 259
Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure 259
Prerequisites 260
Network Requirements for Creating an Identity Connector and KMS Resources 260
Using the Console to Manage Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure 263
Create an Identity Connector from the OCI Console 263
View the details of an Identity Connector 265
Enable or Disable the Azure Key Management 265
Create Azure Key Vault (Managed HSM, Premium, and Standard) and Assign Required Permissions 266
Register Azure Key Vaults in the OCI console 267
Create a Database and Use Azure Key Vault as the Key Management Solution 268
Change the Key Management from Oracle Wallet to Azure Key Vault 269
Rotate the Keys Managed by Azure Key Vault for a Container Database 270
Rotate the Keys Managed by Azure Key Vault for a Pluggable Database 270
Using the API to Manage Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure 271
Updating the Multicloud PKCS#11 Driver 272
Release Notes 274
Release 2.202 (260721.2348) 274
Release 2.196 (260325.2123) 274
Release 2.194 (260227.1414) 274
Release 0.2 (250910.0501) 274
Google Cloud Key Management Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud 275
Prerequisites 275
Using the Console to Manage GCP KMS Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud 276
Verify the Default Identity Connector Attached to the VM Cluster 277
Create a Key Ring in Google Cloud Console 277
Create a Key in Google Cloud Console 278

Copyright © 2023, 2026, Oracle and/or its affiliates.

Grant Permissions in Google Cloud KMS for Key Discovery by Oracle Cloud Infrastructure (OCI) 279
Register GCP Key Ring in Oracle Cloud Infrastructure (OCI) 280
Enable or Disable Google Cloud Key Management 281
Create a Database and Use GCP Customer-Managed Encryption Key (CMEK) as the Key Management Solution 281
Change the Key Management from Oracle Wallet to GCP Customer Managed Encryption Key (CMEK) 282
Rotate the GCP Customer Managed Encryption Key of a Container Database (CDB) 282
Rotate the GCP Customer Managed Encryption Key of a Pluggable Database (PDB) 283
Using the API to Manage GCP KMS Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud 283
Updating the Multicloud PKCS#11 Driver 284
Release Notes 286
Release 2.202 (260721.2348) 287
Release 2.196 (260325.2123) 287
Release 2.185 (251201.0551) 287
AWS Key Management Service Integration for Exadata Database Service on Exascale Infrastructure@AWS 287
Prerequisites 288
Configure OCI Identity Domain 288
Create or Use an Existing ODB Network 288
Create an Exascale VM Cluster 289
Configure Identity Provider 289
Associate an IAM Role to an Exascale VM Cluster 290
Create a Key 290
Register AWS KMS Key 292
OCI IAM Policy Requirements 293
Understanding Key Pools 293
Create a Key Pool 295
Discover and Register Existing AWS Keys in a Key Pool 295
Replicate a Key Pool Across Regions 296
Add or Remove a Key from a Key Pool 296
View Key Usage Count 297
Migrate from an Individual AWS KMS Key to a Key Pool 297
Create a Database with an Individual AWS KMS Key Using dbaascli 298
Using the Console to Manage AWS KMS Integration for Exadata Database Service on Exascale Infrastructure@AWS 299
Enable or Disable AWS Key Management 299
Create a Database and Use AWS KMS as the Key Management Solution 299
Steps to Change the Key Management of Existing Databases from Oracle Wallet to AWS KMS 300
Rotate the AWS KMS Key of a Container Database (CDB) 300
Rotate the AWS KMS Key of a Pluggable Database (PDB) 301

Copyright © 2023, 2026, Oracle and/or its affiliates.

(Optional) Disable a CDB or PDB Key 301
Using the API to Manage AWS KMS Integration for Exadata Database Service on Exascale Infrastructure@AWS 302
OracleDbAwsIdentityConnector 302
OracleDbAwsKey 303
Updating the Multicloud PKCS#11 Driver 303
Release Notes 306
Release 2.203 (260927.1348) 306
Release 2.202 (260721.2348) 306
Release 2.196 (260325.2123) 306
Release 2.185 (251201.0551) 306

6 Reference Guides for Oracle Exadata Database Service on Exascale Infrastructure

EU Data Act Compliance 2
Using the dbaascli Utility on Oracle Exadata Database Service on Exascale Infrastructure 2
About Using the dbaascli Utility on Oracle Exadata Database Service on Exascale Infrastructure 3
Creating Databases Using dbaascli 3
Listing Available Software Images and Versions for Database and Grid Infrastructure 3
Creating Oracle Database Home 5
Creating Oracle Database In the Specified Oracle Database Home 6
Changing the Database Passwords 8
Managing Oracle Exadata Database Service on Exascale Infrastructure Software Images Using the Dbaascli Utility 9
Listing Available Software Images and Versions for Database and Grid Infrastructure 10
To download a software image 11
Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check Using dbaascli 12
Collecting Tooling Log Data Examples 13
Performing a Health Check Examples 16
Updating Cloud Tooling Using dbaascli 17
Creating a Duplicate Database 18
Using dbaascli to Duplicate a Cloud Database 18
Considerations When Using OCI Vault for the Key Management 19
dbaascli Command Reference 20
dbaascli admin updateStack 25
dbaascli cswlib deleteLocal 27
dbaascli cswlib download 28
dbaascli cswlib listLocal 29
dbaascli cswlib showImages 31
dbaascli database addInstance 32
dbaascli database backup 34

Copyright © 2023, 2026, Oracle and/or its affiliates.

dbaascli database bounce 38
dbaascli database changepassword 40
dbaascli database convertToPDB 42
dbaascli database create 42
dbaascli database delete 46
dbaascli database deleteInstance 48
dbaascli database duplicate 49
dbaascli database getDetails 52
dbaascli database getPDBs 53
dbaascli database modifyParameters 54
dbaascli database recover 56
dbaascli database createTemplate 58
dbaascli database start 58
dbaascli database status 60
dbaascli database stop 60
dbaascli database upgrade 62
dbaascli dataguard prepareStandbyBlob 64
dbaascli dataguard updateDGConfigAttributes 65
dbaascli dbhome create 66
dbaascli dbHome delete 68
dbaascli dbhome getDatabases 70
dbaascli dbHome getDetails 72
dbaascli dbimage purge 74
dbaascli diag collect 74
dbaascli diag healthCheck 75
dbaascli grid configureTCPS 78
dbaascli grid removeTCPSCert 81
dbaascli grid rotateTCPSCert 83
dbaascli grid upgrade 85
dbaascli job getStatus 87
dbaascli patch db prereq 88
dbaascli pdb backup 88
dbaascli pdb bounce 91
dbaascli pdb close 92
dbaascli pdb getConnectString 93
dbaascli pdb create 93
dbaascli pdb delete 95
dbaascli pdb getDetails 96
dbaascli pdb list 96
dbaascli pdb localClone 97
dbaascli pdb open 99
dbaascli pdb recover 100

Copyright © 2023, 2026, Oracle and/or its affiliates.

dbaascli pdb refresh 101
dbaascli pdb relocate 101
dbaascli pdb remoteClone 103
dbaascli system getDBHomes 105
dbaascli tde changePassword 106
dbaascli tde addSecondaryHsmKey 108
dbaascli tde enableWalletRoot 110
dbaascli tde encryptTablespacesInPDB 113
dbaascli tde fileToHsm 115
dbaascli tde getHsmKeys 117
dbaascli tde getMkidForKeyVersionOCID 119
dbaascli tde getPrimaryHsmKey 121
dbaascli tde hsmToFile 123
dbaascli tde listKeys 125
dbaascli tde removeSecondaryHsmKey 127
dbaascli tde rotateMasterKey 130
dbaascli tde setKeyVersion 133
dbaascli tde setPrimaryHsmKey 136
dbaascli tde status 139
Database Service Events 141
Overview of Database Service Events 141
Monitor Metrics for VM Cluster Resources 144
View Metrics for VM Cluster 144
View Metrics for a Database 145
View Metrics for VM Clusters in a Compartment 147
View Metrics for Databases in a Compartment 148
Manage Oracle Trace File Analyzer 149
Manage Database Service Agent 150
Metrics for Oracle Exadata Database Service on Exascale Infrastructure in the Monitoring Service 150
Oracle Exadata Database Service on Exascale Infrastructure Events 155
About Event Types on Oracle Exadata Database Service on Exascale Infrastructure 157
Prerequisites for Event Service 157
Oracle Exadata Database Service on Exascale Infrastructure Event Types 158
Oracle Exadata Database Service on Exascale Infrastructure Maintenance Event Types 160
Exadata Cloud Infrastructure Critical and Information Event Types 167
Exadata Cloud Infrastructure VM Cluster Event Types 170
VM Node Subsetting Event Types 182
Data Guard Event Types 187
Oracle Database Home Event Types 190
Database Event Types 192
Pluggable Database Event Types 194

Copyright © 2023, 2026, Oracle and/or its affiliates.

Database Service Events 201
Overview of Database Service Events 201
Receive Notifications about Database Service Events 205
Database Service Event Types 205
Temporarily Restrict Automatic Diagnostic Collections for Specific Events 213
Application VIP Event Types 216
Interim Software Updates Event Types 228
Serial Console Connection Event Types 234
Viewing Audit Log Events 239
Monitor Metrics to Diagnose and Troubleshoot Problems with Pluggable Databases 240
About Database Management 241
Using the Console to Enable Database Management for a Container Database (CDB) 241
Enable Database Management 242
Using the Console to Enable Database Management for a Pluggable Database (PDB) 244
Enable Database Management 245
Using the Console to Edit Database Management for a Pluggable Database (PDB) 247
Edit Database Management 248
Using the Console to Disable Database Management for a Pluggable Database (PDB) 250
Using the Console to View Performance Hub for a Container Database (CDB) 251
Using the Console to View Performance Hub for a Pluggable Database (PDB) 251
Using the API to Enable, Disable, or Update Database Management Service 252
Oracle Cloud Database Metrics 252
Using the Console View Metrics for a Container Database (CDB) 253
Using the Console to View Metrics for a Pluggable Database (PDB) 253
Policy Details for Oracle Exadata Database Service on Exascale Infrastructure 254
About Resource-Types 255
Resource-Types for Exadata Cloud Service Instances 255
Supported Variables 256
Details for Verb + Resource-Type Combinations 256
Database-Family Resource Types 257
Soft Deprecation of Certain Database Service APIs and Fields 257
Permissions and API operation details for DB Backups 258
Permissions and API operation details for Databases (CDBs) 259
Permissions and API operation details for CDB Snapshot 260
Permissions and API operation details for Data Guard Association 260
Permissions and API operation details for DB Nodes 261
Permissions and API operation details for DB Node Snapshots 261
Permissions and API operation details for DB Node Console Connection 262
Permissions and API operation details for DB Homes 262
Permissions and API operation details for Database Software Image 263
Permissions and API operation details for VM Clusters 264
Permissions and API operation details for Exascale DB Storage Vaults 265

Copyright © 2023, 2026, Oracle and/or its affiliates.

Permissions and API operation details for Data Guard Group 266
Permissions and API operation details for Key Stores 266
Permissions and API operation details for DB Connection Bundles 267
Permissions Required for Each API Operation 267
Permissions and API operation details for Pluggable Databases (PDBs) 271
Permissions and API operation details for Oracle DB Azure Connectors 273
Permissions and API operation details for Oracle DB Azure Blob Mounts 274
Permissions and API operation details for Oracle DB Azure Blob Containers 274
Permissions and API operation details for Oracle DB MCI Work Requests 275
Permissions and API operation details for Multicloud Resource Discoveries 275
Permissions and API operation details for Oracle DB Azure Vaults 276
Permissions and API operation details for Oracle DB Azure Keys 276
Permissions and API operation details for Oracle DB Azure Vault Associations 277
Permissions and API operation details for Oracle DB GCP Connectors 277
Permissions and API operation details for Oracle DB GCP Keyrings 278
Permissions and API operation details for Oracle DB GCP Keys 278
Permissions and API operation details for Oracle DB AWS Identity Connectors 278
Permissions and API operation details for Oracle DB AWS Key 279
Permissions and API operation details for Oracle DB AWS Key Pool 279
Oracle Cloud Infrastructure Operations Insights 280
Managing Exadata Resources with Oracle Enterprise Manager Cloud Control 280
Overview of Oracle Enterprise Manager Cloud Control 281
Features of Enterprise Manager Cloud Control 281
Analyzing Exadata Database Service Database Performance 282
Troubleshooting Oracle Exadata Database Service on Exascale Infrastructure Systems 282
Known Issues for Oracle Exadata Database Service on Exascale Infrastructure 283
Troubleshooting Oracle Managed Database Software Updates Issues 283
Troubleshooting Oracle Data Guard 286
Troubleshooting Data Guard using logfiles 287
Troubleshooting the Data Guard Setup Process 289
Obtaining Further Assistance 291
Collecting Cloud Tooling Logs 292
Collecting Oracle Diagnostics 292

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 Oracle 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.
- Support Requirements
- 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

Oracle 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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. Oracle 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.

Oracle Exadata Database Service on Exascale Infrastructure supports both Oracle Database 19c and Oracle AI Database 26ai. However, the Exascale implementation is different: Oracle AI Database 26ai leverages the Exascale smart storage architecture, while Oracle Database 19c uses the Exascale block storage architecture.

Oracle AI Database 26ai is tightly integrated with the Exascale Infrastructure. This integration enables direct access to Exascale storage through the Exascale Storage Vault, which eliminates the need for Oracle Automatic Storage Management (ASM). Storage for database files resides in an Oracle Exadata Exascale Storage Vault. The Storage Vault provides high-performance and scalable Exadata Smart Storage. You can scale Exadata Smart Storage online as needed, with a single command. Unlike Exadata Database Service on Dedicated Infrastructure, Exadata Database Service on Exascale Infrastructure does not require you to manage adding storage servers to the system, or to manage storage allocations.

The following schematic provides a high-level overview of the architecture of an Oracle AI Database 26ai VM Cluster and associated resources:

Figure 1-1 ExaDB-XS Oracle AI Database 26ai Architecture using Exascale smart storage

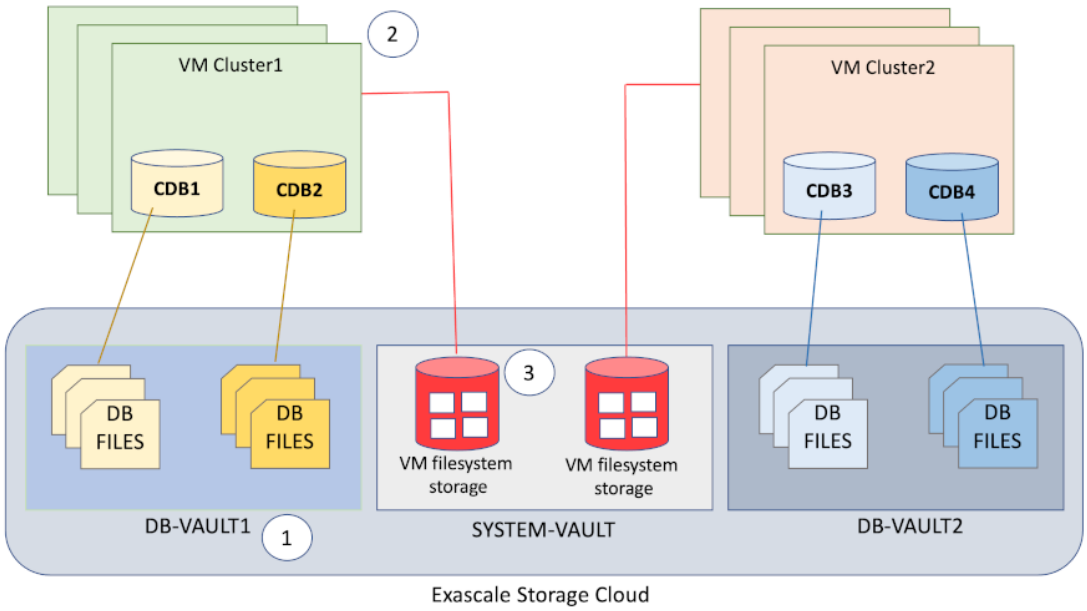

In the diagram, the architecture indicated by the numbered callouts consists of the following elements:

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Oracle Database 19c leverages Exascale block storage. This storage is comprised of Exascale Direct Volumes (EDV), and relies on ASM to manage the database storage. Each database stores its files in dedicated DATA, RECO, and LOG disk groups. During container database creation, you must specify storage allocation (in GB) for the DATA (data storage) and RECO (recovery area storage) disk groups. The LOG (redo) disk group storage allocation is automatically determined based on the number of VMs in the cluster. As with Oracle AI Database 26ai, overall storage is managed using the Exascale storage vault. You can scale this storage online, with additional storage becoming immediately available. However, in the Oracle Database 19c architecture, you must also scale storage at the individual database level by providing updated allocation values for the DATA and RECO disk groups.

The following schematic provides a high-level overview of the architecture of an Oracle Database 19c VM Cluster and associated resources:

Figure 1-2 ExaDB-XS Oracle Database 19c Architecture Using Exascale block storage

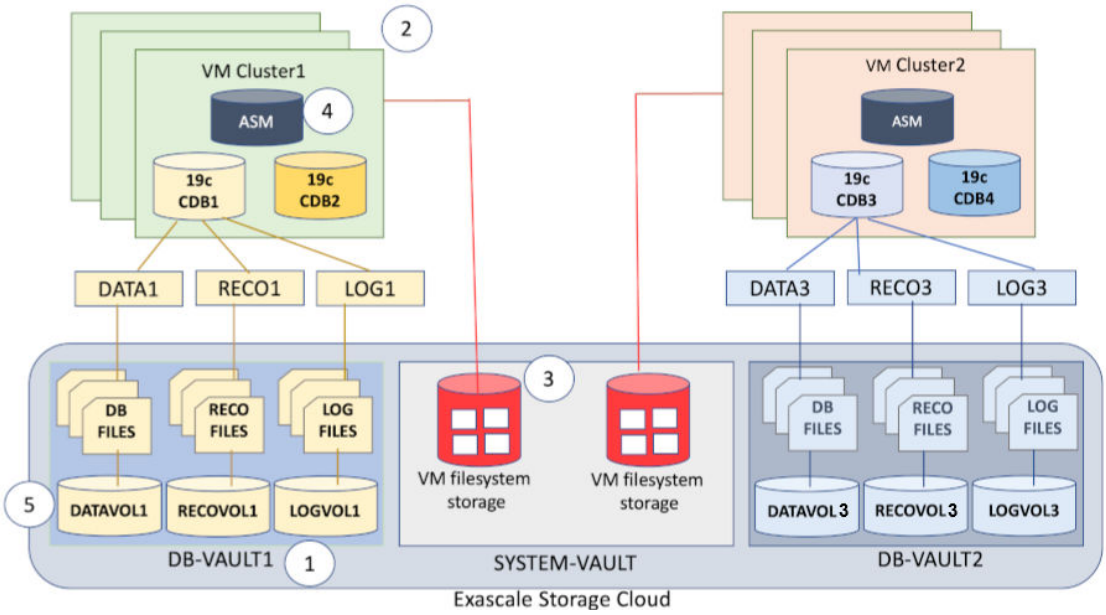

In the diagram, the architecture indicated by the numbered callouts consists of the following elements:

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. An ASM instance to manage database storage of container databases in the VM Cluster using DATA, RECO, and LOG disk groups
5. Exascale Direct Volumes (EDV) for the DATA, RECO, and LOG disk groups for each container database in the VM Cluster

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> The following restrictions apply:
>
> - You cannot provision within the same VM cluster both Oracle AI Database 26ai using Exascale smart storage and Oracle Database 19c using Exascale block storage.
> - An ExaDB-XS VM cluster configured with Exascale smart storage cannot share a storage Vault with another ExaDB-XS VM cluster that uses the Exascale block storage.

The following schematic illustrates the architectural difference between Oracle Database 19c and Oracle AI Database 26ai.

Figure 1-3 ExaDB-XS Architecture Comparison Between Oracle AI Database 26ai and Oracle Database 19c

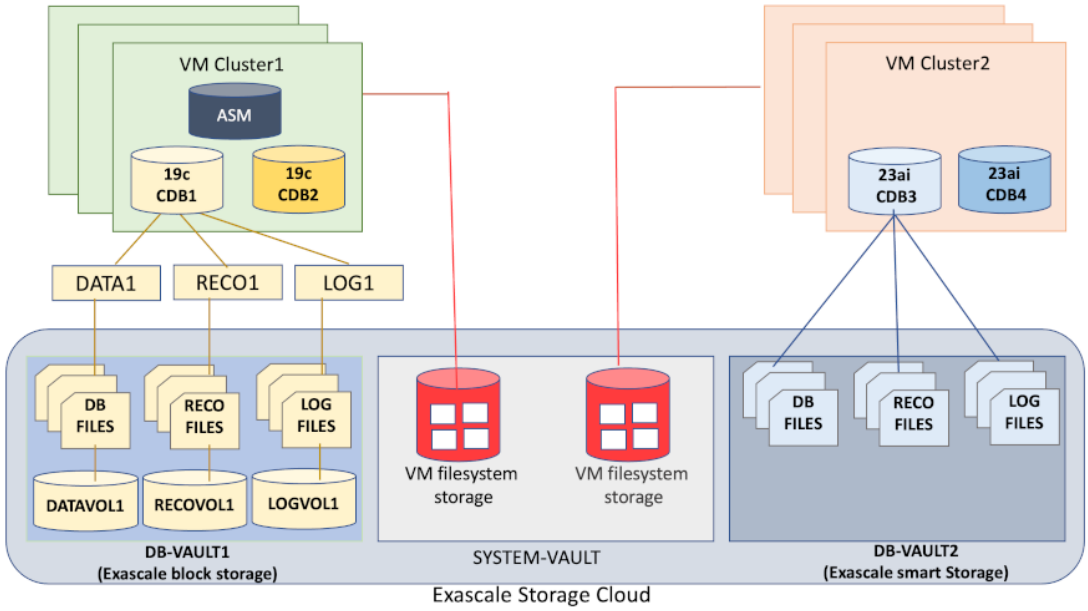

The basic unit of consumption in ExaDB-XS is a VM cluster. To facilitate VM portability, Exascale hosts storage for VM file systems on shared storage that is fully managed by Oracle. Oracle can migrate VMs across a pool of physical servers, because the VM filesystems that host the database binaries do 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.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 1-4 Core Reservation and Scaling

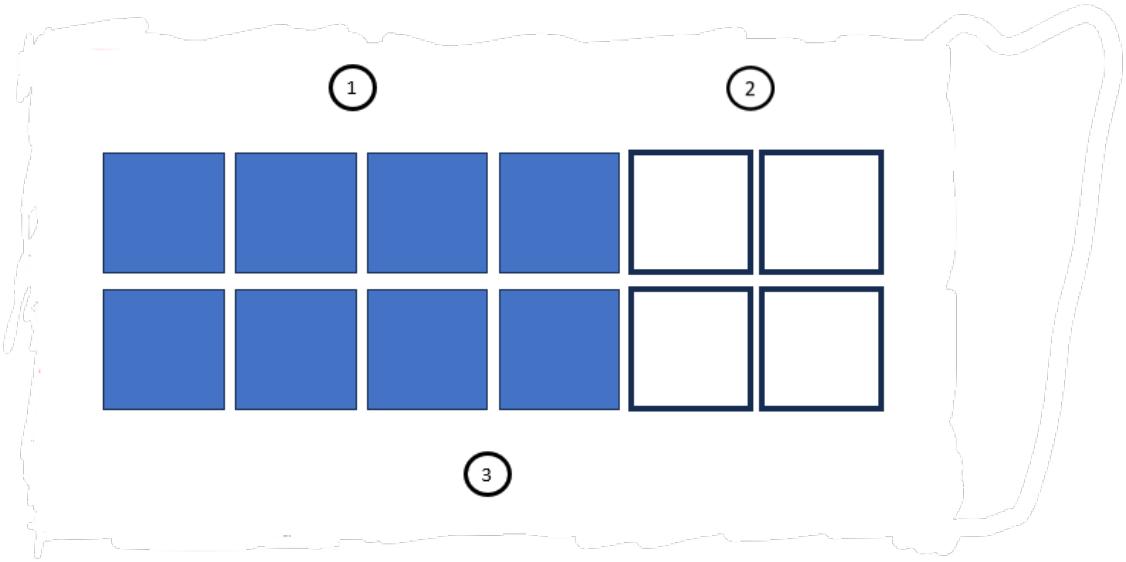

The illustration shows the following active and reserved cores in a VM:

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Related Topics

- Compute Models in Autonomous Database

Accessing the Oracle 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 AI Database**, and then to **Oracle 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

Edition options, such as Oracle Database In-Memory and Oracle Real Application Clusters (Oracle RAC).

Oracle Exadata Database Service on Exascale Infrastructure supports Oracle AI Database 26ai and Oracle Database 19c.

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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> - The values provided for ExaDB-XS in the table represent the current default tenancy limits and do not necessarily reflect the overall available capacity.
> - Using the OCI console, you can request increase in the limits. For more information, see To Request Service Limit Increase.
> - For design specification limits of the ExaDB-XS service, see Capacity Limits for Exadata Database Service on Exascale Infrastructure .

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.

Support Requirements

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

Oracle Applications Unlimited (specifically Oracle E-Business Suite, PeopleSoft, JD Edwards, Siebel and Hyperion applications) require ongoing updates to ensure continued interoperability. As such, Oracle requires customers to have an active Oracle support contract for any such programs running on Oracle Exadata Database Service. Note that third party support does not

Copyright © 2023, 2026, Oracle and/or its affiliates.

satisfy this requirement. This requirement shall not be interpreted as superseding anything in the Oracle Technical Support Policies, including, but not limited to, the Matching Service Level requirement.

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)

Copyright © 2023, 2026, Oracle and/or its affiliates.

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*

Copyright © 2023, 2026, Oracle and/or its affiliates.

- *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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 Exadata Database Service on Exascale Infrastructure
  Oracle continues to add new regions for Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).
- Change in Backup Destination for New Tenancies in Select OCI Regions (Effective November 04, 2026)
- Container Database (CDB) Snapshot and Clone
  This feature allows you to create a snapshot or clone of a container database (CDB).
- Flexible Software Release Selection for Exadata Database
  Lifecycle operations now default to the latest available Database and Grid Infrastructure release updates, and Exadata System Software maintenance releases. Where release selection is supported, you can select certain older release updates or maintenance releases by acknowledging the message displayed.
- Managed MCP Server Support for OCI Database Services
  OCI Database Services now supports managed Model Context Protocol (MCP) servers for connecting MCP-compatible AI agents and applications to supported database resources.
- Key Pools for Oracle AI Database@AWS
  You can now use key pools to manage customer-managed encryption keys for your CDBs and PDBs.
- Automatic Data Safe Registration for New Databases
  Newly provisioned container databases (CDBs) are now automatically registered with Oracle Data Safe by default, and newly created pluggable databases (PDBs) inherit the Data Safe registration policy from their parent CDB.
- AWS Key Management Service Integration for Exadata Database Service on Exascale Infrastructure@AWS
  Exadata Database Service on Exascale Infrastructure@AWS supports integration with AWS Key Management Service (KMS).
- Data Guard Automatic Failover for Exadata Database Service on Exascale Infrastructure
  Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS) now supports automatic failover powered by Oracle Data Guard automation.
- Data Guard Support to Convert Physical Standby to Snapshot Standby and Vice Versa
  With this enhancement, you can temporarily convert a physical standby database to an updatable (read-write) copy, then revert it to its original state.
- Google Cloud Key Management Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud
  Exadata Database Service on Exascale Infrastructure@Google Cloud now supports integration with Google Cloud Platform's Key Management Service (KMS). This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using GCP Customer Managed Encryption Keys (CMEKs).

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure
  Exadata Database Service on Exascale Infrastructure@Azure enables you to store your database's transparent data encryption (TDE) keys, also known as master encryption keys (MEKs) in either a file-based Oracle wallet or in the OCI Vault.
- Change in Backup Destination for New Tenancies in Select OCI Regions (Effective June 23, 2026)
- Enhanced Controls for Managing Unresponsive Virtual Machines
  This release introduces enhanced controls for managing unresponsive virtual machines (VMs), providing clearer behavior and safer handling of reboot and stop operations.
- Oracle Managed Database Software Updates
  Oracle Managed Database Software Updates gives you the option to let Oracle manage the end-to-end database software update lifecycle for enrolled databases.
- Delegate Access Control for Exadata Database Service on Exascale Infrastructure
  Delegate Access Control service is enabled for Oracle Exadata Database Service on Exascale Infrastructure.
- Change in Backup Destination for New Tenancies in Select OCI Regions (Effective May 15, 2026)
- Cross-Service Data Guard Operations for Oracle Database 19c
  You are now able to deploy Oracle Data Guard environments with cross-service primary and secondary databases for Oracle Database 19c.
- Protect Sensitive Data With Oracle Cloud Infrastructure Zero Trust Packet Routing
  Protect your data against unauthorized access and exfiltration with Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR), a data-aware cloud security control.
- Managing Certificates using DBCS Certificate Management
  Database Cloud Services (DBCS) Certificate Management lets you centrally manage X.509 certificates across all DBCS-managed Oracle Database services. This service is HA-aware (High Availability) and includes features that provide connection details to facilitate seamless connectivity to all Data Guard peer databases.
- Manage Serial Console Access in Oracle Exadata Database Service on Exascale Infrastructure
  You can create and delete serial console connections to your Oracle Exadata Database Service on Exascale Infrastructure systems to diagnose and resolve VM guest operating system issues using an SSH connection in case standard SSH access to the VMs is not possible.
- Enhanced Out-of-Place Restore
  ExaDB-XS currently supports database creation from a specific backup and from latest backup. This new functionality enables you to create database from a specified point-in-time (timestamp).
- Enhanced Data Guard Health Insights
  Enhanced Data Guard health status reporting now provides detailed metrics on overall health, switchover/failover readiness, data loss exposure, and disaster recovery, making it especially useful for managing multiple standby databases.
- Virtual Machine (VM) Snapshots
  This feature allows customers to create and manage snapshots of their virtual machines (VM) and store them locally on Exascale storage.
- Auto scaling Exascale Storage Vault
  The Exascale Storage Vaults can now scale automatically as usage grows, with storage expanding incrementally up to a user-defined limit.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Oracle Database 19c Support
  Oracle is excited to announce that Exadata Database Service on Exascale Infrastructure (ExaDB-XS) now supports Oracle Database 19c and Oracle Grid Infrastructure 19c.
- Pluggable Database (PDB) Snapshots
  With this enhancement, you can restore, refresh, and relocate a Pluggable Database (PDB).
- Cross-Service Data Guard Operations
  You are now able to deploy Oracle Data Guard environments with cross-service primary and secondary databases.
- Change in Backup Destination for New Tenancies in Select OCI Regions (Effective August 06, 2025)
  Starting August 06, 2025, the Autonomous Recovery Service is now the exclusive backup destination during automatic backup configuration for newly created tenancies in the following OCI regions: Frankfurt (FRA), Phoenix (PHX), and Tokyo (NRT).
- Expanded Tagging Support Across Database Workflows
  Tagging, previously supported during database creation, is now extended to additional workflows, including creating a standby database and creating a database from a backup.
- Data Guard Setup with Pre-Check Validation
  You are now able to perform a precheck before setting up Oracle Data Guard, and address potential issues before setting up Data Guard.
- Auto Setup (Inline) of Network During Provisioning of VM Clusters
  During deployment of VM clusters on Oracle Exadata Database Service on Exascale Infrastructure, you can choose to deploy a default network during VM Cluster creation.
- Long-Term Retention Backup (LTR)
- Performance Hub & metrics available for databases running in Oracle Exadata Database Service on Exascale Infrastructure
  You can now use the Performance Hub tool and view metrics on cloud databases that run on Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).
- Multiple Standby Databases
  You can now create and manage multiple local and remote Oracle Data Guard standby databases linked to a primary database.
- 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.
- Scale ECPUs to Zero
  You can now scale ECPUs enabled per VM to zero on ExaDB-XS.

New Regions and Realms for Exadata Database Service on Exascale Infrastructure

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> - Where the region supports multiple availability domains (AD), only a single availability domain is supported with Oracle Exadata Database Service on Exascale Infrastructure.
> - Starting November 4, 2025, Oracle Database 19c is available in all regions where the ExaDB-XS service is available.

Release Date: August 18, 2026

ExaDB-XS is now available in the following region:

LEJ: Morocco West (Casablanca)

Release Date: July 28, 2026

ExaDB-XS is now available in the following region:

LIN: Italy Northwest (Milan)

Release Date: June 2, 2026

ExaDB-XS is now available in the following region:

ORD: US Midwest (Chicago)

Release Date: April 3, 2026

ExaDB-XS is now available in the following region:

MAD: Spain Central (Madrid)

Release Date: May 12, 2026

ExaDB-XS is now available in the following region:

ARN: Sweden Central (Stockholm)

Release Date: October 22, 2025

ExaDB-XS is now available in the following region:

MTZ: Israel Central (Jerusalem)

Release Date: September 9, 2025

ExaDB-XS is now available in the following regions:

- DXB: UAE East (Dubai)
- QRO: Mexico Central (Queretaro)
- MTY: Mexico Northeast (Monterrey)

Release Date: August 26, 2025

ExaDB-XS is now available in the following regions:

- AGA: Salt Lake City, UT
- AMS: Netherlands Northwest (Amsterdam)

Copyright © 2023, 2026, Oracle and/or its affiliates.

Release Date: June 10, 2025

ExaDB-XS is now available in the following region:

SCL: Chile Central (Santiago)

Release Date: April 29, 2025

ExaDB-XS is now available in the following regions:

- IAD: US East (Ashburn)
- XSP: Singapore West (Singapore)

Release Date: March 28, 2025

ExaDB-XS is now available in the following region:

CWL: UK West (Newport)

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

ExaDB-XS is now available in the following region:

YUL: Canada Southeast (Montreal)

Release Date: December 19, 2024

ExaDB-XS is now available in the following region:

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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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

Change in Backup Destination for New Tenancies in Select OCI Regions (Effective November 04, 2026)

Release date: November 2026

Effective November 04, 2026, Autonomous Recovery Service is the only available backup destination when enabling automatic backups in the OCI Console for tenancies created on or after November 04, 2026, in the following OCI regions:

LFI: US Gov East (Ashburn), **LUF**: US Gov West (Phoenix), **PIA**: US DoD North (Chicago), **RIC**: US DoD East (Ashburn), **TUS**: US DoD West (Phoenix), **BRS**: UK Gov West (Newport), **LTN**: UK Gov South (London), **STR**: EU Sovereign Central (Frankfurt), **VLL**: EU Sovereign South (Madrid)

For more information, see *Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console*.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Container Database (CDB) Snapshot and Clone

This feature allows you to create a snapshot or clone of a container database (CDB).

Release Date: October 2026

With this enhancement, the OCI Console and API will support CDB-level snapshot functionality, enabling you to create and manage snapshots of a container database. You can use an existing container database or its snapshots to create thick (full) or thin clones of the container database. These snapshot and clone capabilities are powered by Exascale redirect-on-write technology.

Key Benefits

- **Enables independent, fully production-representative development and test environments**
  Using Exascale redirect-on-write technology, you can create space-efficient database level clones from a consistent database baseline for individual developers, development teams, quality assurance, and release engineering. This enables parallel application development, regression testing, upgrade rehearsals, application validation, and troubleshooting without dependency on shared database environments.
- **Accelerates application delivery and release readiness**
  Rapid database snapshot and thin clone creation reduces the time and operational effort required to provision and refresh non-production database environments. Teams can maintain known-good database baselines and create clones as needed to support development, testing, validation, and release activities.

> **Note**
>
> This feature requires dbaastools version 26.3.1 or later.

Related Topics

- To Create and List Container Database Snapshot
  Using the console, you can create and list the snapshots of the container database (CDB) in a VM Cluster.
- To Create a Clone Database from a Snapshot
  Using the console, you can create a clone database from a container database (CDB) snapshot.
- To Create a Clone of an Existing Container Database
  Using the console, you can create a clone of an existing container database.
- To Delete a Container Database Snapshot
  Using the console, you can delete the snapshot of a container database (CDB) in a VM Cluster.
- Permissions and API operation details for CDB Snapshot

Flexible Software Release Selection for Exadata Database

Lifecycle operations now default to the latest available Database and Grid Infrastructure release updates, and Exadata System Software maintenance releases. Where release

Copyright © 2023, 2026, Oracle and/or its affiliates.

selection is supported, you can select certain older release updates or maintenance releases by acknowledging the message displayed.

Release Date: October 2026

Releases outside the supported self-service range require a Service Request.

These choices apply to supported lifecycle operations across Oracle Exadata Database Service on Cloud@Customer, Oracle Exadata Database Service on Dedicated Infrastructure, and Oracle Exadata Database Service on Exascale Infrastructure.

Existing databases and systems aren't automatically upgraded as a result of this change.

Related Topics

- Flexible Software Release Selection for Exadata Database
  Oracle applies release selection controls to Database and Grid Infrastructure release updates that you can select for lifecycle operations.
- Acknowledging Older Release Updates and Maintenance Releases in API Requests
  Some lifecycle operations allow you to use older Database and Grid Infrastructure release updates, and Exadata System Software maintenance releases instead of the latest eligible releases.
- 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
  Follow this procedure to create a database software image 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.
- To Create a Database in an Existing VM Cluster
  Learn how you can create your first or subsequent databases.
- 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
- To Create a Database from a Backup
  Learn how to use a backup to create a database on Oracle Exadata Database Service on Exascale Infrastructure.
- To Create a Database from the Last Backup
  Use this procedure to create a database from the last backup on Oracle Exadata Database Service on Exascale Infrastructure.
- 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 patch updates the Oracle Grid Infrastructure (GI) on the resource.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Managed MCP Server Support for OCI Database Services

OCI Database Services now supports managed Model Context Protocol (MCP) servers for connecting MCP-compatible AI agents and applications to supported database resources.

Release Date: October 2026

You can create an MCP server for Base Database Service, Oracle Exadata Database Service on Cloud@Customer, Oracle Exadata Database Service on Dedicated Infrastructure, and Oracle Exadata Database Service on Exascale Infrastructure and use the generated MCP endpoint to connect a compatible MCP client.

Managed MCP servers use Streamable HTTP and integrate with OCI Identity and Access Management (IAM) for authentication and authorization. Access to OCI resources continues to be governed by your existing IAM policies.

You can create and manage multiple MCP servers, control the resource categories available through each server, and enable or disable a server as needed.

This release supports read-only operations. Operations that modify OCI Database resources aren't available through MCP.

For more information, see Managed MCP Servers for OCI Database Services.

Key Pools for Oracle AI Database@AWS

You can now use key pools to manage customer-managed encryption keys for your CDBs and PDBs.

Release Date: October 2026

> **Note**
>
> - Update the PKCS#11 Driver to the latest version before using the Key Pool functionality. See Release Notes for the latest version.
> - To upgrade the driver, see Updating the Multicloud PKCS#11 Driver.

> **Note**
>
> - To migrate an existing database from an individual AWS KMS key to a Key Pool, see Migrate from an Individual AWS KMS Key to a Key Pool.
> - To create a database using an AWS KMS key with the dbaascli command, see Create a Database with an Individual AWS KMS Key Using dbaascli.

A key pool is an OCI-managed collection of registered AWS keys. It enables separate key for CDBs and PDBs and allows VM Clusters to use the registered keys for encryption and decryption.

With key pools, you can:

- Create an empty pool and define its key-reuse policy before adding keys.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Discover and register customer-owned AWS keys in a key pool.
- Associate existing keys with a key pool.
- Select a key from a key pool when provisioning a CDB.
- Create PDBs without selecting a key manually.
- Automatically assign keys from a key pool during key rotations.
- Review pool membership, key usage, and lifecycle activity.
- Replicate pool information across supported regions for Data Guard deployments.
- For databases created before key pools were introduced, continue using the assigned encryption key or migrate the database to a key pool.

Each AWS key can belong to only one key pool. Key-reuse policies help manage key consumption while maintaining separation between database environments. Audit information is available for key pool creation, changes, membership updates, movement, and deletion.

Related Topics

- Understanding Key Pools
  Key Pools help you organize and control the AWS keys used to encrypt your Oracle Databases.
- Register AWS KMS Key
  To enable AWS KMS for your Exascale VM Clusters, you must first register the AWS KMS key in OCI.
- Create a Database and Use AWS KMS as the Key Management Solution
  This topic describes only the steps for creating a database and using AWS KMS as the key management solution.

Automatic Data Safe Registration for New Databases

Newly provisioned container databases (CDBs) are now automatically registered with Oracle Data Safe by default, and newly created pluggable databases (PDBs) inherit the Data Safe registration policy from their parent CDB.

Release Date: October 2026

Any Oracle database registered with Data Safe for security management is referred to as target database.

Related Topics

- Automatic Registration of Databases with Oracle Data Safe
  You can now automatically register newly provisioned container databases (CDBs) with Oracle Data Safe.
- To Create a Database in an Existing VM Cluster
  Learn how you can create your first or subsequent databases.

AWS Key Management Service Integration for Exadata Database Service on Exascale Infrastructure@AWS

Exadata Database Service on Exascale Infrastructure@AWS supports integration with AWS Key Management Service (KMS).

Release Date: August 2026

Copyright © 2023, 2026, Oracle and/or its affiliates.

This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using AWS customer managed keys.

Related Topics

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

Data Guard Automatic Failover for Exadata Database Service on Exascale Infrastructure

Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS) now supports automatic failover powered by Oracle Data Guard automation.

Release Date: August 2026

This new capability combines Oracle Data Guard with OCI automation to deliver automatic failover, helping reduce downtime during planned maintenance and unplanned outages.

Oracle Data Guard provides robust protection for Exadata Database Service customers against production database and infrastructure outages. Oracle Data Guard Automatic Failover builds on these proven capabilities by further automating and simplifying database resiliency. With single-click setup, continuous health monitoring, and automated failover orchestration, customers can accelerate database recovery and minimize application disruption without requiring changes to application connection strings.

With this release, customers can enable Oracle-managed automatic failover to a standby database when the primary database becomes unavailable. This capability can be enabled and configured using the OCI Console, REST APIs, SDKs, or Terraform.

The failover workflow is automated and transparent to applications. An Observer continuously monitors database health and availability. Connection details for both the primary and standby databases are available for administrators to download and configure. When a failure is detected, Oracle uses Oracle Data Guard Fast-Start Failover technology to automatically initiate failover and reroute application connections to the standby database without requiring manual intervention.

Supported Capabilities

- Automatic failover for ExaDB-XS environments
- Integration with Oracle Data Guard Fast-Start Failover
- Automatic application connection rerouting during failover
- Support for planned switchover and unplanned failover scenarios
- Customer-controlled enablement of automatic failover
- Support for customer-selected failover targets
- Support through OCI Console, API/SDK, and Terraform

Availability and Recovery Improvements

- Reduced recovery time by automating failover operations and connection rerouting

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Reduced operational overhead during disaster recovery and outage events

> **Note**
>
> Before enabling automatic failover, ensure that ingress on port 7095 is allowed for the agent running in the DomU. See To Enable Automatic Failover for the prerequisite.

Related Topics

- Observers in a Fast-start Failover Configuration
- To enable Data Guard on Oracle Exadata Database Service on Exascale Infrastructure
  Learn how to set up an Oracle Data Guard group between databases.
- To Enable Automatic Failover
  Learn how to enable Automatic Failover for an Oracle Data Guard configuration.
- To Add Failover Targets to the Primary or Standby Database
  Learn how to add failover targets to a primary or standby database in an Oracle Data Guard configuration.
- To Disable Automatic Failover
  Learn how to disable Automatic Failover for an Oracle Data Guard configuration.
- Handling Standby Database Termination
  Understand the scenarios that would affect automatic failover when you terminate a standby database.

Data Guard Support to Convert Physical Standby to Snapshot Standby and Vice Versa

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

Release Date: August 2026

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

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

Related Topics

- To Convert a Physical Standby Database to Snapshot Standby Database
  A snapshot standby database is a fully updateable standby database created by converting a physical standby database into a snapshot standby database.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- To Convert a Snapshot Standby Database to Physical Standby Database
  Oracle recommends converting a snapshot standby database back to a physical standby database within 14 days.

Google Cloud Key Management Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud

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

Release Date: July 2026

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

Related Topics

- Google Cloud Key Management Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud
  Exadata Database Service on Exascale Infrastructure@Google Cloud now supports integration with Google Cloud Platform's Key Management Service (KMS).

Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure

Exadata Database Service on Exascale Infrastructure@Azure enables you to store your database's transparent data encryption (TDE) keys, also known as master encryption keys (MEKs) in either a file-based Oracle wallet or in the OCI Vault.

Release Date: July 2026

This feature enables Exadata Database Service on Exascale Infrastructure@Azure users to utilize Azure Key Vault (AKV) Managed HSM, AKV Premium, and AKV Standard for managing TDE MEKs. This integration allows applications, Azure services, and databases to use a centralized key management solution for enhanced security and simplified key lifecycle management.

Related Topics

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

Change in Backup Destination for New Tenancies in Select OCI Regions (Effective June 23, 2026)

Copyright © 2023, 2026, Oracle and/or its affiliates.

Release date June 2026

Effective June 23, 2026, Autonomous Recovery Service is the only available backup destination when enabling automatic backups in the OCI Console for tenancies created on or after June 23, 2026, in the following OCI regions:

AMS: Netherlands Northwest (Amsterdam), **BOG**: Colombia Central (Bogota), **GRU**: Brazil East (Sao Paulo), **HSG**: Indonesia North (Batam), **HYD**: India South (Hyderabad), **IAD**: US East (Ashburn), **ICN**: South Korea Central (Seoul), **JBP**: Malaysia West 2 (Kulai), **JED**: Saudi Arabia West (Jeddah), **LEJ**: Morocco West (Casablanca), **LHR**: UK South (London), **LIN**: Italy Northwest (Milan), **MAD**: Spain Central (Madrid), **MTZ**: Israel Central (Jerusalem), **NRQ**: Italy North (Turin), **ORF**: Spain Central (Madrid 3), **SCL**: Chile Central (Santiago), **VAP**: Chile West (Valparaiso), **VCP**: Brazil South East (Vinhedo), **YUL**: Canada Southeast (Montreal), and **YYZ**: Canada Southeast (Toronto).

For more information, see *Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console*.

Enhanced Controls for Managing Unresponsive Virtual Machines

This release introduces enhanced controls for managing unresponsive virtual machines (VMs), providing clearer behavior and safer handling of reboot and stop operations.

Release Date: June 2026

When a VM is unresponsive, you can select the force option when you reboot or stop it.

Available actions:

- Reboot
- Force reboot
- Stop
- Force stop

When performing a Force Reboot or Force Stop, you must type the VM name to confirm that the action is performed on the intended VM.

If a VM remains unresponsive for 30 minutes after you select Reboot or Stop, the system automatically uses Force reboot or Force stop to complete the operation.

If a Force Reboot does not complete successfully, Oracle recommends that you raise a Service Request through the Oracle Support Center as a last resort.

These enhancements improve control, predictability, and safety when managing unresponsive VMs, while reducing the risk of conflicting actions and accidental disruptions.

Related Topics

- Using the Console to Stop, Start, or Reboot a Virtual Machine
  Use the console to stop, start, or reboot a virtual machine.
- To Check the Status of a VM Cluster Virtual Machine
  Review the health status of a VM cluster virtual machine.
- Terminate or Remove a VM from a VM Cluster
  To terminate or remove a virtual machine from a provisioned cluster, use this procedure.
- VM Node Subsetting Event Types
  Review the list of event types that VM Node Subsetting emits.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Oracle Managed Database Software Updates

Oracle Managed Database Software Updates gives you the option to let Oracle manage the end-to-end database software update lifecycle for enrolled databases.

Release Date: June 2026

You define your scheduling preference for database software updates, while Oracle automatically prepares, validates, and applies the updates.

For databases enrolled in Oracle Managed Database Software Updates, Oracle periodically validates update readiness, stages software in advance, applies updates during the scheduled maintenance window, and provides proactive notifications throughout the update process.

You can enable Oracle Managed Database Software Updates for an existing database or while creating a new database. Once enabled, Oracle will keep database software updated with the latest available Release Update (RU) according to your scheduling preferences

Key benefits and capabilities include:

- **Simplified routine maintenance** – Reduce the operational effort required to plan, prepare, apply, and maintain database software updates.
- **Maintenance scheduling flexibility** – Define your preferred day and time of the week for Oracle to automatically apply database software updates when a new Release Update (RU) is available.
- **Stay current on Database Release Updates** – Enrolled databases are automatically updated to the latest available RU in a rolling manner, one database instance at a time.
- **Proactive readiness checks with remediation** – Oracle periodically checks update readiness, notifies you of any customer-actionable issues, and recommends fixes so you can resolve them before the maintenance window.
- **Database home lifecycle management** – Oracle manages database software homes throughout the update lifecycle, including cleanup of homes that are no longer required.
- **Successful maintenance outcome** – Oracle proactively monitors and resolves issues encountered during planned maintenance. If an issue cannot be fixed during the planned maintenance window, Oracle rolls back the update to ensure the database remains healthy and available.

> **Note**
>
> This feature is now available in all OC1 regions.

Related Topics

- About Oracle Managed Database Software Updates
  Oracle Managed Database Software Updates gives you the option to let Oracle manage the end-to-end database software update lifecycle for enrolled databases.
- To Create a Database in an Existing VM Cluster
  Learn how you can create your first or subsequent databases.
- Enable Oracle Managed Database Software Updates for an Existing Database
  Learn how you can modify an existing database to support Oracle Managed Database Software Updates for Oracle Exadata Database Service on Exascale Infrastructure.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- To Move a Database to Another Database Home
  Learn how to move a database to another Database home.
- To Reschedule Oracle Managed Database Software Updates for a Database
  Learn how to reschedule the Oracle managed database software updates for Oracle Exadata Database Service on Exascale Infrastructure.
- To View and Resolve Oracle Managed Database Software Updates Issues
  Learn how to view and resolve the Oracle managed database software updates issues for Oracle Exadata Database Service on Exascale Infrastructure.
- To Create a Database from a Backup
  Learn how to use a backup to create a database on Oracle Exadata Database Service on Exascale Infrastructure.
- To Create a Database from the Last Backup
  Use this procedure to create a database from the last backup on Oracle Exadata Database Service on Exascale Infrastructure.
- 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.
- Troubleshooting Oracle Managed Database Software Updates Issues
  Learn to identify and resolve issues in databases that are enrolled for Oracle Managed Database Software Updates.

Delegate Access Control for Exadata Database Service on Exascale Infrastructure

Delegate Access Control service is enabled for Oracle Exadata Database Service on Exascale Infrastructure.

Release Date: June 2026

Delegate Access Control is a service that enhances the security and manageability of Oracle Exadata environments by enabling customers to govern how service provider operators access their resources.

It allows customers to retain full control over when, how, and for how long access is granted, while still enabling Oracle service providers to deliver maintenance and support efficiently.

Related Topics

- Overview of Delegate Access Control
- To Create a Database in an Existing VM Cluster
  Learn how you can create your first or subsequent databases.
- Exadata Database Service Dedicated Security Controls
- Oracle Engineered System Deployment and Infrastructure Support

Change in Backup Destination for New Tenancies in Select OCI Regions (Effective May 15, 2026)

Release date May 2026

Copyright © 2023, 2026, Oracle and/or its affiliates.

Effective May 15, 2026, Autonomous Recovery Service is the only available backup destination when enabling automatic backups in the OCI Console for tenancies created on or after May 15, 2026, in the following OCI regions:

ARN: Sweden Central (Stockholm), **AUH**: UAE Central (Abu Dhabi), **BOM**: India West (Mumbai), **CDG**: France Central (Paris), **CWL**: UK West (Newport), **DXB**: UAE East (Dubai), **JNB**: South Africa Central (Johannesburg), **KIX**: Japan Central (Osaka), **MEL**: Australia Southeast (Melbourne), **MRS**: France South (Marseille), **MTY**: Mexico Northeast (Monterrey), **ORD**: US Midwest (Chicago), **QRO**: Mexico Central (Queretaro), **RUH**: Saudi Arabia Central (Riyadh), **SIN**: Singapore (Singapore), **SJC**: US West (San Jose), **SYD**: Australia East (Sydney), **XSP**: Singapore West (Singapore), **YNY**: South Korea North (Chuncheon), and **ZRH**: Switzerland North (Zurich).

For more information, see *Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console*.

Cross-Service Data Guard Operations for Oracle Database 19c

You are now able to deploy Oracle Data Guard environments with cross-service primary and secondary databases for Oracle Database 19c.

Release Date: April 2026

We are pleased to announce cross-service Oracle Data Guard deployment support for Oracle Database 19c. In a cross-service deployment, you set up primary and secondary databases between two services: Exadata Database Service on Dedicated Infrastructure (ExaDB-D) and Exadata Database Service on Exascale Infrastructure (ExaDB-XS). The ability to deploy cross-service Oracle Data Guard provides enhanced availability. The following configurations are available:

- Primary database on ExaDB-D with one or more standby databases on ExaDB-XS or ExaDB-D.
- Primary database on ExaDB-XS with one or more standby databases on ExaDB-D or ExaDB-XS.

> **Note**
>
> For Oracle Database 19c, ExaDB-XS supports only the Data Guard Group resource.

Related Topics

- Cross-Service Interoperability Between ExaDB-D and ExaDB-XS: Data Guard, Backup, and Restore

Protect Sensitive Data With Oracle Cloud Infrastructure Zero Trust Packet Routing

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

Release date February 2026

Copyright © 2023, 2026, Oracle and/or its affiliates.

Related Topics

- To create a cloud VM cluster
  Create a VM cluster in an Oracle Exadata Database Service on Exascale Infrastructure instance.

Managing Certificates using DBCS Certificate Management

Database Cloud Services (DBCS) Certificate Management lets you centrally manage X.509 certificates across all DBCS-managed Oracle Database services. This service is HA-aware (High Availability) and includes features that provide connection details to facilitate seamless connectivity to all Data Guard peer databases.

Release date January 2026

With this release, Oracle installs Certificate Authority (CA) and leaf certificates for the provisioned PDBs and also manages the rotation of CA and leaf certificates. You can download the connection bundle (`DbConnectionBundle`) from the PDB to trust the CA certificates and connect to your PDBs using TCPS.

The Connection Bundle for a PDB is auto-refreshed when its Data Guard group changes. Therefore, when you download the bundle, you have the most up-to-date, HA-aware connection strings for your application.

Manage Serial Console Access in Oracle Exadata Database Service on Exascale Infrastructure

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

Release date January 2026

> **Note**
>
> - The serial console feature requires (at a minimum) Exadata System Software 23.1.13. Once the necessary software is installed via Quarterly Maintenance and a reboot of your VMs takes place, you will be able to use the new serial console feature.
> - If the serial console connection is unexpectedly lost during any operation, please delete the existing console connection and create a new one.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

Enhanced Out-of-Place Restore

ExaDB-XS currently supports database creation from a specific backup and from latest backup. This new functionality enables you to create database from a specified point-in-time (timestamp).

Release date January 2026

Enhanced Data Guard Health Insights

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

Release date November 2025

Virtual Machine (VM) Snapshots

This feature allows customers to create and manage snapshots of their virtual machines (VM) and store them locally on Exascale storage.

Release date November 2025

This new capability introduces self-service VM file system snapshots, enabling customers to capture and preserve point-in-time images of their virtual machine file systems directly through the Oracle Cloud Console or API. The feature is built on Exascale's redirect-on-write snapshot technology that delivers instant data protection and operational agility by efficiently recording the exact state of VM file systems with exceptional storage efficiency.

Auto scaling Exascale Storage Vault

The Exascale Storage Vaults can now scale automatically as usage grows, with storage expanding incrementally up to a user-defined limit.

Release Date: September 2025

Previously, only manual scaling up or down of the Exascale Storage Vault was available. This new feature, auto scaling for the Exascale Storage Vault, adds agility to Exadata Database Service on Exascale Infrastructure environments by allowing storage to expand automatically as usage grows, without manual intervention.

When auto scaling is enabled, if the system detects that the Vault's free space is less than or equal to 5% of the allocated storage capacity, it automatically expands the Vault's capacity. The Vault capacity increases in increments of 10% of the allocated storage capacity until the user-defined auto scale limit is reached.

This eliminates the need for customers to manually increase capacity when the Vault usage nears its limits. The storage scales dynamically, ensuring mission-critical workloads continue to run smoothly and without interruption.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Note

- The feature is not available with Oracle Database 19c.
- Auto scale is primarily intended to accommodate the ongoing, organic growth of databases. For large, sudden data loads or rapid data growth that can exceed 5% of the allocated vault size within a short time period, it is recommended to manually scale storage in advance to avoid potential out-of-space errors.
- Tablespaces with auto-extend enabled can cause out-of-space errors if the auto-extend increment size is greater than 5% of the allocated vault storage. This could be caused if a single, large tablespace auto-extension suddenly exceeds the available space in the vault before crossing the auto scale threshold. Or, it could also be caused by multiple tablespaces extensions occurring concurrently, resulting in the same effect. Oracle recommends that you manage your tablespace auto-extend increment settings to avoid having more than 5% of the allocated vault storage added at once when using auto scale.

Oracle Database 19c Support

Oracle is excited to announce that Exadata Database Service on Exascale Infrastructure (ExaDB-XS) now supports Oracle Database 19c and Oracle Grid Infrastructure 19c.

Release Date: August 2025

Last year, Oracle launched ExaDB-XS with Oracle AI Database 26ai, introducing the industry's only intelligent data architecture for the cloud – a revolutionary leap that combines the best of Exadata and the best of the cloud. ExaDB-XS delivers Exadata benefits at a lower entry cost with granular scalability, making it an ideal platform for a wide range of mission-critical databases. You can easily migrate your on-premises databases to shared Exascale Infrastructure in the cloud. Until now, these benefits were available only to customers ready to adopt to Oracle AI Database 26ai. Many organizations with mission-critical workloads on Oracle Database 19c are not yet ready to upgrade to Oracle AI Database 26ai. These organizations can now immediately realize the benefits of Exascale Infrastructure for Oracle Database 19c with an option to upgrade to Oracle AI Database 26ai at their own pace.

Key benefits

- **Extreme Low Cost** - Benefit from pay-per-use pricing, low upfront costs, and no cost per IO, delivering Exadata performance at extremely low cost
- **Scalable Pooled Resources** - Fully elastic architecture allows you to instantly allocate resources as needed.
- **Powerful** - Exadata performance, reliability, availability, and security capabilities are integrated by default.

Pluggable Database (PDB) Snapshots

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

Release date August 2025

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> This feature is available starting with Oracle AI Database 26ai and later versions.

The PDB level snapshot functionality in the OCI console and API enables you to create and manage PDB level snapshots and create clones (including thin clones). It leverage space efficient Exascale redirect-on-write technology.

Cross-Service Data Guard Operations

You are now able to deploy Oracle Data Guard environments with cross-service primary and secondary databases.

Release Date: August 2025

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

- Primary database on ExaDB-D with one or more standby databases on ExaDB-XS or ExaDB-D.
- Primary database on ExaDB-XS with one or more standby databases on ExaDB-D or ExaDB-XS.

> **Note**
>
> For Oracle Database 19c, ExaDB-XS supports only the Data Guard Group resource.

Related Topics

- Cross-Service Interoperability Between ExaDB-D and ExaDB-XS: Data Guard, Backup, and Restore

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

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

Release date August 2025

For more information, see the following:

Oracle Recommended Options to Perform Backup and Recovery Operations

Effective August 06, 2025, when you enable automatic backups in the OCI Console, the Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after 06 August 2025.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- The database is deployed in OCI regions Frankfurt (FRA), Phoenix (PHX), and Tokyo (NRT).

The Oracle Database version is later than 19.18 or 23.4.

If these conditions are not met, OCI Object Storage will be shown as a Backup destination.

Expanded Tagging Support Across Database Workflows

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

Release Date: July 2025

Data Guard Setup with Pre-Check Validation

You are now able to perform a precheck before setting up Oracle Data Guard, and address potential issues before setting up Data Guard.

Release Date: June 2025

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

Auto Setup (Inline) of Network During Provisioning of VM Clusters

During deployment of VM clusters on Oracle Exadata Database Service on Exascale Infrastructure, you can choose to deploy a default network during VM Cluster creation.

Release Date: May 2025

We are pleased to announce the general availability (GA) of deploying a default network during VM Cluster creation. When you select the option to create a default VCN, a new VCN is created for you with default settings that auto-populate with network configurations for your tenancy. Previously, when creating a new VM Cluster, it was necessary to create a VCN in-advance, meeting various specific requirements. Now, you can select an option to have the cloud automation create a default new VCN setup for you. This is a convenient way to set up a new VM Cluster quickly, without advance planning of networking, which can be especially useful for simple testing use cases.

Long-Term Retention Backup (LTR)

Release Date: May 2025

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

Performance Hub & metrics available for databases running in Oracle Exadata Database Service on Exascale Infrastructure

You can now use the Performance Hub tool and view metrics on cloud databases that run on Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).

Release Date: April 2025

This feature provides additional monitoring and management functions to databases on ExaDB-XS. For more information, see Analyzing Exadata Cloud Service Database Performance and Analyzing Virtual Machine / Bare Metal Database Performance .

Multiple Standby Databases

You can now create and manage multiple local and remote Oracle Data Guard standby databases linked to a primary database.

Release Date April.2025

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

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

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

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

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

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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, 2025

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.

Scale ECPUs to Zero

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

Release Date: January 14, 2025

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.
- Cross-Service Interoperability Between ExaDB-D and ExaDB-XS: Data Guard, Backup, and Restore
  You can now create a cross-service Oracle Data Guard group across database services.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

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

- Getting Started with Policies
- Common Policies
- Policy Details for the Database Services

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

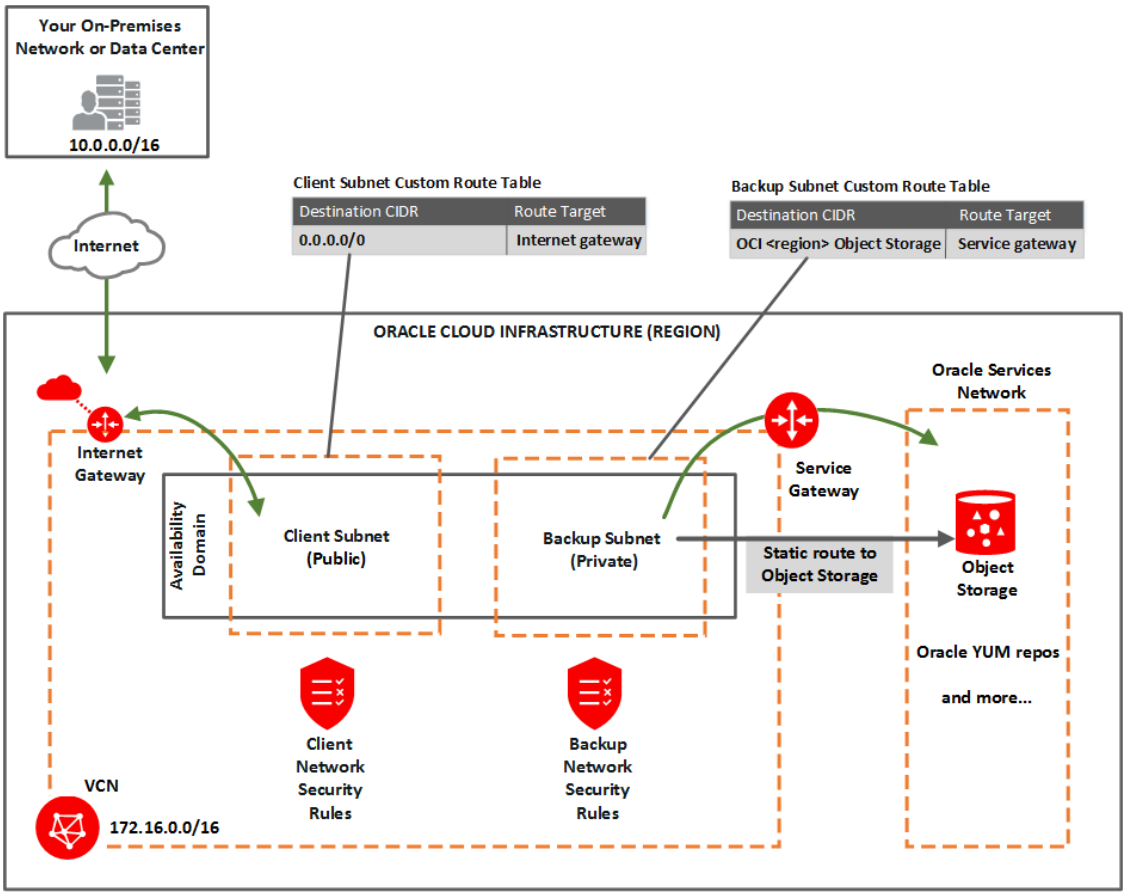

You set up:

- 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 Sevice 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).

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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.

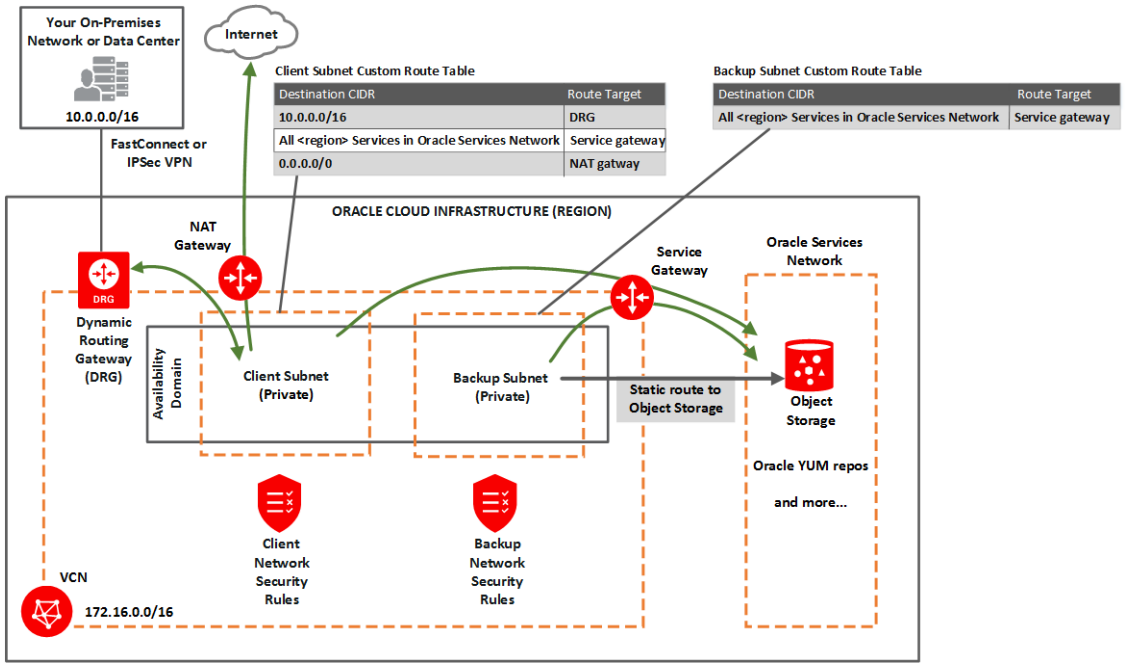

You set up:

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

        * 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 through the client subnet, for example, 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 resource

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

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

For example:

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

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

- Node 1: `exacs-abcde1.clientpvtad1.acmevcniad.oraclevcn.com`
- Node 2: `exacs-abcde2.clientpvtad1.acmevcniad.oraclevcn.com`

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Node 3: `exacs-abcde3.clientpvtad1.acmevcniad.oraclevcn.com`
- And so on

Requirements for the hostname prefix:

- Recommended maximum: 23 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 128 characters, out of which 16 characters are allocated for service-added characters. Here is a safe general rule:

  `<23_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 VM cluster provisioning. For details, see Private DNS resolvers.
- Forwarding to another DNS server should be set up beforehand in the DNS console. This can be done by going to the VCN's resolver, and creating the endpoint and then the rules. For details, see DNS in Your Virtual Cloud Network.
- Private zone's name cannot have more than 4 labels. For example, a.b.c.d is allowed while a.b.c.d.e is not.
- It is also required to add the private view to the resolver of the VCN. For details, see Adding a Private View to a Resolver.
- When provisioning an Exadata VM Cluster using the 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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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, Oracle recommends that you configure Autonomous Recovery Service. Regardless of how you configure the VCN with that access (for example, with a service gateway), if you use Object Storage, 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.

Note

Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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, network setup with a public client subnet:

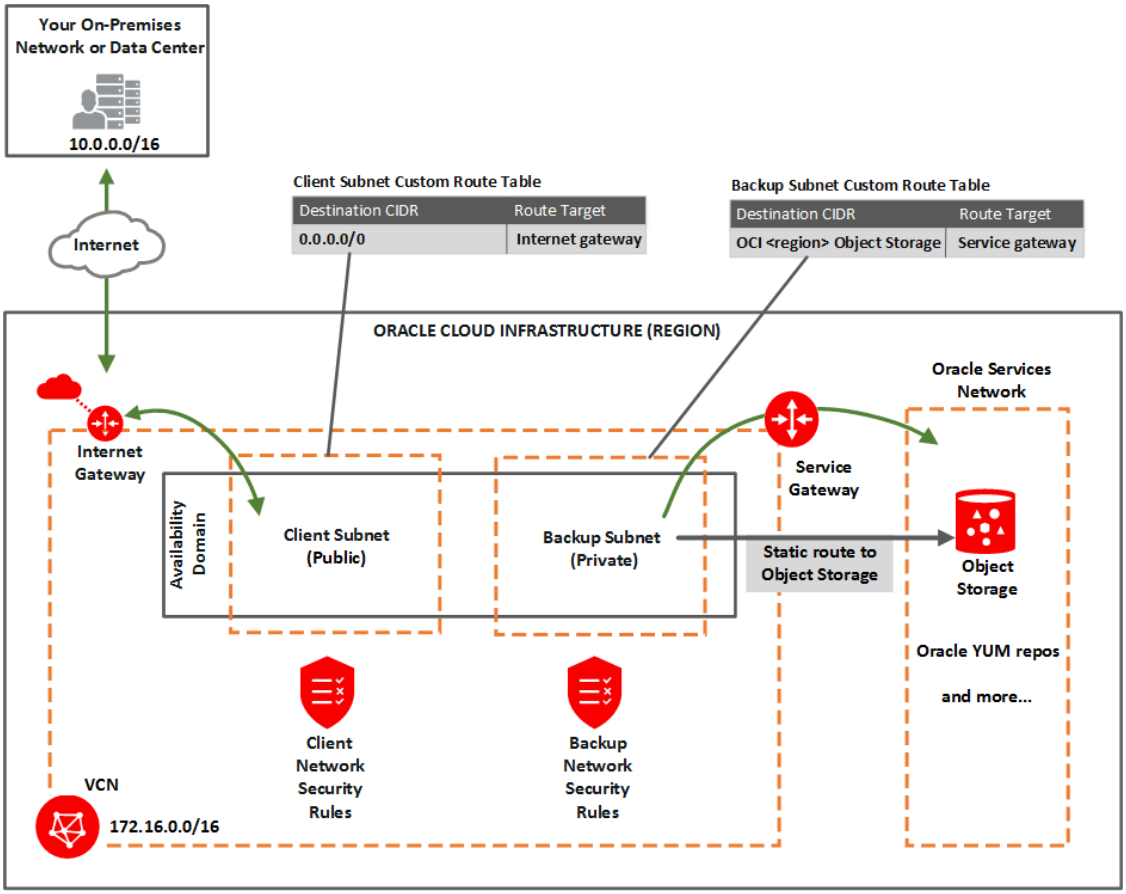

In general, you must:

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 VM cluster to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repositories if the client network doesn't have access to them).

Option 2: Service Gateway Access to 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, a network setup with private subnets:

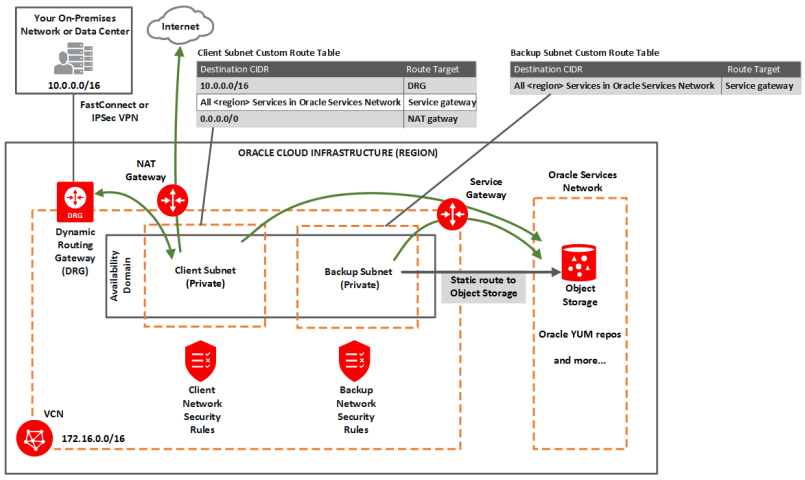

In general, you must:

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

> **Note**
>
> Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

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

Rules Required for Both the Client Network and Backup Network

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **Description:** Optionally, add a meaningful description of the rule. For example: "Allow access to Exadata Fleet Update private endpoint within the subnet."

Client Ingress Rule 4: Allows automatic failover traffic to the DomU agent

This rule is required when Fast-Start Failover is enabled for any Data Guard-enabled database. It allows ingress traffic on port 7095 to reach the agent running in the DomU.

- **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:** 7095
- **Description:** Allow ingress traffic to the DomU agent for automatic failover.

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 2 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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 VM cluster to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repos if the client network doesn't have access to them). It is redundant with the general egress rule in this topic (and in the default security list). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

Backup egress rule: Allows access to Object Storage

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

- Rules Required for Both the Client Network and Backup Network
  This topic has several general rules that enable essential connectivity for hosts in the VCN.
- Rules Required Specifically for the Client Network
  The following security rules are important for the client network.
- Rule Required Specifically for the Backup Network
  The following security rule is important for the backup network because it enables the VM cluster to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repositories if the client network doesn't have access to them).
- Rules Required for Events Service
  The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Events service.
- Rules Required for Monitoring Service
  The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Monitoring service.

Rules Required for Both the Client Network and Backup Network

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

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

- General ingress rule 1: Allows SSH traffic from anywhere
- General ingress rule 2: Allows Path MTU Discovery fragmentation messages

Copyright © 2023, 2026, Oracle and/or its affiliates.

- General ingress rule 3: Allows connectivity error messages within the VCN
  This rule enables the hosts in the VCN to receive connectivity error messages from each other.
- General egress rule 1: Allows all egress traffic

Related Topics

- default security list

General ingress rule 1: Allows SSH traffic from anywhere

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

General ingress rule 2: Allows Path MTU Discovery fragmentation messages

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

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

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

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

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

General egress rule 1: Allows all egress traffic

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Rules Required Specifically for the Client Network

The following security rules are important for the client network.

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

- Client ingress rule 1: Allows ONS and FAN traffic from within the client subnet
  The first rule is recommended and enables the Oracle Notification Services (ONS) to communicate about Fast Application Notification (FAN) events.
- Client ingress rule 2: Allows SQL*NET traffic from within the client subnet
  This rule is for SQL*NET traffic and is required in these cases:
- Client egress rule 1: Allows all TCP traffic inside the client subnet
  This rule is for SQL*NET traffic as noted.
- Client egress rule 2: Allows all egress traffic (allows connections to the Oracle YUM repos)
  Client egress rule 3 is important because it allows connections to the Oracle YUM repos.

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

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

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

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

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

- If you need to enable client connections to the database

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

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

This rule is for SQL*NET traffic as noted.

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

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

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

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

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

Related Topics

- default security list

Rule Required Specifically for the Backup Network

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

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

Related Topics

- General egress rule 1: Allows all egress traffic
- default security list

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 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- If you use network security groups
- If you use security lists
  If you choose to use security lists, follow this 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 Oracle 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, follow this 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

> **Warning**
>
> **Do not remove the default egress rule from the default security list**. If you do, make sure to instead include the following replacement egress rule in the client subnet's security list:
>
> - **Stateless:** No (all rules must be stateful)
> - **Destination Type:** CIDR
> - **Destination CIDR:** 0.0.0.0/0
> - **IP Protocol:** All

Network Requirements for Oracle Database Autonomous Recovery Service

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

To use Recovery Service for backups, follow the steps outlined in Onboarding Oracle Database to Recovery Service.

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

Create a Service Gateway to Object Storage

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

> **Note**
>
> Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Cross-Service Interoperability Between ExaDB-D and ExaDB-XS: Data Guard, Backup, and Restore

You can now create a cross-service Oracle Data Guard group across database services.

With this feature, you can establish cross-service backup and restore capability for the following configurations:

- Primary database on ExaDB-D with one or more standby databases on ExaDB-XS or ExaDB-D.
- Primary database on ExaDB-XS with one or more standby databases on ExaDB-D or ExaDB-XS.

This capability enables you to leverage the potential of your database services across your tenancy.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> To ensure maximum availability, Oracle recommends that you place your peer VM clusters in a different Exadata infrastructure from the primary VM cluster.

Standby Database Configuration Options

When adding a standby database, you can specify the following details:

- **Peer VM Cluster Details:** If the target is ExaDB-D, you can select the Exadata Infrastructure.
- **Target Service Selection:** Choose either ExaDB-D or ExaDB-XS. By default, the service initiating the Data Guard group is selected. If a service is unavailable in the selected region, it is disabled with the message: "Service is not available in this region".
- **Data Guard Type:** The group can be configured with either Data Guard or Active Data Guard, and each standby database can have a different type.
- **Data Guard Group Limitation:** A primary database is limited to one Data Guard group.
- **Standby Database Creation:** Only one standby database can be added at a time. However, standby databases can be created in any of the following categories without restrictions on their number:
  - Within the same service
  - Across services
  - Within the same Availability Domain (AD)
  - Across Availability Domains within the same region
  - Across regions

Key Considerations for Cross-service Primary and Standby Databases

- Both the primary and standby databases must use the same key management solution.
- Data Guard can be configured in Max Performance or Max Availability protection mode with Sync or Async transport type, subject to the following rules:
  - For the first standby database:
    - Defaults to Max Performance mode with Async transport.
    - Protection mode and transport type cannot be changed.
  - For the second and subsequent standby databases:
    - Inherits the protection mode from the first standby.
    - Defaults to Async transport.
    - Protection mode and transport type cannot be changed.

Viewing and Editing Data Guard Configurations

- View all databases that are part of the Data Guard group in the Data Guard Group table, irrespective of whether you are on the primary or standby database pages.
- Edit the Data Guard configuration to update:
  - **Data Guard type:** Can be different for each standby database. This can be changed only from a standby database.

Copyright © 2023, 2026, Oracle and/or its affiliates.

  - **Database administrator password:** Can be edited from any database.
  - **Protection mode:** Can be switched between Max Performance and Max Availability from the primary or any standby database.
  - **Transport type:** Can be switched between Async and Sync only from a standby database.

> **Note**
>
> If the Protection Mode is set to Max Availability, then the system verifies that at least one standby database uses Sync transport. If not found, then an error message is displayed:
>
> "To achieve zero data loss, you require at least one standby with Sync transport. It is recommended to have a standby with Sync transport in the same service as the primary database."

Switchover and Failover of Databases

- **Switchover (on any standby database)**
  - Switchover is performed without enforcing a major version or Release Update (RU) check. However, a warning appears if the standby database has asymmetric configuration, such as differing node counts, memory, or service type.
- **Failover (on any standby database)**
  - Failover is performed without enforcing a major version or Release Update (RU) check. However, a warning appears if the standby database has an asymmetric configuration, such as differing node counts, memory, or service type.

Database Backup and Restore Options

> **Note**
>
> Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

- **Scheduled and Automatic Backups**
  - You can schedule automatic backups on the primary database, standby database, or both.
  - Both Object Storage and Recovery Service backups are supported.
  - If backups are configured on both primary and standby databases, they must use the same backup destination type.
- **In-Place Restore of the Same Database within ExaDB-XS**
  Restore and recover a database using a backup taken from the same database in ExaDB-XS:
  - Restore primary using a backup taken on the primary database.
  - Restore standby using a backup taken on the standby database.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **In-Place Restore of a Peer Database within ExaDB-XS**
  Restore and recover a peer database (which has no backups configured) using a backup from the source database with Recovery Service:
  - **Scenario 1:** Restore the primary using the standby backup.
    - **Primary database:** ExaDB-XS (no backups configured)
    - **Standby database:** ExaDB-XS (backups configured to Recovery Service)
  - **Scenario 2:** Restore the standby using the primary backup.
    - **Primary database:** ExaDB-XS (backups configured to Recovery Service)
    - **Standby database:** ExaDB-XS (no backups configured)
- **In-Place Restore of a Peer Database Across Services**
  Restore and recover a database in ExaDB-XS or ExaDB-D using a backup from the source database in the opposite service (ExaDB-D or ExaDB-XS) with Recovery Service:
  - **Scenario 1:** Restore a peer database in ExaDB-XS using a backup from ExaDB-D
    - **Use Case 1:** Restore primary using the standby backup.
    - **Use Case 2:** Restore standby using the primary backup.
  - **Scenario 2:** Restore a peer database in ExaDB-D using a backup from ExaDB-XS
    - **Use Case 1:** Restore primary using the standby backup.
    - **Use Case 2:** Restore standby using the primary backup.

Out-of-place Restore (Creating a New Database)

- **Within ExaDB-XS**You can create a new database in ExaDB-XS using a backup from the same service.
  - Restore within:
    - The same Availability Domain (AD)
    - A different AD within the same region
    - A different region
  - Supports Object Storage or Recovery Service as the backup destination.
- **Across Services**
  - **Scenario 1:** Create a New Database in ExaDB-D Using a Backup from ExaDB-XS
    - **Source:** ExaDB-XS (database and backups)
    - **Target:** ExaDB-D (any region or AD)
    - **Backup Destination:** Object Storage or Recovery Service
  - **Scenario 2:** Create a New Database in ExaDB-XS Using a Backup from ExaDB-D
    - **Source:** ExaDB-D (database and backups)
    - **Target:** ExaDB-XS (any region or AD)
    - **Backup Destination:** Object Storage or Recovery Service

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

- Regional Availability for Exadata Database Service on Exascale Infrastructure
  Learn what OCI regions support Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).
- 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).

Regional Availability for Exadata Database Service on Exascale Infrastructure

Learn what OCI regions support Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).

In the tables in this topic, you can find which Oracle Cloud Infrastructure (OCI) regions offer Exadata Database Service on Exascale Infrastructure.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

Asia Pacific (APAC)

Region Name	Region Key	Realm Key	Region Location
Australia East	SYD	OC1	Sydney, Australia
Australia Southeast	MEL	OC1	Melbourne, Australia
India South	HYD	OC1	Hyderabad, India
India West	BOM	OC1	Mumbai, India
Japan Central	KIX	OC1	Osaka, Japan
Japan East	NRT	OC1	Tokyo, Japan
Singapore	SIN	OC1	Singapore, Singapore
Singapore West	XSP	OC1	Singapore, Singapore
South Korea Central	ICN	OC1	Seoul, South Korea
South Korea North	YNY	OC1	Chuncheon, South Korea

Europe, Middle East, Africa (EMEA)

Region Name	Region Key	Realm Key	Region Location
EU Sovereign Central	STR	OC19	Frankfurt, Germany
EU Sovereign South	VLL	OC19	Madrid, Spain
France Central	CDG	OC1	Paris, France
France South	MRS	OC1	Marseille, France
Germany Central	FRA	OC1	Frankfurt, Germany
Israel Central	MTZ	OC1	Jerusalem, Israel
Italy Northwest	LIN	OC1	Milan, Italy
Morocco West	LEJ	OC1	Casablanca, Morocco
Netherlands Northwest	AMS	OC1	Amsterdam, Netherlands
Saudi Arabia West	JED	OC1	Jeddah, Saudi Arabia
South Africa Central	JNB	OC1	Johannesburg, South Africa
Spain Central	MAD	OC1	Madrid, Spain
Sweden Central	ARN	OC1	Stockholm, Sweden
Switzerland North	ZRH	OC1	Zurich, Switzerland
UAE Central	AUH	OC1	Abu Dhabi, UAE
UAE East	DXB	OC1	Dubai, UAE
UK Gov South	LTN	OC4	London, United Kingdom
UK Gov West	BRS	OC4	Newport, United Kingdom
UK South	LHR	OC1	London, United Kingdom
UK West	CWL	OC1	Newport, United Kingdom

Copyright © 2023, 2026, Oracle and/or its affiliates.

Latin America (LATAM)

Region Name	Region Key	Realm Key	Region Location
Brazil East	GRU	OC1	Sao Paulo, Brazil
Brazil Southeast	VCP	OC1	Vinhedo, Brazil
Chile Central	SCL	OC1	Santiago, Chile
Colombia Central	BOG	OC1	Bogota, Colombia

North America (NA)

Region Name	Region Key	Realm Key	Region Location
Canada Southeast	YUL	OC1	Montreal, Canada
Canada Southeast	YYZ	OC1	Toronto, Canada
US East	IAD	OC1	Ashburn, VA
US West	PHX	OC1	Phoenix, AZ
US West	SJC	OC1	San Jose, CA
US Midwest	ORD	OC1	Chicago, IL
Mexico Central	QRO	OC1	Queretaro, Mexico
Mexico Northeast	MTY	OC1	Monterrey, Mexico

> **Note**
>
> To provision Exadata Database Service on Exascale resources in a supported region, your OCI tenancy must be subscribed to that region. For example, to provision resources for Exadata Database on Exascale Infrastructure in the East US region, your OCI tenancy must be subscribed to US East (Ashburn). See Managing Regions and Subscribing to an Infrastructure Region for more information.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

> **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 defined tags that can apply to most resources, which are available from the **Namespace** field list. 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

reporting and governance, OCI service automation can deliver workload-specific optimizations based on standard tag values. You can also create freeform tags with your own key names in the **Key** field, and your own values in the **Value** field. When you have provided your settings, click **Add tag** to add the tag.

Figure 4-1 Tagging Example

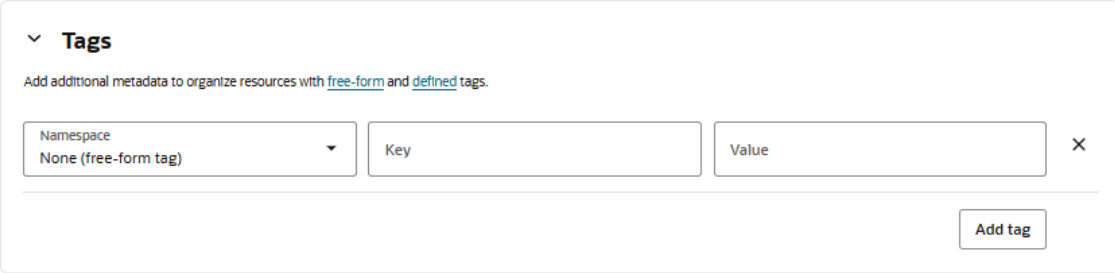

Oracle Standard Tags

Your tenancy governance administrators can deploy the standard defined 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 Defined Tags

Tag Key	Tag Value Options	Description
`OracleStandard.Criticality`	• Extreme • High • Medium • Low	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. For example, a database resource with `OracleStandard.Criticality` set to Extreme or High may require the highest availability SLA and may need to be configured with Autonomous Data Guard.
`OracleStandard.Environment`	• Dev • Test • Prod • Pre-Prod • Staging • Trial • Sandbox • User Testing	Denotes a resource lifecycle. In the case of databases, it helps determine consolidation density, database distribution across containers, set maintenance plans, and manage clones.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 4-1 (Cont.) Oracle Standard Defined Tags

Tag Key	Tag Value Options	Description
`OracleStandard.Sensitivity`	• Public • Internal • Sensitive • Highly Sensitive • Extremely Sensitive	An application or database classification tag. `OracleStandard.Sensitivity` set to Highly Sensitive may indicate that an access control list or certain Network Security Group (NSG) enforcement is mandatory to restrict access.
`OracleStandard.Regulation`	Refer to *List of Compliance Regulations* 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 *Using Predefined Values* for more details.
`OracleStandard.TargetUsers`	• Public • Customers • Partners • Company • Division • Department • Workgroup	Denotes the end users of a resource. Another form of resource classification that helps determine target users and allows governance teams to set corporate standards based on user or application type.
`OracleStandard.EndUserCount`	• 1 • 10 • 100 • 1000 • 10000 • 100000 • 1000000 • 10000000	An approximate count of end-users. This tag helps determine the number of users impacted or the blast radius during an availability or security event. This also helps prioritize recovery efforts in the event of major outages affecting a large number of cloud resources.
`OracleStandard.OwnerEmail`	Free form tag. For example *john.smith@example.com* or *app_support_grp@example.com*	Denotes the email address of the resource owner.
`OracleStandard.Org`	• HR • Finance • Marketing • Sales • Legal • R&D • Customer Support • Internal Support • Manufacturing	Identifies the customer's line of business or department that owns or uses the resource. This may help with cost aggregation reports and determining usage across business units. Tag administrators may add relevant values to the list from the OCI Governance and Administration console. Refer to *Using Predefined Values* for more details.
`OracleStandard.CostCenter`	• 12345 • WebMarketing	Freeform field for cost center.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 4-1 (Cont.) Oracle Standard Defined Tags

Tag Key	Tag Value Options	Description
`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**
>
> You will generally be notified of planned maintenance at least two weeks in advance of the scheduled date, except for time-sensitive situations such as urgent security updates.
>
> 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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/

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

```
srvctl 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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.
- 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 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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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. This can be either the database system identifier (SID), or a service name.

Connecting to a Database with SQL Developer

You can connect to a database with SQL Developer by using one of the following methods:

- Create a temporary SSH tunnel from your computer to the database. This method provides access only for the duration of the tunnel. (When you are done using the database, be sure to close the SSH tunnel by exiting the SSH session.)
- Open the port used as the Oracle SCAN listener by updating the security list used for the cloud VM cluster resource in the Exadata Cloud Service instance. The default SCAN listener port is 1521. This method provides more durable access to the database. For more information, see Updating the Security List.

After you've created an SSH tunnel or opened the SCAN listener port as described above, you can connect to an 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 VM cluster's private SCAN IP addresses, and then the application can use a connection string that includes the external SCAN name. In the

Copyright © 2023, 2026, Oracle and/or its affiliates.

following `tnsnames.ora` example, `extscanname.example.com` is defined in the on-premises DNS server.

```
testdb =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = <extscanname.example.com>)(PORT =
1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

To connect using public IP addresses

You can use the node's public IP address to connect to the database if the client and database are in different VCNs, or if the database is on a VCN that has an internet gateway.

However, there are important implications to consider:

- When the client uses the public IP address, the client bypasses the SCAN listener and reaches the node listener, so server side load balancing is not available.
- When the client uses the public IP address, it cannot take advantage of the VIP failover feature. If a node becomes unavailable, new connection attempts to the node will hang until a TCP/IP timeout occurs. You can set client side sqlnet parameters to limit the TCP/IP timeout.

The following `tnsnames.ora` example shows a connection string that includes the CONNECT_TIMEOUT parameter to avoid TCP/IP timeouts.

```
test=
  (DESCRIPTION =
    (CONNECT_TIMEOUT=60)
    (ADDRESS_LIST=
      (ADDRESS = (PROTOCOL = TCP)(HOST = <publicIP1>)(PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP)(HOST = <publicIP2>)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

Connecting to a Database Using SCAN

To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches.

- Connecting to a Database Using a Connect Descriptor that References All of the SCAN VIPs
  You can set up a connect descriptor for Oracle Exadata Database Service on Exascale Infrastructure System using multiple SCAN listeners.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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)))
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

```
exa1scan.example.com:1521/S12C1
```

Or

```
exa1scan.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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
```

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 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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. This can be 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.

To get the IP addresses of private client and backup subnets

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

```
exa1scan.example.com:1521/S12C1
```

Or

```
exa1scan.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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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
- 220 GB (for 26ai) and 260 GB (for 19c) 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 restarting the VM.
- You can scale memory in 2.75 GB increments per total ECPU.
- You can perform hot additions or removals of VMs
- You can scale VM storage for `/u02`.
- 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: 1100 GB
- 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 Oracle Advanced Cluster File System (ACFS). The remaining space is available for the first database, as shown in the following table:

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 220 GB (for 26ai) and 260 GB (for 19c). 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) for 26ai ExaDB-XS VM Cluster	Minimum Total Capacity (GB) for 19c ExaDB-XS VM Cluster
`/`	15	15
`/` (mirror)	15	15
`/home`	4	4
`/var`	5	5
`/var` (mirror)	5	5
`/var/log`	18	18
`/var/log/audit`	2	2
`/tmp`	10	10
`/u01`	42	82
`/u02`	62	62
Internal file systems	42	42
Total File system size per VM (GB)	**220**	**260**

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 yourOracle 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.
- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5

How-to Guides

A collection of tasks and procedures for managing Oracle Exadata Database Service on Exascale 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 Oracle 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
- Flexible Software Release Selection for Exadata Database
  Oracle applies release selection controls to Database and Grid Infrastructure release updates that you can select for lifecycle operations.
- 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
- Interim Software Updates
  For authorized environments, learn how to download interim 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 Oracle 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.
- Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure
  Exadata Database Service on Exascale Infrastructure@Azure enables you to store your database's transparent data encryption (TDE) keys, also known as master encryption keys (MEKs) in either a file-based Oracle wallet or in the OCI Vault.
- Google Cloud Key Management Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud
  Exadata Database Service on Exascale Infrastructure@Google Cloud now supports integration with Google Cloud Platform's Key Management Service (KMS).
- AWS Key Management Service Integration for Exadata Database Service on Exascale Infrastructure@AWS
  Exadata Database Service on Exascale Infrastructure@AWS supports integration with AWS Key Management Service (KMS). This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using AWS customer managed keys.
- Managed MCP Servers for OCI Database Services
  Managed Model Context Protocol (MCP) servers enable MCP-compatible AI agents and applications to access supported information about OCI Database resources.

Managed MCP Servers for OCI Database Services

Managed Model Context Protocol (MCP) servers enable MCP-compatible AI agents and applications to access supported information about OCI Database resources.

OCI Database Services supports managed MCP servers for Base Database Service, Oracle Exadata Database Service on Cloud@Customer, Oracle Exadata Database Service on Dedicated Infrastructure, and Oracle Exadata Database Service on Exascale Infrastructure. A managed MCP server provides a standardized MCP interface to OCI Database control plane APIs. MCP-compatible clients can discover available operations as MCP tools and use these tools to retrieve information about OCI Database resources.

How Managed MCP Servers Work

A managed MCP server provides an interface between an MCP-compatible client and supported OCI Database control plane APIs.

When you create an MCP server, you associate the server with an identity domain and select the applicable database service and resource categories. OCI generates an MCP endpoint and MCP client configuration for the server. You can use this information to configure a compatible MCP client.

Copyright © 2023, 2026, Oracle and/or its affiliates.

When you submit a natural-language request through an MCP-compatible AI agent or client, the client connects to the managed MCP server by using MCP over Streamable HTTP. The managed MCP server authenticates the client and authorizes access based on OCI Identity and Access Management (IAM) policies. If the request is authorized, the MCP server invokes the applicable supported OCI Database operation and returns information about the requested database resources to the client.

For example, you can ask an MCP-compatible AI agent to retrieve information about databases in a compartment. The agent can identify an appropriate MCP tool, invoke the tool, and use the information returned by OCI Database Services to respond to the request.

Using an MCP server doesn't change the permissions required to access OCI Database resources. The authenticated identity can access only the resources and operations permitted by the applicable IAM policies.

MCP Servers

An MCP server is an OCI resource that provides a remote endpoint for MCP-compatible clients.

You can create MCP servers for the following OCI Database services:

- Base Database Service
- Oracle Exadata Database Service on Cloud@Customer
- Oracle Exadata Database Service on Dedicated Infrastructure
- Oracle Exadata Database Service on Exascale Infrastructure

You can create multiple MCP servers in a tenancy for different database services.

An MCP server configuration includes information such as the associated identity domain, database service, resource categories, and tags.

OCI provides a unique MCP endpoint for each MCP server. OCI also provides MCP JSON configuration that you can use to configure compatible MCP clients.

You can enable or disable an MCP server. Disabling an MCP server prevents clients from accessing the server without deleting its configuration. When you no longer require an MCP server, you can delete it.

MCP Tools

MCP tools are supported OCI Database operations that an MCP-compatible client can discover and invoke through an MCP server.

The MCP server makes supported read-only control plane operations available as MCP tools. An AI agent can determine which available tool is appropriate for a request and invoke that tool with the required parameters.

The tools available through an MCP server depend on the database service and resource categories configured for the server.

Supported resource categories can include:

- Database infrastructure
- Virtual machines
- Databases Backups
- Data Guard
- Related database operational metadata

Copyright © 2023, 2026, Oracle and/or its affiliates.

The availability of an MCP tool doesn't grant permission to use the corresponding OCI Database operation. OCI IAM policies determine whether the authenticated identity can access the requested resource.

MCP Endpoint and Client Configuration

Each MCP server has an MCP endpoint that compatible clients use to connect to the server by using Streamable HTTP.

After creating an MCP server, you can obtain the generated endpoint from the MCP server details. You can also obtain the MCP JSON configuration required to configure a compatible MCP client.

After the client is configured and authenticated, it can connect to the endpoint and discover the MCP tools available through the server.

The way that you add an MCP server configuration to a client depends on the MCP client that you use.

Authentication

MCP servers use OAuth-based authentication through OCI Identity. When you create an MCP server, you associate it with an identity domain.

To access an MCP server, the MCP client authenticates with the associated identity domain and obtains an OAuth access token. The client uses the access token when connecting to the MCP server. The MCP server validates the token before allowing the client to access the supported MCP tools.

Authorization

OCI IAM controls access to OCI Database resources accessed through an MCP server.

IAM policies can control access based on users, groups, tenancies, compartments, and other supported OCI IAM constructs. The permissions of the authenticated identity determine which resources the identity can access through the MCP tools.

An MCP server doesn't provide additional permissions or bypass existing IAM policies. For example, if an identity has permission to view databases in one compartment but not another, the identity can retrieve information only for databases that the identity is authorized to view. Connecting through an MCP server doesn't extend that access to databases in other compartments.

Access through an MCP server is therefore subject to both the MCP server configuration and the permissions assigned to the authenticated identity.

Read-Only Operations

Managed MCP servers support approved read-only OCI Database control plane operations.

MCP-compatible AI agents can use these operations to retrieve supported information about database resources and their operational state. Operations that create, update, or delete OCI Database resources aren't available through the MCP server.

For example, an authorized agent can retrieve information about a database, database infrastructure, backup, or Data Guard configuration when a corresponding read-only MCP tool is available. The agent can't use the MCP server to modify those resources.

The read-only restriction applies to operations exposed through the managed MCP server and doesn't change the operations available through other OCI interfaces.

Using Multiple MCP Servers

Copyright © 2023, 2026, Oracle and/or its affiliates.

You can create multiple MCP servers to support different access and configuration requirements.

For example, you can configure separate MCP servers for different database services or expose different resource categories through different servers. Each MCP server has its own configuration and endpoint.

The permissions of the authenticated identity continue to apply independently of the MCP server that the client uses.

MCP Server States

An MCP server can be enabled, disabled, or deleted.

- When an MCP server is enabled, clients can connect to the MCP endpoint and use the supported MCP tools, subject to authentication and authorization.
- When an MCP server is disabled, clients can't connect to the MCP endpoint or use the supported MCP tools. The MCP server configuration is retained, and you can enable the server again when access is required.
- When an MCP server is deleted, clients can no longer connect to the MCP endpoint or use the supported MCP tools, and the MCP server configuration is permanently removed.

Security Considerations

Consider the following when using managed MCP servers:

- Use IAM policies to grant only the permissions required by users and groups that access OCI Database resources through MCP.
- Use separate MCP servers when different services or organizational requirements require different configurations.
- Disable an MCP server when MCP access is temporarily not required.
- Delete MCP servers that are no longer required.
- Review the permissions assigned to identities that use MCP-compatible clients.

Managed MCP servers provide another interface to supported OCI Database control plane operations. Existing OCI authentication, authorization, tenancy, and compartment boundaries continue to apply when those operations are accessed through MCP.

- Create MCP Server
  Create an MCP server to provide MCP-compatible clients with access to supported OCI Database services.
- View the List of MCP Servers
  View the MCP servers that you have access to in a compartment.
- View the Details of an MCP Server
  View the configuration, endpoint, and other details of an MCP server.
- Enable or Disable an MCP Server
  Enable or disable an MCP server to control client access. Disabling an MCP server prevents clients from accessing the server without deleting its configuration.
- Move an MCP Server to Another Compartment
  Move an MCP server from its current compartment to another compartment.
- Delete an MCP Server
  Move an MCP server from its current compartment to another compartment.
- Edit the Details of an MCP Server
  Edit an MCP server to update its name or associated database services.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Create MCP Server

Create an MCP server to provide MCP-compatible clients with access to supported OCI Database services.

By default, the MCP server is configured for all supported database services. You can limit the server to specific database services when you create it.

1. Open the navigation menu. Under **Oracle AI Database**, click any one of the following services, and then click **AI Capabilities**:
   - Base Database Service
   - Oracle Exadata Database Service on Cloud@Customer
   - Oracle Exadata Database Service on Dedicated Infrastructure
   - Oracle Exadata Database Service on Exascale Infrastructure
2. Click **View MCP Servers**.
   (or)

   Click the **MCP Servers** tab.
3. Click **Create MCP server**.
4. Specify the following information:
   - **Name**: Enter a name for the MCP server.
   - **Compartment**: Select the compartment in which to create the MCP server.
   - **Compartment in Domain**: Select the compartment containing the identity domain.
   - **Domain**: Select the identity domain to associate with the MCP server.
   - **Services**: Select the database services that the MCP server can access. All supported services are selected by default. To limit the service scope, select one or more of the following:
     - Base Database Service
     - Oracle Exadata Database Service on Cloud@Customer
     - Oracle Exadata Database Service on Dedicated Infrastructure
     - Oracle Exadata Database Service on Exascale Infrastructure
   - **Tags**: Optionally, add tags to the MCP server.
5. Click **Create**.

After the MCP server is created, it appears on the MCP Servers page.

View the List of MCP Servers

View the MCP servers that you have access to in a compartment.

1. Open the navigation menu. Under **Oracle AI Database**, click any one of the following services, and then click **AI Capabilities**:
   - Base Database Service
   - Oracle Exadata Database Service on Cloud@Customer
   - Oracle Exadata Database Service on Dedicated Infrastructure

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Oracle Exadata Database Service on Exascale Infrastructure

2. Click **View MCP Servers**.
   (or)

   Click the **MCP Servers** tab.

The MCP Servers page displays the MCP servers that you have access to, with the following information:

- **Name**: Name of the MCP server.
- **State**: Current state of the MCP server, either Active or InActive.
- **Services**: Database services associated with the MCP server.
- **Created**: Date and time when the MCP server was created.

View the Details of an MCP Server

View the configuration, endpoint, and other details of an MCP server.

1. Open the navigation menu. Under **Oracle AI Database**, click any one of the following services, and then click **AI Capabilities**:
   - Base Database Service
   - Oracle Exadata Database Service on Cloud@Customer
   - Oracle Exadata Database Service on Dedicated Infrastructure
   - Oracle Exadata Database Service on Exascale Infrastructure
2. Click **View MCP Servers**.
   (or)

   Click the **MCP Servers** tab.
3. From the list of MCP servers, use one of the following methods to view the details of an MCP server:
   - Click the name of the MCP server.
   - Click the **Actions** menu (three dots) for the MCP server, and then select **View details**.

   The MCP server details page displays the endpoint and configuration details of the selected MCP server.
4. To copy the MCP endpoint, click **Copy endpoint**.
5. To view or manage the tags associated with the MCP server, click the **Tags** tab.

Enable or Disable an MCP Server

Enable or disable an MCP server to control client access. Disabling an MCP server prevents clients from accessing the server without deleting its configuration.

1. Open the navigation menu. Under **Oracle AI Database**, click any one of the following services, and then click **AI Capabilities**:
   - Base Database Service
   - Oracle Exadata Database Service on Cloud@Customer
   - Oracle Exadata Database Service on Dedicated Infrastructure
   - Oracle Exadata Database Service on Exascale Infrastructure

Copyright © 2023, 2026, Oracle and/or its affiliates.

2. Click the **MCP Servers** tab.
3. From the list of MCP servers, use one of the following methods:
   - Click the Actions menu (three dots) for the MCP server that you want to enable or disable, and then select **Enable** or **Disable**, as applicable.
   - Click the name of the MCP server to open its details page, and then click **Enable** or **Disable**, as applicable.
4. In the confirmation dialog box, click **Enable** or **Disable** to confirm the action.

The state of the MCP server changes to **Active** or **Inactive**, depending on the selected action.

MCP Server States

An MCP server can have one of the following states:

- **Active**: The MCP server is enabled. Authenticated and authorized MCP clients can connect to the server endpoint and use the supported MCP tools.
- **Inactive**: The MCP server is disabled. MCP clients can't access the server endpoint or use its MCP tools. The server configuration is retained and can be enabled again.

You can enable or disable an MCP server without changing its configuration. When you enable an inactive MCP server, its state changes to **Active**. When you disable an active MCP server, its state changes to **Inactive**.

Move an MCP Server to Another Compartment

Move an MCP server from its current compartment to another compartment.

1. Open the navigation menu. Under **Oracle AI Database**, click any one of the following services, and then click **AI Capabilities**:
   - Base Database Service
   - Oracle Exadata Database Service on Cloud@Customer
   - Oracle Exadata Database Service on Dedicated Infrastructure
   - Oracle Exadata Database Service on Exascale Infrastructure
2. Click the **MCP Servers** tab.
3. From the list of MCP servers, use one of the following methods:
   - Click the Actions menu (three dots) for the MCP server that you want to move, and then select **Move resource**.
   - Click the name of the MCP server to open its details page, click the Actions menu, and then select **Move resource.**.
4. On the Move resource page, select the destination compartment, and then click **Move resource**.

The MCP server is moved to the selected compartment.

Delete an MCP Server

Move an MCP server from its current compartment to another compartment.

1. Open the navigation menu. Under **Oracle AI Database**, click any one of the following services, and then click **AI Capabilities**:
   - Base Database Service

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Oracle Exadata Database Service on Cloud@Customer
- Oracle Exadata Database Service on Dedicated Infrastructure
- Oracle Exadata Database Service on Exascale Infrastructure

2. Click the **MCP Servers** tab.
3. From the list of MCP servers, use one of the following methods:
   - Click the Actions menu (three dots) for the MCP server that you want to delete, and then select **Delete**.
   - Click the name of the MCP server to open its details page, click the Actions menu, and then select **Delete**.
4. In the confirmation dialog box, click **Delete** to confirm the action.

The MCP server is deleted and is no longer available on the MCP Servers page. MCP clients can no longer access the server through its endpoint.

Edit the Details of an MCP Server

Edit an MCP server to update its name or associated database services.

1. Open the navigation menu. Under **Oracle AI Database**, click any one of the following services, and then click **AI Capabilities**:
   - Base Database Service
   - Oracle Exadata Database Service on Cloud@Customer
   - Oracle Exadata Database Service on Dedicated Infrastructure
   - Oracle Exadata Database Service on Exascale Infrastructure
2. Click the **MCP Servers** tab.
3. From the list of MCP servers, click the name of the MCP server to open its details page, click the **Actions** menu, and then select **Edit**.
4. On the Edit MCP server page, update the following information as needed:
   - **Name**: Edit the name of the MCP server.
   - **Services**: Select the database services that the MCP server can access.
5. Click **Update**.

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
- Automatic Registration of Databases with Oracle Data Safe
  You can now automatically register newly provisioned container databases (CDBs) with Oracle Data Safe.

About Oracle Data Safe

Copyright © 2023, 2026, Oracle and/or its affiliates.

Your corporate policy requires that you monitor your databases and retain audit records. Your developers are asking for copies of production data for that new application, and you're wondering what kinds of sensitive information it will contain. Meanwhile, you need to make sure that recent maintenance activities haven't left critical security configuration gaps on your production databases and that staff changes haven't left dormant user accounts on the databases. Oracle Data Safe assists you with these tasks and is included with your Exadata Database Service*.

Oracle Data Safe is a unified control center, that helps you to manage the day-to-day security and compliance requirements of Oracle Databases no matter if they are running in the Oracle Cloud Infrastructure, at Cloud@Customer, on-premises or in any other cloud.

Data Safe supports you to evaluate security controls, assess user security, monitor user activity, and address data security compliance requirements for your database by evaluating the sensitivity of your data as well as masking sensitive data for non-production databases.

Data Safe provides the following features:

- **Security Assessment**: Configuration errors and configuration drift are significant contributors to data breaches. Use security assessment to evaluate your database's configuration and compare it to Oracle and industry best practices. Security assessment reports on areas of risk and notifies you when configurations change.
- **User Assessment**: Many breaches start with a compromised user account. User Assessment helps you spot the riskiest database accounts - those accounts which, if compromised, could cause the most damage - and take proactive steps to secure them. User Assessment Baselines make it easy to know when new accounts are added, or an account's privileges are modified. Use OCI events to receive proactive notifications when a database deviates from its baseline.
- **Activity Auditing**: Understanding and reporting on user activity, data access, and changes to database structures supports regulatory compliance requirements and can aid in post-incident investigations. Activity auditing collects audit records from databases and helps you manage audit policies. Audit insights make it easy to identify inefficient audit policies, while alerts based on audit data proactively notify you of risky activity.
- **Sensitive Data Discovery**: Knowing what sensitive data is managed in your applications is critical for security and privacy. Data discovery scans your database for over 150 different types of sensitive data, helping you understand what types and how much sensitive data you are storing. Use these reports to formulate audit policies, develop data masking templates, and create effective access control policies.
- **Data Masking**: Minimizing the amount of sensitive data your organization maintains helps you meet compliance requirements and satisfy data privacy regulations. Data masking helps you remove risk from your non-production databases by replacing sensitive information with masked data. With reusable masking templates, over 50 included masking formats, and the ability to easily create custom formats for your organization's unique requirements, data masking can streamline your application development and testing operations.
- **SQL Firewall Management**: Protect against risks such as SQL injection attacks or compromised accounts. Oracle SQL Firewall is a new security capability built into the Oracle AI Database 26ai kernel and offers best-in-class protection against these risks. The SQL Firewall feature in Oracle Data Safe lets you centrally manage and monitor the SQL Firewall policies for your target databases. Data Safe lets you collect authorized SQL activities of a database user, generate and enable the policy with allowlists of approved SQL statements and database connection paths and provides a comprehensive view of any SQL Firewall violations across the fleet of your target databases.

Copyright © 2023, 2026, Oracle and/or its affiliates.

**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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 Oracle Database 23ai 26ai target database: Enable SQL Firewall On Your Target Database.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

After 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 Database 23ai 26ai.

Automatic Registration of Databases with Oracle Data Safe

You can now automatically register newly provisioned container databases (CDBs) with Oracle Data Safe.

If you have the required Data Safe privileges, registration is enabled by default when you create a new CDB. You can choose not to enable registration during database provisioning.

Any Oracle database registered with Data Safe for security management is referred to as target database. For more information, see:

- Target Database Registration
- Target Database Information Stored in Oracle Data Safe
- Manage Database Security with Oracle Data Safe

When Data Safe registration is enabled for a CDB, Oracle Database creates a registration policy for that CDB in Data Safe. Any pluggable databases (PDBs) created in the CDB automatically inherit the registration policy and are registered with Data Safe.

As part of the automatic registration process, a Data Safe managed service account is created or updated as needed, and the required privileges are granted by default for the following Data Safe features:

- Security Assessment
- User Assessment
- Data Discovery
- Activity Auditing

The privileges required for Data Masking and SQL Firewall are not granted automatically and must be configured in the Oracle Data Safe console.

After registration, Security Assessment and User Assessment run automatically and collect metadata about the database configuration, users and privileges, and database details. During registration, Data Safe stores connection information for the Data Safe service account. Data Safe automatically manages the service account credentials and periodically rotates the account password to help maintain security.

Table 5-1 Registration behavior matrix

Actions	Data Safe	CDB	Future PDBs
Create a CDB with Data Safe enabled	ON	Registered automatically	Automatically registered
Create a CDB with Data Safe disabled	OFF	Not registered	Not automatically registered

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-1 (Cont.) Registration behavior matrix

Actions	Data Safe	CDB	Future PDBs
Toggle Data Safe from OFF to ON after provisioning	ON	CDB and all existing PDBs registered at toggle time	Automatically registered
Toggle Data Safe from ON to OFF	OFF	Unchanged	Not automatically registered

For more information, see:

- Preregistration Tasks
- Register an Oracle Cloud Database
- Create IAM Policies for Oracle Data Safe Users
- Permissions to Register an Oracle Cloud Database with Oracle Data Safe
- Permissions for an Oracle Data Safe Private Endpoint
- Permissions for an Oracle Data Safe On-Premises Connector
- Permissions to Discover Sensitive Data
- Permission to Mask Sensitive Data
- Prerequisites for SQL Firewall
- Manage Database Security with Oracle Data Safe

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

  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
  ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Create a temporary SSH tunnel from your computer to the database. This method provides access only for the duration of the tunnel. (When you are done using the database, be sure to close the SSH tunnel by exiting the SSH session.)
- Open the port used as the Oracle SCAN listener by updating the security list used for the cloud VM cluster resource in the Exadata Cloud Service instance. The default SCAN listener port is 1521. This method provides more durable access to the database. For more information, see Updating the Security List.

After you've created an SSH tunnel or opened the SCAN listener port as described above, you can connect to an 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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. This can be 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.

To get the IP addresses of private client and backup subnets

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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)))
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

```
exa1scan.example.com:1521/S12C1
```

Or

```
exa1scan.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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
```

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 Exadata Cloud Infrastructure virtual machines (VMs) and databases on Oracle Exadata Database Service on Exascale Infrastructure (ExaDB-XS).

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle 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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
   The Exadata VM Clusters page is selected by default.
2. 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

3. Click the **Virtual Machines** tab to display all the virtual machines of the cloud service instance. Click the Actions menu (three dots) for a node and then click one of the following actions:
   - **Start:** Restarts a stopped virtual machine. After the virtual machine is restarted, the **Stop** action is enabled.
   - **Stop:** Shuts down the virtual machine. After the virtual machine is powered off, the **Start** action is enabled.
   - **Reboot:** Shuts down the virtual machine, and then restarts it.

> **Note**
>
> - For billing purposes, the **Stop** state has no effect on the resources you consume. Billing continues for virtual machines 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 start or restart a virtual machine, 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 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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster you want to manage:
   Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on 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 **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**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Management Tasks for the Oracle Cloud Infrastructure Platform

- To Request Service Limit Increase
  Learn how to request increase in service limits for Exadata Database Service on Exascale Infrastructure resources.
- To view a work request for your Oracle Exadata Database Service on Exascale Infrastructure resources
- To move an Oracle Exadata Database Service on Exascale Infrastructure resource to another VM cluster
  You can use the console to move resources to another VM cluster on Oracle Exadata Database Service on Exascale Infrastructure.
- To manage tags for your Oracle Exadata Database Service on Exascale Infrastructure resources
  You can use the console to view, edit, or add tags for a VM cluster on Oracle Exadata Database Service on Exascale Infrastructure.

To Request Service Limit Increase

Learn how to request increase in service limits for Exadata Database Service on Exascale Infrastructure resources.

1. Open the navigation menu and click **Governance & Administration**.
2. Click **Limits, Quotas and Usage**.
3. Select **Request a service limit increase**.
4. **General information**:
   a. Enter the request name and the reason.
   b. Add tags as applicable.
5. Click **Next**.
6. **Requested items:**
   a. Click **Add**.
   b. Select the service that requires the limit increase.
   c. Click **Next**.
   d. Select the limit name from the drop-down list.
   e. Enter the limit value.
   f. Click **Save**.

   You can repeat the process and add multiple services at a time.
7. Select the check mark beside the requested service and click **Next**.
8. Review the entered details.
9. Click **Submit**.
   After submitting the request, your service limit will be increased within 48 hours.

Related Topics

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

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

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The Exadata VM Clusters page is selected by default.
2. Find the Cloud Exadata infrastructure, cloud VM cluster, or database resource you're interested in, and click the name.
3. Click the **Work Requests** tab. The status of all work requests appears on the page.
4. To see the log messages, error messages, and resources that are associated with a specific work request, click the operation name.

   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 Oracle Exadata Database Service on Exascale Infrastructure resource to another VM cluster

You can use the console to move resources to another VM cluster on Oracle Exadata Database Service on Exascale Infrastructure.

> **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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The Exadata VM Clusters page is selected by default.
2. 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.
3. Click the **Actions** button, and select **Move resource**.
4. Select the new compartment.
5. Click **Move Resource**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Related Topics

- Details for the Database Service
- Security Zone Policies

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

You can use the console to view, edit, or add tags for a VM cluster on Oracle Exadata Database Service on Exascale Infrastructure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
   The **Exadata VM Clusters** page is selected by default.
2. Find the Cloud Exadata infrastructure, cloud VM cluster, or database resource you're interested in, and click the name.
3. Click the **Tags** tab.
4. On the Tags page, click **Add** to add additional tags, or select a tag name to manage that tag, and click the **Actions** button. Select the appropriate action for the tag that you want to manage.

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
  You can use the console to update the license type for a VM cluster on 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.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
   The Exadata VM Clusters page is selected by default.
2. 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.
3. Click **Scale VM Cluster** and then specify a new number of CPU cores.
4. Click **Scale**.

To change the license type of a cloud VM cluster

You can use the console to update the license type for a VM cluster on Oracle Exadata Database Service on Exascale Infrastructure.

Copyright © 2023, 2026, Oracle and/or its affiliates.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
   The Exadata VM Clusters page is selected by default.
2. 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.
3. Click the **Actions** button, and select **Update license type**.

   The dialog displays the options with your current license type selected.
4. Select the new license type.
5. 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

- CreateExascaleDbStorageVault
- GetExascaleDbStorageVault
- ListExascaleDbStorageVaults

Exadata VM cluster resource

- CreateExadbVmCluster
- GetExadbVmCluster
- ListExadbVmClusters

Databases

- GetDatabase
- ListDatabases

Database Versions

- ListDbVersions

Database Homes

- CreateDbHome
- GetDbHome
- ListDbHomes

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
- RemoveVirtualMachineFromExadbVmCluster
- UpdateExadbVmCluster

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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).
- Troubleshooting Virtual Machines Using Console Connections
  You can troubleshoot malfunctioning virtual machines using console connections. For example, a previously working Guest VM stops responding.

Related Topics

- Application Checklist for Continuous Service for MAA Solutions

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.
- To add security attribute to an Exadata VM Cluster
  To add a security attribute to Exadata VM Cluster, use this procedure
- To edit a security attribute
  To edit a security attribute of an Exadata VM Cluster, use this procedure.
- To remove a Security Attribute
  To remove a security attribute of an Exadata VM Cluster, use this procedure.
- To scale database storage capacity of a cluster database in the VM cluster
  This topic describes how to scale storage capacity of container databases in the VM cluster configured with Exascale Block Storage.
- To add database server or storage server capacity to a cloud VM cluster
  This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.
- To Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 another user which requires another SSH key.
- To Add or Delete SSH Keys After VM Cluster Creation
  You can use the console to add or delete SSH keys for an existing VM cluster on Oracle Exadata Database Service on Exascale Infrastructure.
- Using the Console to Stop, Start, or Reboot a Virtual Machine
  Use the console to stop, start, or reboot a virtual machine.
- To Check the Status of a VM Cluster Virtual Machine
  Review the health status of a VM cluster virtual machine.
- To Create a Snapshot of Virtual Machines
  Enable users to initiate self-service Virtual Machine (VM) file system snapshot.
- To Delete a Snapshot of Virtual Machines
  Use the console to delete the snapshot of a Virtual Machine (VM).
- To Mount a Snapshot of Virtual Machines
  Use the console to mount the snapshot of a Virtual Machine (VM) in read-only mode.
- To Unmount a Snapshot of Virtual Machines
  Use the console to unmount the snapshot of a Virtual Machine (VM).
- To Create VM File System Snapshots for VM Cluster
  Use the console to create snapshots of all the VMs in a VM Cluster.
- 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.
- 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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The Exadata VM Clusters page is selected by default.
2. Click **Create VM Cluster**.

   The **Create VM Cluster** page is displayed. Provide the required information to configure the VM cluster.
3. **Display name**: Enter a user-friendly display name for the VM cluster. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the VM cluster. Avoid entering confidential information.
4. **Compartment**: Select a compartment for the VM cluster resource.
5. **Provide the cluster name**: Select the name of the VM cluster.

Copyright © 2023, 2026, Oracle and/or its affiliates.

6. **Select an availability domain**: Select the availability domain from the displayed options available.
7. **Configure the VM cluster**: Provide the following information:
   In the **ECPU** section, 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. To reserve additional ECPUs, click **Reserve additional ECPU**.

     The **Total memory across VM Cluster (GB)** (read-only) 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.

     > **Note**
     >
     > For information about reserved and enabled cores, and an overview of the ExaDB-XS architecture, see "About Exadata Database Service on Exascale Infrastructure"
8. (Optional) If you select **Reserve additional ECPU**, then the Reserve Additional ECPU dialog displays.

   The number of VMs in the cluster is displayed. The current configuration of ECPUs enabled per VM and additional ECPUs reserved per VM.

   Total ECPUs per VM: Input a number to change the total number of ECPUs reserved per VM. The total number of ECPUs must be a number between 8 and 200, in multiples of 4 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 ECPUs.
   - **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)**: Automatically updates to show you the reserved memory per VM
   - **Total Memory across VM cluster (GB)**: Shows the total reserved memory across the VM cluster.
9. The **Storage** section contains the input fields **Exascale smart storage** (the default), and **Exascale block storage**, which is required for Oracle Database 19c deployments. Note the following:
   - If your VM cluster uses Exascale smart storage, then the following rules apply:
     - You can only have an Oracle Grid Infrastructure 26ai Grid home on this VM. You cannot place an Oracle Database 19c Grid home on the VM.
     - You can only provision Oracle AI Database 26ai databases on this VM cluster configured with Exascale smart storage. All database files will be stored directly on

Copyright © 2023, 2026, Oracle and/or its affiliates.

the Exascale storage through the Exascale storage vault, eliminating the need for an Automatic Storage Management (ASM).

- If your VM cluster uses Exascale block storage compatible with Oracle Database 19c, then the following rules apply:
  - You can provision either an Oracle Grid Infrastructure 19c Grid home (the default option) or an Oracle Grid Infrastructure 26ai Grid home.
  - You can provision only Oracle Database 19c Oracle homes on this VM cluster configured with Exascale block storage. All database files will be stored on the database-specific `+DATA`, `+RECO`, and `+LOG` ASM disk groups created on Exascale Block volumes.

10. Under the **VM file system storage** section, provide the following information:

    VM File system storage capacity (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 file system storage capacity across VM Cluster (Read-only)** automatically updates to show the total memory allocated across the VM cluster.

11. In the **Database storage** section, provide the following information:

    Exascale Database Storage Vault: Select either **Create new Storage Vault** or **Select existing Storage Vault**. If you select an existing vault, then select the vault in the compartment. Click **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. **Storage Capacity for Databases (GB)**: Specify storage capacity for the databases. You can either specify storage limits manually, or set an auto scaling limit. With auto scale, the system will scale storage capacity by 10 percent up to a user-defined limit. If you specify storage space manually, then the space you reserve must be between 300 to 100,000 GB.
    - **Storage Capacity for Databases (GB)**: Specify storage capacity for the databases between 300 to 100,000.
    - **Storage auto scaling**: For Exascale smart storage, this option is enabled by default.
    - **Auto scaling limit (GB)**: Specify the maximum limit for auto scale.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> – The storage auto scaling feature is not available with Oracle Database 19c.
>
> – Auto scale is primarily intended to accommodate the ongoing, organic growth of databases. For large, sudden data loads or rapid data growth that can exceed 5% of the allocated vault size within a short time period, it is recommended to manually scale storage in advance to avoid potential out-of-space errors.
>
> – Tablespaces with auto-extend enabled can cause out-of-space errors if the auto-extend increment size is greater than 5% of the allocated vault storage. This could be caused if a single, large tablespace auto-extension suddenly exceeds the available space in the vault before crossing the auto scale threshold. Or, it could also be caused by multiple tablespaces extensions occurring concurrently, resulting in the same effect. It is recommended that you manage your tablespace auto-extend increment settings to avoid having more than 5% of the allocated vault storage added at once when using auto scale.

- (Optional) Click **Smart flash options** to configure some of your storage as smart flash cache:

  Smart flash cache to be added (GB).

  – **Add smart flash cache 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. The minimum value is 34, and the maximum value is 300.

  – **Smart flash cache to be added (GB)** The amount of smart flash cache in GB that will be added is specified in the read-only field. The minimum value is 100

13. **Configure Grid Infrastructure**: Select the Grid Infrastructure release from the list of available releases. **Grid Infrastructure version**: Select the release update version from the list of available versions.
14. **Add SSH key**: Add the public key portion of each key pair that you want to use for SSH access to the VM cluster:

    - **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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

15. **Network settings**: You can have Oracle provision networks for you by selecting **Create default VCN**, or you can configure networks manually.

    - **Create Default VCN**

      When you select this option, a new VCN is created for you with default settings. The default settings use an overall network range of 10.2.0.0/24, with 10.2.0.0/26 for the Client subnet and 10.2.0.64/26 for the backup subnet. These default settings will auto-populate for this new VCN and automatically be used for your VM Cluster. You provide only the display name for this VCN. If you do not want to use the default network settings, then do not alter the generated settings. Instead, follow the instructions in the following option, "**Create VCN Manually**".

      – **Display name**: Provide a display name for the network. You must provide a name.

      – **Compartment**: By default, Virtual Cloud Networks in the compartment that you have selected for the VM cluster resource are displayed. Select a Virtual Cloud Network in that compartment from the drop-down list. To select a Virtual Cloud Network from a different compartment (a non-default choice), select **Change compartments**.

      – **CIDR (Read-only)**: The Classless Inter-Domain Routing (CIDR) range of IP addresses is displayed. CIDR addresses are allocated during the configuration of your tenancy. The CIDR notation includes the IP address and a suffix that indicates how many bits of the address are used for the network prefix. This method of allocating IP addresses and routing internet traffic more efficiently than traditional methods. For more information, see:

        CIDR Blocks and OCI Functions

      After you finish selecting the VCN configuration options, click **Create**

      The Create Default VCN dialog box opens, and displays the status of your network configuration. If you want to use the automatically generated VCN, do not alter the generated settings. When the network configuration is complete, click **Close**.

    - **Create VCN Manually**

      If you do not click **Create default VCN**, then select your configuration information in the following input fields:

      – **Virtual cloud network**: Select the virtual cloud network (VCN) for the compartment in which you want to create the VM cluster. Click **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 **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 **Compartment** to select a subnet in a different compartment, if applicable.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 select this option to use 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.
  * **Network security groups**: From the list of available NSGs, select an NSG to use with the network. Click **Another network security group** if you require more than one NSG. If you're not sure, contact your network administrator.

- **Hostname prefix**: Provide your choice of hostname for the Exadata VM cluster. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for an Exadata VM cluster is 12.

  > **Caution**
  >
  > The hostname must be unique within the subnet. If it is not unique, 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 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.

  > **Note**
  >
  > Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

- **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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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. **Diagnostics collection**: By default, diagnostic collection is enabled. Enabling diagnostics collection and notifications allows you and Oracle Cloud operations to identify, investigate, track, and resolve guest VM issues quickly and effectively. You must subscribe to events to receive notification.
18. (Optional) Click **Show Advanced Options** to specify advanced options for the VM cluster. These options are located under three tabbed sections:
    - **Management**: This option provides timezone options. The default time zone for the VM cluster is UTC, but you can specify either the time zone detected by your browser client, or select a different time zone. The time zone options are those supported in both the **Java.util.TimeZone** class and the Oracle Linux operating system.

      > **Note**
      >
      > If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting the **Select another time zone**, option, then selecting "Miscellaneous" in the **Region or country** list and searching the additional **Time zone** selections.

    - **Network**: This option enables you to specify the single client access network (SCAN) port. You can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. The default is 1521.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> Manually changing the SCAN listener port of a VM cluster after provisioning using the backend software is not supported. This change can cause Data Guard provisioning to fail.

- **Zero Trust Packet Routing (ZPR)**: This option is located in the Security attributes tab. Select a namespace, and provide the key and value for the security attribute. To complete this step during configuration, you must already have set up security attributes with Oracle Cloud Infrastructure Zero Trust Packet Routing. You can also add security attributes after configuration, and add them later. For more information about adding Oracle Exadata Database Service on Dedicated Infrastructure specific policies, see Policy Template Builder.
- **Security**: This option helps you control access for your resources. You can select the Namespace from the list, its key, and the value. Click **Add security attribute** if you require more than one security attribute.
- **Tags**: Tags can assist you with managing resources. You can select a tag namespace, provide a tag key, and provide a tag value. 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.

19. Click **Create**.
20. (Optional) **Provide a contact for your VM Cluster**. Oracle 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.

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.

To add security attribute to an Exadata VM Cluster

To add a security attribute to Exadata VM Cluster, use this procedure

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

3. In the list of cloud VM clusters, click the name of the cluster to which you want to add security attributes.
4. On the VM Cluster Details page, click **More actions**, and then select **Add security attributes**.

   or

   On the VM Cluster Details page, click the **Security attributes** tab.
5. Click **Add security attributes**.
6. Select the **Namespace** in which the required security attribute is available.
7. Select the **Key** and **Value** of the required security attribute.
8. To use additional security attributes, click **Add security attribute**.

   > **Note**
   >
   > A maximum of 3 security attributes can be specified for an Exadata VM Cluster.
9. Click **Add security attributes**.

To edit a security attribute

To edit a security attribute of an Exadata VM Cluster, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. In the list of cloud VM clusters, click the name of the cluster to which you want to edit a security attributes.
3. Click the **Security attributes** tab.

   A list of existing security attributes are displayed.
4. Click the icon beside the name of the security attribute to view its details.
5. Choose the required value of the security attribute.
6. Click **Save**.

To remove a Security Attribute

To remove a security attribute of an Exadata VM Cluster, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. In the list of cloud VM clusters, click the name of the cluster for which you want to edit a security attribute.
3. Click the **Security attributes** tab.

   A list of existing security attributes are displayed.
4. Click the icon beside the name of the security attribute to view its details.
5. Click **Remove security attribute**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

To scale database storage capacity of a cluster database in the VM cluster

This topic describes how to scale storage capacity of container databases in the VM cluster configured with Exascale Block Storage.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the name of the VM cluster where the specific container database is located.
3. Click the **Container databases** tab.
4. On the Container Databases page, select the container database where you want to add storage.
5. Click the **Database information** tab.
6. On the **Database Information** pane, click the **Actions** drop-down list and select **Scale database storage**.

> **Note**
>
> This section applies only to Oracle Database 19c container databases (CDBs) running on VM Clusters configured with Exascale Block Storage. There is no need to allocate any storage for Oracle AI Database 26ai CDBs running on VM Clusters configured with Exascale Smart Storage.

7. In the **Scale database storage** dialog, provide the following:
   - **Storage vault**: The name of the storage vault associated with the VM cluster is displayed
   - **Available vault storage capacity (GB)**: (read-only) The available GB of storage capacity in the vault is displayed

     > **Note**
     >
     > If you don't have enough available vault storage capacity for database storage scale up operation, you need to first scale up the storage capacity of the storage vault as described in "Scaling Exascale Database Storage Vault" in Manage Exascale Database Vaults on Oracle Exadata Database Service on Exascale Infrastructure

   - **Data storage (GB)** (`DATA`): The existing data storage size is displayed. You can scale up or scale down the storage that you want to allocate to the database.

     > **Note**
     >
     > The minimum size for Data storage is 56 GB. It cannot be less than the usable space of the primary database.

   - **Recovery area storage size (GB)** The existing recovery area storage size is displayed. You can scale up or scale down the storage as needed.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> The minimum size for the Recovery area is 44 GB for a 2-node cluster.

- **Log storage (GB)** (read-only) The log storage size is calculated automatically based on the number of nodes in the VM cluster. For a 2-node VM cluster, the Log storage (LOG) size is 33 GB.
- **Total Allocated storage capacity (GB)** (read-only) This field displays the total allocated storage for data files, the recovery area storage, and log storage.

8. Click **Scale**.

The Scale database storage dialog closes. On the container database page, the **Updating** status indicator shows that the storage capacity database is being scaled to the new allocation.

When that process is complete, the **Available** status indicator is displayed.

To add database server or storage server capacity to a cloud VM cluster

This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.

If you have added additional database (compute) or storage servers to the service instance, then you must add the additional capacity to the cloud VM cluster to utilize the additional resources.

1. Navigate to **Exadata VM Clusters**
2. In the list of cloud VM clusters, click the name of the cluster to which you want to add capacity.
3. On the VM Cluster Details page, click **Scale VM Cluster**.
4. If you have additional capacity available as a result of scaling the cloud Exadata infrastructure resource, a banner at the top of the **Scale VM Cluster** panel provides a message telling you the type and amount of additional capacity available to the VM cluster. Check the **Add Capacity** box.
5. Select either the **Add Database Server** or the **Add Storage** radio button, depending on which type of capacity you want to add to the cloud VM cluster.
6. Click **Update**. The cloud VM cluster goes into the Updating state. When the capacity has been successfully added, the cluster returns to the Available state.

> **Note**
>
> If you have added additional database servers to the cluster, then you can allocate the new storage after the cluster is in the **Available** state by clicking the **Scale VM Cluster** button again.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

configuration to all the resources such as DB home, Database, and so on under the VM cluster.

> **Note**
>
> - You are opting in with the understanding that the list of events, metrics, and log files collected can change in the future. You can opt-out of this feature at any time.
> - Oracle may add more metrics in the future, but if you have already chosen to collect metrics, you need not update your opt-in value. It will remain enabled/disabled based on your current preference.
> - If you have previously opted in for incident log and trace file collection and decide to opt out when Oracle Cloud operations run a log collection job, then the job will run its course and will not cancel. Future log collections won't happen until you opt-in again to the incident logs and trace file collection option.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI 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 **Exadata 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 resulting Update license type window, choose one of the following license types and then click **Save**.
   - **Bring Your Own License (BYOL):** Select this option if your organization already owns Oracle Database software licenses that you want to use on the VM cluster.
   - **License Included:** Select this option to subscribe to Oracle Database software licenses as part of 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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 **Scale VM Resources** window opens, and displays the current configuration of your VM cluster.
3. In the Configure section, make selections as needed.
   - **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. The **Total ECPUs enabled across VM Cluster** field automatically updates as you make selections.

     > **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.

   - **Additional ECPUs**: Expand these fields as needed.

     The read-only fields **ECPUs additional reserved per VM** and **Total additional ECPUs reserved across VM Cluster**: display the reserved ECPUs for each VM, and for the VM cluster. The number of additional ECPUs in these fields will be automatically updated as you add or remove 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.
   - **Total memory across VM Cluster (GB)**: This 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**: Expand these fields as needed.
   - **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`.
   - **Total file system storage capacity across VM Cluster**: This read-only field automatically updates to display the total file system storage capacity.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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 another user which requires another SSH key.

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

   The Exadata VM Clusters page is selected by default.
2. In the list of VM clusters, find the cluster you want to manage and click its highlighted name to view the details page for the cluster.
3. Click the **Actions** button and select **Add SSH keys** from the list.
4. 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.
5. Click **Save**.

To Add or Delete SSH Keys After VM Cluster Creation

You can use the console to add or delete SSH keys for an existing VM cluster on Oracle Exadata Database Service on Exascale Infrastructure.

SSH public keys provide secure administrative access to the virtual machines in a VM cluster. The VM cluster stores the public key. You use the corresponding private key when you connect.

You can add SSH public keys when you create a VM cluster or after the VM cluster is created. For information about adding SSH keys during VM cluster creation, see To create a cloud VM cluster.

> **Note**
>
> Before you add SSH keys, note the following:
>
> - If you select Generate SSH key pair, save the private key before you close the dialog. You cannot download the private key again later.
> - For pasted keys, enter each key on a single line. The combined length of all pasted keys cannot exceed 10,000 characters.

Copyright © 2023, 2026, Oracle and/or its affiliates.

To add SSH keys:

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM Clusters page is selected by default.
2. Click the name of the VM cluster to which you want to add SSH keys.
3. On the VM cluster details page, click the **Actions** menu, and then select **Add SSH Keys**.
4. In the **Add SSH Keys** dialog, select one of these options:
   a. **Generate SSH key pair**: Create a public and private key pair. Click **Save Private Key** to download the private key. Optionally, click **Save Public Key** to download the public key.
   b. **Upload SSH key files**: Upload one or more .pub files.
   c. **Paste SSH keys**: Paste one or more public keys. To add more keys, click **Another SSH Key**.
5. Click **Save**.

> **Note**
>
> Before you delete SSH keys, note the following:
>
> - At least one SSH key must remain associated with the VM cluster.
> - If the VM cluster has only one SSH key, you cannot delete that key.
> - If the VM cluster has multiple SSH keys, you can delete one or more keys as long as at least one SSH key remains.
> - Deleting an SSH key prevents new SSH connections that use the corresponding private key. Existing SSH sessions that use the deleted key are not affected.
> - Before you delete a key, check whether administrators, automation tools, scripts, or other operational processes use it.

To delete SSH keys:

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM Clusters page is selected by default.
2. Click the name of the VM cluster from which you want to delete SSH keys.
3. On the **VM Cluster Details** page, click the **Actions** menu, and then select **Delete SSH Keys**.
4. Select one or more keys that you want to delete.
5. Click **Delete**.

Related Topics

- Managing Key Pairs on Linux Instances

Using the Console to Stop, Start, or Reboot a Virtual Machine

Use the console to stop, start, or reboot a virtual machine.

Copyright © 2023, 2026, Oracle and/or its affiliates.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle 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 **Exadata 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 virtual machines, click the Actions icon (three dots) for a virtual machine you want to manage, and then choose one of the following actions:
   a. **Start**: Restarts a stopped virtual machine. After the virtual machine is restarted, the Stop action is enabled.
   b. **Stop**: When a virtual machine (VM) is unresponsive, you can stop it using either a graceful or force option.
      i. Choose one of the following options:
         - **Stop** to allow the VM to shut down gracefully.
         - **Force Stop** to immediately stop the VM without waiting for it to shut down cleanly.
      ii. When prompted, type the name of the VM in the confirmation text box to verify that you are performing the action on the intended VM.
      iii. Confirm the action to stop the VM.

      > **Note**
      >
      > If you select Stop and the VM remains unresponsive for 30 minutes, the system automatically initiates a Force Stop.

   c. **Reboot**: When a virtual machine (VM) is unresponsive, you can reboot it using either a graceful or force option.
      i. Choose one of the following options:
         - **Reboot** to allow the VM to shut down gracefully.
         - **Force Reboot** to immediately reboot the VM without waiting for it to shut down cleanly.
         - When prompted, type the name of the VM in the confirmation text box to verify that you are performing the action on the intended VM.
         - Confirm the action to proceed with the reboot.

      > **Note**
      >
      > If you select Reboot and the VM remains unresponsive for 30 minutes, the system automatically initiates a Force Reboot.

Limitations

Copyright © 2023, 2026, Oracle and/or its affiliates.

When a VM is already in an updating state (for example, stopping or rebooting), the following behaviors apply:

- Graceful actions are disabled because the VM cannot respond while another workflow is in progress.
- Only Force Reboot or Force Stop options are available to override the existing workflow.
- Patching operations and graceful shutdowns cannot be overridden.
- A graceful action in progress cannot be overridden by another graceful action.

To Check the Status of a VM Cluster Virtual Machine

Review the health status of a VM cluster virtual machine.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on 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. Click the **Virtual Machines** tab.

   The list of virtual machines displays the name, state, and client IP address for each virtual machine in the VM cluster. In the list, locate the virtual machine you want to manage and check its state:

   - **Available**: The virtual machine is running and operational.
   - **Starting**: The virtual machine is starting as a result of a start or reboot action initiated from the Console or API.
   - **Stopping**: The virtual machine is stopping as a result of a stop or reboot action initiated from the Console or API.
   - **Stopped**: The virtual machine is not running.
   - **Failed**: An error condition prevents the virtual machine from operating.

To Create a Snapshot of Virtual Machines

Enable users to initiate self-service Virtual Machine (VM) file system snapshot.

You can mount a VM snapshot in read-only mode to access and copy files exactly as they existed at the time the snapshot was taken. This is useful when you need to restore specific configurations of your VM to a previous state—for example, recovering earlier SSH keys, and so on.

1. Navigate to the VM Cluster main page.
2. In the list of VM clusters, find the cluster you want to manage and click its highlighted name to view the details page for the cluster.
3. Click the **Virtual Machines** tab.
4. From the Virtual Machines table, click the name of the virtual machine for which you want to create a snapshot.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5. Click the **VM File System Snapshots** tab.
6. Click **Create Snapshot**.
7. Enter a name for your snapshot, and click **Create**.
   When the snapshot creation is complete, the state changes from Creating to Available.

To Delete a Snapshot of Virtual Machines

Use the console to delete the snapshot of a Virtual Machine (VM).

1. Navigate to the VM Cluster main page.
2. In the list of VM clusters, find the cluster you want to manage and click its highlighted name to view the details page for the cluster.
3. Click the **Virtual Machines** tab.
4. From the Virtual Machines table, click the name of the virtual machine for which you want to delete a snapshot.
5. Click the **VM File System Snapshots** tab.
6. In the row corresponding to the snapshot that you want to delete, click the Action menu.
7. Click **Delete**.
   Read the message and acknowledge that you want to delete the VM snapshot.

To Mount a Snapshot of Virtual Machines

Use the console to mount the snapshot of a Virtual Machine (VM) in read-only mode.

1. Navigate to the VM Cluster main page.
2. In the list of VM clusters, find the cluster you want to manage and click its highlighted name to view the details page for the cluster.
3. Click the **Virtual Machines** tab.
4. Click the **VM File System Snapshots** tab.
5. Select the snapshot to mount and click the Action menu.
6. Click **Mount**.
7. Read the message and acknowledge that you want to mount the snapshot.

When the snapshot mount is complete, the mount point is displayed in the Resources list of the VM filesystem snapshots. You can then use this mount point to copy the files from selected VM snapshot.

> **Note**
>
> Snapshots are mounted as read-only filesystems.

To Unmount a Snapshot of Virtual Machines

Use the console to unmount the snapshot of a Virtual Machine (VM).

1. Navigate to the VM Cluster main page.
2. In the list of VM clusters, find the cluster you want to manage and click its highlighted name to view the details page for the cluster.

Copyright © 2023, 2026, Oracle and/or its affiliates.

3. Click the **Virtual Machines** tab.
4. From the Virtual Machines table, click the name of the virtual machine for which you want to unmount a snapshot.
5. Click the **VM File System Snapshots** tab.
6. In the row corresponding to the snapshot that you want to unmount, click the Action menu.
7. Click **Unmount**.
8. Read the message and acknowledge that you want to unmount the snapshot.

To Create VM File System Snapshots for VM Cluster

Use the console to create snapshots of all the VMs in a VM Cluster.

1. Navigate to the VM Cluster main page.
2. In the list of VM clusters, find the cluster you want to manage and click its highlighted name to view the details page for the cluster.
3. Click the **Virtual Machines** tab to view the list of Virtual Machines present in the cluster.
4. Click **Create VM file system snapshot for entire VM Cluster**.
5. Enter a name for your snapshot, and click **Create**.
   When the snapshot creation is complete, the state of the Virtual Machines changes from Updating to Available.

> **Note**
>
> - These snapshots are not available at the VM Cluster level, but are created synchronously for all individual Virtual Machines in the VM cluster.
> - To restore the entire cluster, you must restore each individual Virtual Machine located in the cluster. For more information, see *To Mount Snapshots of Virtual Machines*.

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 AI Database**, click **Oracle 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 **More actions**, then click **Move resource**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

6. In the resulting dialog, choose the new compartment for the VM cluster, and click **Move**.

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 AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.

   The Exadata VM Clusters page is selected by default.
2. Click the name of the VM cluster that you want to terminate.

   The VM Cluster Details page displays information about the selected VM cluster.
3. Click **Actions**, and then click **Terminate**.
4. In the resulting dialog:
   - Review the message about the backup retention policy
   - Enter the name of the VM cluster
   - Click **Terminate** to confirm the action.

> **Note**
>
> The database stays in a terminated state with backups listed until all backups are expired.
>
> 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 **Oracle AI Database**, click **Oracle 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 or Remove a VM from a VM Cluster
  To terminate or 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 AI Database**, click Oracle Exadata Database Service on Exascale Infrastructure.
   The Exadata VM Clusters page is selected by default.
2. On the Exadata VM Clusters page, ensure that the compartment filter shows the compartment that contains the VM Cluster where you want to add the VM cluster, and click the name of the VM cluster to which you want to add a virtual machine.
3. Click the **Virtual Machines** tab. The VM Cluster details page for the cluster is displayed.
4. Click **Add Virtual Machines**.
5. In the Add Virtual Machines dialog, click to select the number of virtual machines that you want to add.

   > **Note**
   >
   > The VM that is added will have the same resources as the other VMs in the cluster.

> **Note**
>
> Add a VM to a VM Cluster is NOT supported using Terraform.

Terminate or Remove a VM from a VM Cluster

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

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

The Exadata VM Clusters page is selected by default.

2. Choose the **Region** and **Compartment** that contains the VM Cluster for which you want to terminate or remove the virtual machine.
3. Click the name of the VM cluster for which you want to remove a virtual machine.
4. Click the **Virtual Machines** tab.
5. In the list of virtual machines, click the Actions icon (three dots) for a virtual machine that you want to terminate, and then select **Terminate**.
   Alternatively, click the name of the virtual machine to open its details page. From the **Actions** menu, select **Terminate**.
6. On the Terminate Virtual Machine dialog, review the warning messages, enter the name of the virtual machine, 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 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.
>
> TFA is shipped with scripts that run when a particular component is requested, for example, for CRS component, `crscollect.pl` will run a number of `crsctl` commands and gather the input. By default, TFA does not redact collected logs.

Table 5-2 Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
0S: Operating system logs	`oscollect.pl`	• `/var/log/messages` • OSWatcher archive • **Exadata Only:** ExaWatcher archive `/opt/oracle.ExaWatcher/archive/`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-2 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
CRS: Grid Infrastructure and cluster logs	`crscollect.pl`	• `/etc/oracle` • `GIHOME/crf/db/HOSTNAME1` • `GIHOME/crs/log` • `GIHOME/css/log` • `GIHOME/cv/log` • `GIHOME/evm/admin/log` • `GIHOME/evm/admin/logger` • `GIHOME/evm/log` • `GIHOME/log/-/client` • `GIHOME/log/HOSTNAME1` • `GIHOME/log/HOSTNAME1/admin` • `GIHOME/log/HOSTNAME1/client` • `GIHOME/log/HOSTNAME1/crflogd` • `GIHOME/log/HOSTNAME1/crfmond` • `GIHOME/log/HOSTNAME1/crsd` • `GIHOME/log/HOSTNAME1/cssd` • `GIHOME/log/HOSTNAME1/ctssd` • `GIHOME/log/HOSTNAME1/diskmon` • `GIHOME/log/HOSTNAME1/evmd` • `GIHOME/log/HOSTNAME1/gipcd` • `GIHOME/log/HOSTNAME1/gnsd` • `GIHOME/log/HOSTNAME1/gpnpd` • `GIHOME/log/HOSTNAME1/mdnsd` • `GIHOME/log/HOSTNAME1/ohasd` • `GIHOME/log/HOSTNAME1/racg` • `GIHOME/log/HOSTNAME1/srvm` • `GIHOME/log/HOSTNAME1/xag` • `GIHOME/log/diag/asmtool`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-2 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
		• GIHOME/log/diag/clients • GIHOME/log/procwatcher/PRW_SYS_HOSTNAME1 • GIHOME/network/log • GIHOME/opmn/logs • GIHOME/racg/log • GIHOME/scheduler/log • GIHOME/srvm/log • GRIDBASE/crsdata/@global/cvu • GRIDBASE/crsdata/HOSTNAME1/core • GRIDBASE/crsdata/HOSTNAME1/crsconfig • GRIDBASE/crsdata/HOSTNAME1/crsdiag • GRIDBASE/crsdata/HOSTNAME1/cvu • GRIDBASE/crsdata/HOSTNAME1/evm • GRIDBASE/crsdata/HOSTNAME1/output • GRIDBASE/crsdata/HOSTNAME1/ovmmwallets • GRIDBASE/crsdata/HOSTNAME1/scripts • GRIDBASE/crsdata/HOSTNAME1/trace • GRIDBASE/diag/crs/-/crs/cdump • GRIDBASE/diag/crs/HOSTNAME1/crs/cdump • GRIDBASE/diag/crs/HOSTNAME1/crs/incident • GRIDBASE/diag/crs/HOSTNAME1/crs/trace

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-2 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
Database: Oracle Database logs	No DB Specific Script - runs `opatch lsinventory` for the ORACLE_HOME the DB runs from TFA will run ipspack based on the time range for certain DB incidents.	• `ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/cdump` • `ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/trace` • `ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/incident`

Cloud Tool Logs

- **Creg files:** `/var/opt/oracle/creg/*.ini` files with masked sensitive info
- **Cstate file:** `/var/opt/oracle/cstate.xml`
- **Database related tooling logs:**

  If `dbName` specified, `/var/opt/oracle/log/<dbName>`, else collect logs for all databases `/var/opt/oracle/log/`

  If `dbName` specified, `/var/opt/oracle/dbaas_acfs/log/<dbName>`, else collect logs for all databases `/var/opt/oracle/log/<dbName>`
- **Database env files:** If `dbName` specified, `/home/oracle/<dbName>.env`, else collect logs for all databases `/home/oracle/*.env`
- **Pilot logs:** `/home/opc/.pilotBase/logs`
- **List of log directories:**
  - `/var/opt/oracle/log`
  - `/var/opt/oracle/dbaas_acfs/log`
  - `/var/opt/oracle/dbaas_acfs/dbsystem_details`
  - `/var/opt/oracle/dbaas_acfs/job_manager`
  - `/opt/oracle/dcs/log`

DCS Agent Logs

- `/opt/oracle/dcs/log/`

Tooling-Related Grid Infrastructure/Database Logs

- **Grid Infrastructure:** `GI_HOME/cfgtoollogs`
- **Database alertlog:** `/u02/app/oracle/diag/rdbms/*/*/alert*.log`

Health Metrics

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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 5-3 Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
`CpuUtilization`	CPU Utilization	Percentage	Mean	One minute	Five minutes	The CPU utilization is expressed as a percentage, which is aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of 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.
`BlockChanges`	DB Block Changes	Changes per second	Mean	One minute	Five minutes	The Average number of blocks changed per second.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-3 (Cont.) Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-3 (Cont.) Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Guest VM Health Metrics List - Non-Database Metrics

Table 5-4 Guest VM Health Metrics List - Non-Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
`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: `cat /proc/meminfo`.
`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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- CreateExascaleDbStorageVault
- DeleteExascaleDbStorageVault
- GetExascaleDbStorageVault
- ListExascaleDbStorageVaults
- UpdateExascaleDbStorageVault

Exadata VM cluster

- ChangeExadbVmClusterCompartment
- CreateExadbVmCluster
- DeleteExadbVmCluster
- GetExadbVmCluster
- ListExadbVmClusters
- RemoveVirtualMachineFromExadbVmCluster
- UpdateExadbVmCluster

Troubleshooting Virtual Machines Using Console Connections

You can troubleshoot malfunctioning virtual machines using console connections. For example, a previously working Guest VM stops responding.

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

To connect to a running instance for administration and general use, use a Secure Shell (SSH). For more information, see Connecting to a Virtual Machine with SSH.

To make an SSH connection to the serial console, follow these configuration steps.

1. Ensure that you have the correct permissions.
2. Complete the prerequisites, including creating your SSH key pair (in case you don't have one yet).
3. Create the Virtual Machine Serial Console.
4. Connect to the serial console via SSH.

To check the DB server version installed, follow these steps:

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.
3. Click **Exadata Infrastructure** under **Oracle Exadata Database Service on Exascale Infrastructure**.
4. Click the name of the infrastructure that you are interested in.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5. In the resulting Infrastructure Details page, go to the Version section to find the DB Server version installed.

- Required IAM Policies
  An administrator must grant you secure access to the virtual machine console on the Oracle Exadata Database Service on Exascale Infrastructure system through an IAM policy.
- Prerequisites
  You must install an SSH client and create SSH key pairs.
- Create the Virtual Machine Serial Console Connection
  Before you can make a local connection to the serial console, you need to create the virtual machine console connection.
- Make an SSH Connection to the Serial Console
- Using Cloud Shell to Connect to the Serial Console
- Displaying the Console History for a Virtual Machine
- Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems
- Exiting the Virtual Machine Serial Console Connection

Required IAM Policies

An administrator must grant you secure access to the virtual machine console on the Oracle Exadata Database Service on Exascale Infrastructure system through an IAM policy.

This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tools. If you get a message that you don’t have permission or are unauthorized, verify with your administrator what type of access you have and which compartment to work in.

To create virtual machine console connections, an administrator needs to grant user access to read and manage virtual machine console connections through an IAM policy. The resource name for virtual machine console connections is `dbnode-console-connection`. The resource name for the virtual machine is `db-nodes`. The following policies grant users the ability to create virtual machine console connections:

```
Allow group <group_name> to manage dbnode-console-connection in tenancy
Allow group <group_name> to read db-nodes in tenancy
```

Prerequisites

You must install an SSH client and create SSH key pairs.

Install an SSH Client and a Command-line Shell (Microsoft Windows)

Microsoft Windows does not include an SSH client by default. If you are connecting from a Windows client, you need to install an SSH client. You can use PuTTY plink.exe with Windows PowerShell or software that includes a version of OpenSSH such as:

- Git for Windows
- Windows Subsystem for Linux

The instructions in this topic frequently use PuTTY and Windows PowerShell.

Copyright © 2023, 2026, Oracle and/or its affiliates.

If you want to make the console connection from Windows with Windows PowerShell, PowerShell might already be installed on your Windows operating system. If not, follow the steps at the link. If you are connecting to the instance from a Windows client using PowerShell, plink.exe is required. plink.exe is the command link connection tool included with PuTTY. You can install PuTTY or install plink.exe separately. For installation information, see https://www.putty.org.

Create SSH Key Pairs

To create the secure console connection, you need an SSH key pair. The method to use for creating key pairs depends on your operating system. When connecting to the serial console, you must use an RSA key. The instructions in this section show how to create an RSA SSH key pair.

Create the SSH key Pair for Linux

If you're using a UNIX-style system, you probably already have the `ssh-keygen` utility installed. To determine whether the utility is installed, type `ssh-keygen` on the command line. If the utility isn't installed, you can download OpenSSH for UNIX from https://www.openssh.com/portable.html and install it.

1. Open a shell or terminal for entering the commands.
2. At the prompt, enter `ssh-keygen` and provide a name for the key when prompted. Optionally, include a passphrase.
   The keys will be created with the default values: RSA keys of 2048 bits.

   Alternatively, you can type a complete `ssh-keygen` command, for example:

   ```
   ssh-keygen -t rsa -N "" -b 2048 -C "<key_name>" -f <path/root_name>
   ```

Argument	Description
`-t rsa`	Use the RSA algorithm.
`-N "<passphrase>"`	A passphrase to protect the use of the key (like a password). If you don't want to set a passphrase, don't enter anything between the quotes. A passphrase is not required. You can specify one as a security measure to protect the private key from unauthorized use. If you specify a passphrase, when you connect to the instance you must provide the passphrase, which typically makes it harder to automate connecting to an instance.
`-b 2048`	Generate a 2048-bit key. You don't have to set this if 2048 is acceptable, as 2048 is the default. A minimum of 2048 bits is recommended for SSH-2 RSA.
`-C "<key_name>"`	A name to identify the key.
-f <*path/root_name*>	The location where the key pair will be saved and the root name for the files.

Create the SSH Key Pair for Windows Using PuTTY

If you are using a Windows client to connect to the instance console connection, use an SSH key pair generated by PuTTY.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> Ensure that you are using the latest version of PuTTY. See https://www.putty.org.

1. Find `puttygen.exe` in the PuTTY folder on your computer, for example, `C:\Program Files (x86)\PuTTY`. Double-click `puttygen.exe` to open it.
2. Specify a key type of SSH-2 RSA and a key size of 2048 bits:
   - In the **Key** menu, confirm that the default value of **SSH-2 RSA key** is selected.
   - For the **Type of key to generate**, accept the default key type of **RSA**.
   - Set the **Number of bits in a generated key** to 2048 if not already set.
3. Click **Generate**.
4. To generate random data in the key, move your mouse around the blank area in the PuTTY window.
   When the key is generated, it appears under **Public key for pasting into OpenSSH authorized_keys file**.
5. A **Key comment** is generated for you, including the date and timestamp. You can keep the default comment or replace it with your own more descriptive comment.
6. Leave the **Key passphrase** field blank.
7. Click **Save private key**, and then click **Yes** in the prompt about saving the key without a passphrase.
   The key pair is saved in the PuTTY Private Key (PPK) format, which is a proprietary format that works only with the PuTTY tool set.

   You can name the key anything you want but use the `ppk` file extension. For example, `mykey.ppk`.
8. Select all of the generated keys that appear under **Public key** for pasting into OpenSSH `authorized_keys` file, copy it using **Ctrl + C**, paste it into a text file, and then save the file in the same location as the private key.

   > **Note**
   >
   > Do not use the **Save public key** option because it does not save the key in the OpenSSH format.

   You can name the key anything you want, but for consistency, use the same name as the private key and a file extension of `pub`. For example: `mykey.pub`.
9. Write down the names and locations of your public and private key files. You need the public key when creating an instance console connection. You need the private key to connect to the instance console connection using PuTTY. For example: `$HOME\Documents\mykey.ppk`.

To create a connection using the SSH key pair generated using PuTTY

For more information about generating SSH key pairs, see Create the SSH Key Pair for Windows Using PuTTY

Do the following on the Create serial console access window:

Copyright © 2023, 2026, Oracle and/or its affiliates.

1. Paste the SSH Key generated from OpenSSH format or choose **Upload SSH key file** and provide the path of the public key saved at step 8 in Create the SSH Key Pair for Windows Using PuTTY.
2. Once the connection is **Active**, click **Copy serial console connection for Windows**.
3. Paste the connection string copied from the previous step into a text file.
4. In the text file, replace `<PATH_FILE_PUTTY_PRIVATE.ppk>` to point to your PuTTY Private Key (PPK) file path on your computer. For example, if you have saved `.ppk` file at `$HOME\Documents\mykey.ppk`.
5. Paste the modified connection string into the PowerShell window, and then press **Enter** to connect to the console.

Sign in to a Virtual Machine From the Serial Console

If you want to sign in to a virtual machine using a virtual machine console connection, you can use Secure Shell (SSH) connection to sign in. If you want to sign in with a username and password, you need a user account with a password. Oracle Exadata Cloud does not set a default password for the `opc` or `root` users. Therefore, if you want to sign in as the `opc` or `root` user, you need to create a password for the `opc` or `root` user. Otherwise, add a different user with a password and sign in as that user. This should be completed in advance, before a potential situation that might require you to log in to the serial console.

Connect Through Firewalls

If the client you will use to access the serial console is behind a firewall, you must ensure that this client can reach the required endpoint to access the serial console of the virtual machine. The client system connecting to the serial console must be able to reach the serial console server (for example, `vm-console-ad1.exacs.us-ashburn-1.oci.oraclecloud.com`) over SSH using port 443, directly or through a proxy.

Create the Virtual Machine Serial Console Connection

Before you can make a local connection to the serial console, you need to create the virtual machine console connection.

Virtual machine console connections are limited to one client at a time. If the client fails, the connection remains active for approximately five minutes. During this time, no other client can connect. After five minutes, the connection is closed, and a new client can connect. During the five-minute timeout, any attempt to connect a new client fails with the following message:

```
channel 0: open failed: administratively prohibited: console access is
limited to one connection at a time
```

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the VM Cluster that you're interested in.
3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
   Under **Resources**, **Console connections** is selected by default.
4. Click **Create console connection**.
5. In the resulting Create serial console access window, you have three options for adding the SSH key

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **Generate a key pair for me**: You can have Oracle Cloud Infrastructure generate an SSH key pair to use. If you are using PowerShell or PuTTY to connect to the instance from a Windows client, you cannot use the generated SSH key pair without first converting it to a `.ppk` file.
- **Upload public key file**: Browse to a public key file on your computer. If you followed the steps in Creating SSH Key Pairs in the Prerequisites section to create a key pair, use this option to navigate to the `.pub` file.
- **Paste public key**: Paste the content of your public key file into the text box.

6. Click **Create console connection**.
   When the console connection has been created and is available, the state changes to **Active**.

Make an SSH Connection to the Serial Console

After you create the console connection for the virtual machine, you can connect to the serial console using a Secure Shell (SSH) connection. When making an SSH connection to the serial console, you must use an RSA key. You can use the same SSH key for the serial console that was used when you launched the instance, or you can use a different SSH key.

When you are finished with the serial console and have terminated the SSH connection, you should delete the serial console connection. If you do not disconnect from the session, Oracle Cloud Infrastructure terminates the serial console session after 24 hours and you must reauthenticate to connect again.

Validate Server Host Keys

When you first connect to the serial console, you're prompted to validate the fingerprint of the server host key. The fingerprint of the server host key is the SHA256 hash of the server host's public SSH key. The server SSH handshake response is signed with the associated private key. Validating the server host key's fingerprint protects against potential attacks.

When you make a manual connection to the serial console, the fingerprint of the server host key is not automatically validated. To manually validate the fingerprint, compare the fingerprint value displayed in the Oracle Cloud Infrastructure Console to the value of the RSA key fingerprint that appears in the terminal when you connect.

To find the fingerprint of the server host key in the Console, on the Virtual Machine details page, under **Resources**, click **Console connection**. The table displays the fingerprint of the server host key. The fingerprint in the Console should match the value of the **RSA key fingerprint** shown in the terminal when you connect to the serial console.

The server host keys are periodically rotated for security purposes. Key rotation reduces the risk posed when keys are compromised by limiting the amount of data encrypted or signed by one key version. When your key is rotated and you try to connect to the serial console, a warning appears indicating a potential attack. The warning includes an `Host key verification failed` error and a line number in your `.ssh/known_hosts` file. Delete that line in your `.ssh/known_hosts` file and then reconnect to the serial console. You are then prompted to accept a new server host key fingerprint.

Connect from Mac OS X and Linux Operating Systems

Use an SSH client to connect to the serial console. Mac OS X and most Linux and UNIX-like operating systems include the SSH client OpenSSH by default.

To connect to the serial console using OpenSSH on Mac OS X or Linux:

Copyright © 2023, 2026, Oracle and/or its affiliates.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the VM Cluster that you're interested in.
3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
4. On the Virtual Machine details page in the Oracle Cloud Infrastructure Console, under **Resources**, click **Console connection**.
5. Click the Actions menu (three dots), and then click **Copy serial console connection for Linux/Mac**.
6. Paste the connection string into a terminal window on a Mac OS X or Linux system, and then press Enter to connect to the console.
   If you are not using the default SSH key or ssh-agent, modify the serial console connection string to include the identity file flag, `-i` to specify the private key portion for the SSH key to use, for example, `id_rsa`. Specify this flag for both the SSH connection and the SSH Proxy Command, as shown in the following line:

   ```
   ssh -i /<path>/<ssh_key> -o ProxyCommand='ssh -i /<path>/<ssh_key> -W
   %h:%p -p 443...
   ```

7. If prompted, validate and accept the fingerprint of the server host key.
   If you have previously accepted a fingerprint for the server host key but the key has been rotated, a warning appears indicating a potential attack. The warning includes an `Host key verification failed` error and a line number in your `.ssh/known_hosts` file. Delete the specified line in your `.ssh/known_hosts` file and then reconnect to the serial console. Validate and accept the new server host key fingerprint.
8. Press **Enter** again to activate the console.
   If the connection is active, a message appears in the console:

   ```
   ================================================
   IMPORTANT: You are now connected to the serial console for this VM. This should be
   used in emergency situations only.
   See product documentation for more details and alternative connectivity options for
   normal operations
   ================================================
   ```

9. Reboot your virtual machine.
   You do not need to enter a user name or password. If the Virtual Machine is functional and the connection is active, the serial output appears in your console. If the serial output does not appear in the console, the Guest VM operating system is not booting.

   For more troubleshooting options, see *Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems*.

   a. Go to the ExaDB-XS VM Cluster Details page.
   b. Under **Resources**, click **Virtual Machines**.
   c. Select **Reboot** from the Actions menu (three dots) for the virtual machine that you want to reboot.

Connect from Windows Operating Systems

The steps to connect to the serial console from Microsoft Windows PowerShell are different from the steps for OpenSSH. The following steps do not work in the Windows terminal.

If you are connecting to the instance from a Windows client using PowerShell, `plink.exe` is required. `plink.exe` is the command link connection tool included with PuTTY. You can install

Copyright © 2023, 2026, Oracle and/or its affiliates.

PuTTY or install `plink.exe` separately. For more information, see *Installing an SSH Client and a Command-line Shell (Windows)*.

To connect to the serial console on Microsoft Windows:

1. On the Virtual Machine details page in the Oracle Cloud Infrastructure Console, under **Resources**, click **Console connection**.
2. Click the Actions menu (three dots).
   Depending on which SSH client you are using, do one of the following:
   - If you are using Windows PowerShell, click **Copy serial console connection for Windows**.
   - If you are using OpenSSH, click **Copy serial console connection for Linux/Mac**.

   > **Note**
   >
   > The copied connection string for Windows contains the parameter `-i` specifying the location of the private key file. The default value for this parameter in the connection string references an environment variable that might not be configured on your Windows client, or it might not represent the location where the private key file is saved. Verify the value specified for the `-i` parameter and make any required changes before proceeding to the next step.
3. Paste the connection string copied from the previous step into a text file so that you can add the file path to the private key file.
4. In the text file, replace `$env:homedrive$env:homepath\oci\console.ppk` with the file path to the `.ppk` file on your computer. This file path appears twice in the string. Replace it in both locations.

   > **Note**
   >
   > For PuTTY versions 0.82 and above, add parameter -legacy-stdio-prompts to the first call to plink.
5. Paste the modified connection string into the PowerShell window or your OpenSSH client, and then press **Enter** to connect to the console.
6. If prompted, validate and accept the fingerprint of the server host key.
   If you have previously accepted a fingerprint for the server host key, but the key has been rotated, a warning appears indicating a potential attack. The warning includes a Host key verification failed error and a line number in your `.ssh/known_hosts` file. Delete the specified line in your `.ssh/known_hosts` file and then reconnect to the serial console. Validate and accept the new server host key fingerprint.
7. Press **Enter** again to activate the console.
8. Reboot your virtual machine.
   You do not need to enter a user name or password. If the Virtual Machine is functional and the connection is active, the serial output appears in your console. If the serial output does not appear in the console, the Guest VM operating system is not booting.

   For more troubleshooting options, see Troubleshooting Virtual Machines from Guest VM Console Connections.

   a. Go to the ExaDB-XS VM Cluster Details page.
   b. Under **Resources**, click **Virtual Machines**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

c. Select **Reboot** from the Actions menu (three dots) for the virtual machine that you want to reboot.

To create a connection using the SSH key pair generated using the OCI Console:

Do the following on the Create serial console access window:

1. Click **Generate a key pair for me**.
2. Click **Save Private Key**.
3. Click **Create console connection**.

> **Note**
>
> Ensure that you are using the latest version of PuTTY, see http://www.putty.org.

4. Find `puttygen.exe` in the PuTTY folder on your computer, for example, `C:\Program Files (x86)\PuTTY. Double-click puttygen.exe` to open it.
5. On the PuTTY Key Generator, click the **Conversions** menu and then click **Import**.
6. On the Windows Explorer, select OCI Console generated SSH key (step 1) and then click **Open**.
   PuTTY imports the key and displays information about the key on the PuTTY Key Generator window.
7. Click **Save private key**.
8. Click **Yes** when prompted about saving the key without a passphrase.
   The key pair is saved in the PuTTY Private Key (PPK) format, which is the proprietary format that works only with the PuTTY tool set.

   You can name the key anything you want but use the `.ppk` file extension. For example, `$HOME\Desktop\key-vm-console.ppk`.
9. Use a text editor to change the command to point to your PuTTY Private Key (PPK) path. Replace *<PATH_FILE_PUTTY_PRIVATE.ppk>* to point to your PuTTY Private Key (PPK) file path on your computer. For example, if you have saved `.ppk` file at `$HOME\Desktop\key-vm-console.ppk`.
10. Paste the modified connection string into the PowerShell window, and then press **Enter** to connect to the console.

To convert a generated .key private key file:

1. Open PuTTYgen.
2. Click **Load**, and select the private key generated when you created the instance.
   The extension for the key file is `.key`.
3. Click **Save private key**.
4. Specify a name for the key.
   The extension for the new private key is `.ppk`.
5. Click **Save**.

Using Cloud Shell to Connect to the Serial Console

You can connect to the serial console quickly and easily using the Cloud Shell integration. Cloud Shell is a web browser-based terminal accessible from the Console. The Cloud Shell integration automatically creates the instance console connection and a temporary SSH key.

Copyright © 2023, 2026, Oracle and/or its affiliates.

The only prerequisite for connecting to the serial console from Cloud Shell is granting users the correct permissions. For an introductory walkthrough of using Cloud Shell, see Using Cloud Shell.

> **Note**
>
> - By default, Cloud Shell limits network access to OCI internal resources in your tenancy home region only unless you have enabled the Cloud Shell managed Public Network. Your administrator must configure an Identity policy to enable Cloud Shell Public Network. For more information, see Cloud Shell Networking.
> - You cannot concurrently connect to more than one DB node using Cloud Shell. As an example, if you have an open connection to *DBnode1* and want to connect to *DBnode2*, you must first exit the active Cloud Shell from *DBnode1* and then establish a connection to *DBnode2*.

When you are finished with the serial console and have terminated the SSH connection, you should delete the serial console connection. If you do not disconnect from the session, Oracle Cloud Infrastructure terminates the serial console session after 24 hours and you must re-authenticate to connect again.

- To connect to the serial console using Cloud Shell

Related Topics

- Cloud Shell
- Using Cloud Shell
- Required IAM Policies

To connect to the serial console using Cloud Shell

1. Sign in to the Console.
2. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
3. On the instance details page in the Oracle Cloud Infrastructure Console, under **Resources**, click **Console connection**.
4. Click **Launch Cloud Shell connection**.
   This action displays the Cloud Shell in a "drawer" at the bottom of the Console.
5. If a console connection already exists, you are asked if you want to delete the existing resource. Press **y**, and then press **Enter**.
6. When you are done, exit the instance console connection.

Related Topics

- Exiting the Virtual Machine Serial Console Connection

Displaying the Console History for a Virtual Machine

You can capture and display recent serial console data for a Virtual Machine. The data includes configuration messages that occur when the Virtual Machine boots, such as kernel and BIOS messages, and is useful for checking the status of the Virtual Machine or diagnosing and troubleshooting problems.

Copyright © 2023, 2026, Oracle and/or its affiliates.

The console history captures up to a megabyte of the most recent serial console data for the specified Virtual Machine. Note that the raw console data, including multi-byte characters, is captured.

The console history is a point-in-time record. To troubleshoot a malfunctioning Virtual Machine using an interactive console connection, use a serial console connection.

- Managing Console History Data

Managing Console History Data

You can use the Console or API to manage console history captures. Console history lets you see serial output from your Virtual Machine without having to connect to the instance remotely. The console history can be used to audit previous access and actions taken with the serial console.

On the instance details page in the Console, you can capture and download console histories, view and edit metadata details, and delete console history captures.

- Using the Console to Capture the Console History
- Using the Console to Download Console History Captures
- Using the Console to View Console History Captures
- Using the Console to View and Edit the Metadata Details of a Console History Capture
- Using the Console to Delete Console History Captures
- Using the API to Manage the Console History Data

Using the Console to Capture the Console History

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
   Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. Click the name of the history that you're interested in.
6. On the resulting window, click **Download** to download a copy of the console history.
7. Click **Save and close** to save the history and close the window.

Using the Console to Download Console History Captures

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
   Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. Click the name of the history that you're interested in.
6. On the resulting window, click **Download** to download a copy of the console history.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Using the Console to View Console History Captures

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
   Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. Click the name of the history that you're interested in.
6. On the console history list, for the console history capture that you want to view, click the **Actions** menu, and then click **View details**.

Using the Console to View and Edit the Metadata Details of a Console History Capture

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
   Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. On the console history list, for the console history capture that you want to view, click the **Actions** menu, and then click **View details**.
6. Optionally, edit the name for the console history. Avoid entering confidential information.
7. To view or edit tags, click **Show tagging options**.
8. To edit or remove tags, click the edit icon next to the tag. To edit a tag, in the **Edit Tag** dialog, make any changes, and then click **Save**. To remove a tag, click **Remove Tag**.
9. Click **Save and close**.

Using the Console to Delete Console History Captures

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
   Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. On the console history list, for the console history capture that you want to view, click the **Actions** menu, and then click **Delete**.
6. On the confirmation dialog, click **Delete console history**.

Using the API to Manage the Console History Data

Review the list of API calls to manage console history data.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

For the complete list of APIs, see Database Service API.

Use the following API operations to manage the console history data.

- To capture the console history, use the **createDbNodeConsoleHistory** method.
- To get details of console history metadata, use the **getDbNodeConsoleHistory** method.
- To get the details of console history content, use the **getDbNodeConsoleHistoryContent** method.
- To edit console history metadata, use the **updateDbNodeConsoleHistory** method.
- To list console history captures, use the **listDbNodeConsoleHistories** method.
- To delete console history captures, use the **deleteDbNodeConsoleHistory** method.

Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems

After you are connected with an instance console connection, you can perform various tasks, such as:

- Edit system configuration files.
- Add or reset the SSH keys for the `opc` user.
- Reset the password for the `opc` user.

These tasks require you to boot into a Bash shell in maintenance mode.

To boot into maintenance mode

> **Note**
>
> **Default user and password:**
>
> - **Account**: Grub boot loader
> - **Username**: root
> - **Default Password**: sos1Exadata
> - **Account Type**: Operating system user
>
> For more information, see *Default User Accounts for Oracle Exadata*.

1. Reboot the VM from the VM Cluster.
2. For virtual machines running Oracle Linux 7.x or Oracle Linux 8.x, when the reboot process starts, switch back to the terminal window, and you see Console messages start to appear in the window. As soon as the **GRUB boot menu** appears, use the **up/down arrow key** to stop the automatic boot process, enabling you to use the boot menu.
3. In the boot menu, highlight the top item in the menu, and press **e** to edit the boot entry.
4. In edit mode, use the **down arrow key** to scroll down through the entries until you reach the line that starts with **linux16**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5. At the end of that line, add the following:

   ```
   init=/bin/bash
   ```

6. Reboot the instance from the terminal window by entering the keyboard shortcut **CTRL**+**X**. When the instance has rebooted, you see the Bash shell command-line prompt, and you can proceed with the following procedures.

To edit the system configuration files

1. From the Bash shell, run the following command to load the SElinux policies to preserve the context of the files you are modifying:

   ```
   /usr/sbin/load_policy -i
   ```

2. Run the following command to remount the root partition with read/write permissions:

   ```
   /bin/mount -o remount, rw /
   ```

3. Edit the configuration files as needed to try to recover the instance.
4. After you have finished editing the configuration files, to start the instance from the existing shell, run the following command:

   ```
   exec /usr/lib/systemd/systemd
   ```

   Alternatively, to reboot the instance, run the following command:

   ```
   /usr/sbin/reboot -f
   ```

To add or reset the SSH key for the opc user

1. From the Bash shell, run the following command to load the SElinux policies to preserve the context of the files you are modifying:

   ```
   /usr/sbin/load_policy -i
   ```

2. Run the following command to remount the root partition with read/write permissions:

   ```
   /bin/mount -o remount, rw /
   ```

3. From the Bash shell, run the following command to change to the SSH key directory for the `opc` user: `cd ~opc/.ssh`
4. Include your public key entry to the `authorized_keys` file.

> **Note**
>
> You can edit the file and remove your previous key if you want to. However, make sure to keep the cloud automation keys to prevent cloud automation from breaking.
>
> ```
> echo '<contents of public key file>' >> authorized_keys
> ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

5. Restart the instance by running the following command:

   ```
   /usr/sbin/reboot -f
   ```

To reset the password for the opc user

1. From the Bash shell, run the following command to load the SElinux policies to preserve the context of the files you are modifying.
   This step is necessary to sign in to your instance using SSH and the Console.

   ```
   /usr/sbin/load_policy -i
   ```

2. Run the following command to remount the root partition with read/write permissions:

   ```
   /bin/mount -o remount, rw /
   ```

3. Run the following command to reset the password for the `opc` user:

   ```
   sudo passwd opc
   ```

4. Restart the instance by running the following command:

   ```
   sudo reboot -f
   ```

> **Note**
>
> Setting a `root` password would be an acceptable alternative to setting an `opc` password.

Exiting the Virtual Machine Serial Console Connection

To exit the serial console connection

When using SSH, the ~ character at the beginning of a new line is used as an escape character.

1. To exit the serial console, enter:

   ```
   ~.
   ```

2. To suspend the SSH session, enter:

   ```
   ~^z
   ```

   The ^ character represents the **CTRL** key.

3. To see all the SSH escape commands, enter:

   ```
   ~?
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

To delete the serial console connection for a Virtual Machine

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the VM Cluster that you're interested in.
3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
   Under **Resources**, **Console connection** is selected by default.
4. Click the Actions menu, and then click **Delete**. Confirm when prompted.

Manage Exascale Database Vaults on Oracle 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. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**
2. Click **Exascale Storage Vaults**.
3. Select the Vault for which you want to view information.

Deleting Exascale Database Storage Vaults

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**
2. Click **Exascale Storage Vaults**.
3. Select the Vault that you want to delete.
4. Click **Actions** and 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 Vault

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**
2. Click **Exascale Storage Vaults**.
3. Select the Vault that you want to scale.
4. Click **Scale Storage Vault**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5. On the Scale Storage Vault dialog, scale the storage manually. You will also have the option to auto scale the storage up to a predefined limit.
   - Scale the storage capacity manually: Enter storage capacity. 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.
   - Scale the storage using storage auto scaling: Enable auto scale on the Vault by selecting **Storage auto scaling**.
     If enabled, specify **Auto scaling limit (GB)**.

   > **Note**
   > – The storage auto scaling feature is not available with Oracle Database 19c.
   > – Auto scale is primarily intended to accommodate the ongoing, organic growth of databases. For large, sudden data loads or rapid data growth that can exceed 5% of the allocated vault size within a short time period, it is recommended to manually scale storage in advance to avoid potential out-of-space errors.
   > – Tablespaces with auto-extend enabled can cause out-of-space errors if the auto-extend increment size is greater than 5% of the allocated vault storage. This could be caused if a single, large tablespace auto-extension suddenly exceeds the available space in the vault before crossing the auto scale threshold. Or, it could also be caused by multiple tablespaces extensions occurring concurrently, resulting in the same effect. It is recommended that you manage your tablespace auto-extend increment settings to avoid having more than 5% of the allocated vault storage added at once when using auto scale.

6. Click **Scale**.

Manage Software Images

- Using Software Images in Oracle Cloud Infrastructure
- Using a Software Image with an Oracle Exadata Database Service on Exascale Infrastructure Instance
  Provision database homes, patch custom images, or set up Data Guard standby database using custom database images.
- 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
  Database software images are resources within your tenancy that you create prior to provisioning or patching a Oracle Exadata Database Service on Exascale Infrastructure instance, Database Home, or database.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 software. You can find the `opatch` utility in the `$ORACLE_HOME/Opatch` directory.

Creation and Storage of Software Images

Database software images are resources within your tenancy that you create prior to provisioning or patching a Oracle Exadata Database Service on Exascale Infrastructure instance, Database Home, or database.

There is no limit on the number of database software images you can create in your tenancy, and you can create your images with any Oracle Database software version and update supported in Oracle Cloud Infrastructure.

Database software images are automatically stored in Oracle-managed Object Storage and can be viewed and managed in the Oracle Cloud Infrastructure Console. Note that database software images incur Object Storage usage costs. Database software image are regional-level resources and can be accessed from any availability domain within their region.

See To create a database software image for information on creating an image.

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 software. 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. This file can then be uploaded to the OCI Console during the creation of a custom Database Software Image to add the exact set of patches used by the source Database Home to the list of patches included in the software image. You can find 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
```

2. Use the `lsinventory` output file to extract the additional One Off Patches applied to a specific Oracle Home.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Using a Software Image with an Oracle Exadata Database Service on Exascale Infrastructure Instance

Provision database homes, patch custom images, or set up Data Guard standby database using custom database images.

Provisioning: After you create a database software image, you can use it to create an Oracle Database Home in an Oracle Exadata Database Service on Exascale Infrastructure instance. For more information, see To create a new Database Home in an existing Exadata Cloud Service instance.

Patching: To patch a database in an Oracle Exadata Database Service on Exascale Infrastructure instance using a custom database software image, create the Database Home using the image, and then move the database to that Database Home. For more information, see Patching Individual Oracle Databases in an Exadata Cloud Service Instance.

Setting up Data Guard: When creating an Oracle Data Guard association, you can use a custom database software image to create a new Database Home for the new standby database. For more information, see To enable Oracle Data Guard on an Exadata Cloud Service instance database.

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 Database Software Image
  Use the following instructions to delete a database 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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Under **Resources**, click **Software Images**.
3. Click **Create Software Image**.
4. In the resulting Create software image page, click **Database Software image**.
5. In the **Display name** field, provide a display name for your image. Avoid entering confidential information.
6. Choose your **Compartment**.
7. Choose a **Database release**.
8. Choose the **Database version** for your image. The latest approved Database Release Update is available by default. You can also select one of the two preceding Release

Copyright © 2023, 2026, Oracle and/or its affiliates.

Updates (N-1 or N-2) after acknowledging the associated security risks. To use a Release Update older than N-2, you must submit a Service Request.

9. Optionally, you can enter a comma-separated list of one-off (interim) patch numbers.
10. Optionally, you can upload an Oracle Home inventory file from an existing Oracle Database. See Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Home for instructions on creating an inventory file using OPatch.
11. Click **Show Advanced Options** to add **tags** to your database software image. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see Resource Tags. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
12. Click **Create Software Image**.

To create a database software image from a Database Home

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on 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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The Exadata VM Clusters page is selected by default.
2. Click the name of the VM cluster that you want to update the database software image.
3. Click **Database homes**.
4. Click the name of the Database Home that you want to update.
5. Click **Updates**.
6. Choose a **Compartment**.
7. Click **Type**, select **Custom**, and then click **Apply Filter**.
8. 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

9. Click the Actions button (three dots) for the update you're interested in, and select **Precheck**.
10. On the resulting Confirm precheck dialog, click **Precheck** to continue.
11. After running the precheck successfully, select **Apply Database Home update** from the Actions button (three dots).
12. On the resulting Confirm apply dialog, click **Apply** to continue.

To Delete a Database Software Image

Use the following instructions to delete a database software image.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle 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**.

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

Flexible Software Release Selection for Exadata Database

Oracle applies release selection controls to Database and Grid Infrastructure release updates that you can select for lifecycle operations.

By default, lifecycle operations use or offer the latest available release update or maintenance release (N). This helps ensure that newly created or updated resources use software containing the latest applicable security fixes.

For lifecycle operations that support release selection, you can select an older release update or maintenance release within the supported self-service range after acknowledging the message. The supported self-service range is generally N-1 and N-2 for Database and Grid Infrastructure release updates. Releases earlier than the applicable self-service range require approval through a Service Request with an approved exception.

The availability of release selection depends on the service and lifecycle operation.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> Use the following guidelines to determine which Oracle Database (DB) and Oracle Grid Infrastructure (GI) you can request for. These guidelines are subject to change.
>
> - **Recommended:** The latest Oracle AI Database 26ai Release Updates for DB and GI as recommended in the KB article KB153930.
> - **Available through self-service acknowledgment:** The two previous DB and GI Release Updates.
> - **Older releases:** Require an SR with an approved exception. Oracle Database 19c requests are limited to release 19.25 or later. Once approved, the requested release is enabled only for creating a custom database software image, which customers must use for provisioning or updates. GI releases older than the two previous Release Updates are not available.

How Release Selection Controls Work

Release selection controls apply when a lifecycle operation creates a new software home, provisions a new resource, or updates an existing resource to a selected release update or maintenance release.

The following behavior generally applies:

- N: The latest approved release update or maintenance release is available by default and is the recommended release.
- N-1 and N-2: Where supported, you can select either of the two preceding release updates or maintenance releases after review and acknowledgment.
- Releases earlier than N-2: Submit a Service Request to request access to the required release update or maintenance release.

Release selection controls can apply to operations such as:

- Creating Database and Grid Infrastructure custom software images.
- Creating Database homes and Grid Infrastructure homes.
- Creating databases when the operation creates a new Database home.
- Patching Database homes and Grid Infrastructure homes.
- Provisioning VM clusters where Grid Infrastructure release update selection is supported.
- Updating or upgrading Guest VM OS.

Operations That Use Existing Database Homes

Lifecycle operations that use an existing Database home aren't subject to the new release selection restrictions.

For example, this behavior applies when:

- Creating a database in an existing Database home.
- Restoring a database to an existing Database home.
- Creating a Data Guard standby in an existing Database home, where applicable.
- Moving a database to an existing Database home, subject to the requirements for the move operation.

Copyright © 2023, 2026, Oracle and/or its affiliates.

If the selected Database home is running a Database Release Update earlier than the latest approved Release Update, the Console displays a message about using the older release update. You can continue the operation using the existing Database home.

Selecting an Older Release Update or Maintenance Release

When a lifecycle operation supports self-service selection of older releases, you can select either of the two release updates or maintenance releases preceding the latest available release.

When you select an eligible older release update or maintenance release, the Console displays information about it. Review the information and acknowledge to continue.

If the required release update or maintenance release is earlier than those available for self-service selection, submit a Service Request to request access.

Messages for Older Releases

When you select an eligible older release update or maintenance release, the Console displays a message about this and recommends using the latest approved release update or maintenance release.

To know more about the message and referenced My Oracle Support note for an older Database or Grid Infrastructure Release Update, see Recommendations to Help Protect Oracle Databases from Emerging AI-enabled Security Threats

For lifecycle operations that support self-service selection of older release updates or maintenance releases, you must acknowledge the message before you can proceed.

For operations that use an existing Database home, release selection restrictions don't apply. If the selected Database home is running a Database Release Update earlier than the latest approved Release Update, the Console displays information about it, but you can continue the operation using that Database home.

Older Custom Software Images

The release updates available for custom software image selection depend on the lifecycle operation and applicable release selection controls.

For operations where an older custom software image isn't available for self-service selection, the image isn't displayed as an available selection. To use an older custom software image that isn't available, submit a Service Request.

Oracle recommends using custom software images based on the latest applicable Release Update to receive the most recent security fixes. For more information, see Recommendations to Help Protect Oracle Databases from Emerging AI-enabled Security Threats.

Existing Resources

These release selection controls don't automatically upgrade existing databases, Database homes, Grid Infrastructure homes, VM clusters, or other existing resources that are running older Database or Grid Infrastructure release updates.

Existing resources can continue to operate according to the applicable support and lifecycle policies. When you perform a lifecycle operation on an existing resource, the release selection controls depend on the operation. For example, an operation that creates a new Database home is subject to the applicable release update selection rules, while an operation that uses an existing Database home isn't restricted based solely on the release update installed in that home.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Acknowledging Older Release Updates and Maintenance Releases in API Requests
  Some lifecycle operations allow you to use older Database and Grid Infrastructure release updates, and Exadata System Software maintenance releases instead of the latest eligible releases.

Acknowledging Older Release Updates and Maintenance Releases in API Requests

Some lifecycle operations allow you to use older Database and Grid Infrastructure release updates, and Exadata System Software maintenance releases instead of the latest eligible releases.

When an operation requires acknowledgment for an older release update or maintenance release, set the `isOlderVersionAcknowledged` property to true in the request.

Request Parameter

Use the following request property when you select or apply an older Database or Grid Infrastructure release update, or Exadata System Software maintenance release:

`isOlderVersionAcknowledged`

Set `isOlderVersionAcknowledged` to `true` to acknowledge the use of the older release update or maintenance release and continue with the operation.

This property applies to supported create, update, and apply operations where you can select or apply an older release update or maintenance release.

The property doesn't apply to operations that don't select or apply a release update or maintenance release, including:

- Retrieving resource details
- Viewing patch, update, or upgrade history
- Deleting resources
- Updating only tags
- Generating software image PAR URLs
- Changing compartments

Response Attribute

For applicable release-list and other read operations, the response can include the following attribute:

`isOlderVersionAcknowledgmentRequired`

This attribute indicates whether acknowledgment is required when you select an older release update or maintenance release.

When `isOlderVersionAcknowledgmentRequired` indicates that acknowledgment is required, set `isOlderVersionAcknowledged` to true in the corresponding create, update, or apply request to use the older release update or maintenance release.

Applicable Operations

Acknowledgment for older release updates or maintenance releases applies to supported lifecycle operations such as:

- Creating Database or Grid Infrastructure software images

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Creating Database Homes
- Patching a Database Home or applying a custom Database software image
- Creating or updating DB Systems when a release update is selected
- Creating VM Clusters
- Applying Grid Infrastructure release updates or Exadata System Software maintenance releases to VM Clusters
- Performing database operations that select or move to a target release update or Database Home
- Creating a Data Guard association when the operation creates a new standby Database Home or DB System
- Creating Grid Infrastructure Homes

For applicable release and update list operations, use the `isOlderVersionAcknowledgmentRequired` response attribute to determine whether acknowledgment is required before submitting the corresponding lifecycle request.

Using the Acknowledgment Property

When an operation requires acknowledgment:

1. Retrieve or select the target Database or Grid Infrastructure release update, Exadata System Software maintenance release, patch, or software image.
2. Check whether acknowledgment is required for the selected release update or maintenance release.
3. In the create, update, or apply request, set `isOlderVersionAcknowledged` to `true`.
4. Submit the request.

The acknowledgment applies to the target release update, maintenance release, patch, or software image selected for the lifecycle operation.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the name of the VM cluster where you want to create the database home.

   The VM Cluster details page opens.
3. Click the **Database Homes** tab.

   A list of Database Homes is displayed.
4. Click **Create database home**.
5. In the **Create database home** dialog, enter the following:
   - **Database Home display name:** The display name for the Database Home. Avoid entering confidential information.
     Database image: Determines what Oracle Database version is used for the Database home. You can have databases with different minor versions in the same Database home. The major versions must remain the same. By default, the latest Oracle-published database software image is selected.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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-published 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 **Compartment** filter and **Database version** selectors to limit the list of custom database software images to a specific compartment or Oracle Database software major release version.

  The Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.

> **Note**
>
> By default, you can create a Database Home using the latest approved Database Release Update (N). You can also select either of the two preceding Release Updates (N-1 or N-2) after acknowledging the associated security risks.
> To create a Database Home using a Release Update older than N-2, submit a Service Request.

> **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.

6. 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

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

   The VM Clusters page is selected by default.
2. 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.
3. Click the **Database Homes** tab.
4. 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.
5. Click **Create Image from Database Home**.
6. In the **Create Database Software Image** panel, enter a **Display name** and select a compartment for the software image.

   If the Database Home uses the latest approved Database Release Update, you can create the software image by default. If the Database Home uses one of the two preceding Release Updates (N-1 or N-2), you can create the software image after acknowledging the associated security risks. To use a Database Home with a Release Update older than N-2, you must submit a Service Request.
7. 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.
- To Move a Database to Another Database Home
  Learn how 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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
2. 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.
3. On the VM Cluster details page, click the **Database Homes** tab.
4. In the list of Database Homes, find the Database Home you are interested in, and then click its name to display details about it.

To delete a database home

You cannot delete a Database Home that contains databases. You must first terminate the databases to empty the Database Home. See To terminate a database to learn how to terminate a database.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
2. 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.
3. On the VM Cluster Details page, click the **Database Homes** tab.
4. In the list of Database Homes, find the Database Home you want to delete, and then click its name to display details about it.
5. Click the **Actions** button, and select **Delete** from the list.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **⚠ Caution**
>
> You must not use the console to delete a database marked as "Failed" if you previously attempted to delete the database using the Database as a service command-line interface API (`dbaascli`), and afterward reuse the diskgroups in that database to create another database using the `dbaascli` API.
>
> If you delete the Failed status database from the Console after attempting to delete it with the command-line API, then this can result in detaching and deleting the underlying reused Exascale Direct Volumes (EDV), which can cause the failure of the newly created disk groups.

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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
2. 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.
3. On the VM Cluster Details page, click the **Database Homes** tab.
4. In the list of Database Homes, find the Database Home where you want to add tags, and then click its name to display details about it.
5. Click the **Tags** tab.
6. On the Tags page, click **Add** to add additional tags, or select a tag name to manage that tag, and click the Actions button. Select the appropriate action for the tag that you want to manage.

To Move a Database to Another Database Home

Learn how to move a database to another Database home.

> **ⓘ Note**
>
> If your container database has Oracle managed software updates enabled, then you cannot move your database to another home. You will see a dialog box when you try to move a database enrolled in Oracle managed database software updates. You must disable Oracle managed software updates for your database.

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

   The Exadata VM Clusters page is selected by default.
2. Choose your **Compartment** that contains the VM cluster that hosts the database that you want to move.
3. Click the name of the VM cluster that contains the database that you want to move.
4. On the Cloud VM Cluster details page, click the **Container Databases** tab.
5. Click the name of the container database that you want to move.

Copyright © 2023, 2026, Oracle and/or its affiliates.

The Database Details page displays information about the selected database.

6. If your database is enabled by Oracle managed software updates, follow these steps to disable it:
   a. Click **More Actions** and select **Manage database software update**.
   b. In the resulting window, choose **Customer managed software updates**.
   c. Review the warning message.
   d. Click **Save**.
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**.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 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.
- To download the PDB Connection Bundle
  The PDB Connection Bundle is HA-aware and provides connection information for the PDB on all Data Guard peers.
- 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Oracle Exadata Database Service on Exascale Infrastructure supports Oracle AI Database 26ai and Oracle Database 19c.

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)

Provisioning and Managing Exadata Databases

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

- About Oracle Managed Database Software Updates
  Oracle Managed Database Software Updates gives you the option to let Oracle manage the end-to-end database software update lifecycle for enrolled databases.
- 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.

About Oracle Managed Database Software Updates

Oracle Managed Database Software Updates gives you the option to let Oracle manage the end-to-end database software update lifecycle for enrolled databases.

You define your scheduling preference for database software updates, while Oracle automatically prepares, validates, and applies the updates.

For databases enrolled in Oracle Managed Database Software Updates, Oracle periodically validates update readiness, stages software in advance, applies updates during the scheduled maintenance window, and provides proactive notifications throughout the update process.

Update scheduling and notifications

You can enable Oracle Managed Database Software Updates for an existing database or during database provisioning. Define your preferred day and time of the week for database software updates, and Oracle automatically schedules updates when a new Release Update (RU) is available. Once enabled, you can view Oracle Managed Software Updates details, including the next scheduled maintenance, target software version, and maintenance method from the Database Details page.

Subscribe to Oracle Managed Database Software Updates events for database and database home to receive notifications for scheduled maintenance, readiness check results, software update lifecycle changes, and maintenance completion. To enable these notifications, use the Notifications service to create a subscription topic and add the required subscribers. Then, configure an OCI Events rule for Oracle Managed Database Software Updates events related to databases and database homes. When creating the event rule, select the Database service as the event source and look for event types with the **Oracle managed** prefix to identify database and database home events specific to Oracle Managed Database Software Updates. Configure the rule action to send matching events to the Notifications topic so that subscribers are notified whenever the selected events occur.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Maintenance readiness checks

Before the scheduled maintenance window, Oracle periodically performs checks to ensure the enrolled database is ready for the software update. If an issue is detected, the Database Details page displays the issue along with recommended resolution steps. For issues requiring your action, Oracle provides guidance to resolve them before the scheduled maintenance window. For issues that Oracle can remediate, Oracle investigates and resolves the issues before the scheduled maintenance. For more information, see *Troubleshooting Oracle Managed Database Software Updates Issues*.

Successful maintenance outcome

During planned maintenance, Oracle monitors update progress and investigates issues encountered while applying the database software update. Oracle uses delegated access control to the database VMs to investigate and resolve eligible issues without requiring your intervention. If an issue cannot be fixed during the planned maintenance window, Oracle rolls back the update to ensure the database remains healthy and available. Oracle engages with you to resolve the issue before rescheduling the update to apply it during a future maintenance window.

Related Topics

- 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.
- Delegate Access Control
- Notifications
- Creating an Events Rule
- Troubleshooting Oracle Managed Database Software Updates Issues
  Learn to identify and resolve issues in databases that are enrolled for Oracle Managed Database Software Updates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 VM Cluster
  Learn how you can create your first or subsequent databases.
- To scale database storage capacity of a cluster database in the VM cluster
  This topic describes how to scale storage capacity of container databases in the VM cluster configured with Exascale Block Storage.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 Oracle Exadata Database Service on Exascale Infrastructure.
- To Create a Database from the Last Backup
  Use this procedure to create a database from the last backup on Oracle Exadata Database Service on Exascale Infrastructure.
- Enable Oracle Managed Database Software Updates for an Existing Database
  Learn how you can modify an existing database to support Oracle Managed Database Software Updates for 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 Reschedule Oracle Managed Database Software Updates for a Database
  Learn how to reschedule the Oracle managed database software updates for Oracle Exadata Database Service on Exascale Infrastructure.
- To View and Resolve Oracle Managed Database Software Updates Issues
  Learn how to view and resolve the Oracle managed database software updates issues for Oracle Exadata Database Service on Exascale Infrastructure.
- 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 change the encryption management configuration.
- To Create and List Container Database Snapshot
  Using the console, you can create and list the snapshots of the container database (CDB) in a VM Cluster.
- To Create a Clone Database from a Snapshot
  Using the console, you can create a clone database from a container database (CDB) snapshot.
- To Create a Clone of an Existing Container Database
  Using the console, you can create a clone of an existing container database.
- To Delete a Container Database Snapshot
  Using the console, you can delete the snapshot of a container database (CDB) in a VM Cluster.

To Create a Database in an Existing VM Cluster

Learn how you can create your first or subsequent databases.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the name of the VM cluster where you want to create the database.
3. Click the **Container databases** tab.
4. On the Container Databases page, click **Create container database**.
5. In the **Create container 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**
  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>
- **PDB name**: Optionally, specify a PDB name of your choosing. The name must be unique in the VM cluster. If you do not choose a unique PDB name, then Oracle generates one for you.

Copyright © 2023, 2026, Oracle and/or its affiliates.

To avoid potential service name collisions when using Oracle Net Services to connect to the PDB, ensure that the PDB name is unique across the entire VM cluster. If you do not provide the name of the first PDB, then a system-generated name is used.

- **Database release**: The Oracle Database release that you want to deploy on the database.
- **Database Software Management**:
  - **Software Management Preference**: Choose one of these options:
    - **Customer managed software updates**: This option is selected by default for all existing databases.
    - **Oracle managed software updates**: Select this option to enable Oracle-managed database home and software version updates.
      If you choose **Oracle managed software updates**:
      - **Maintenance method**: (read-only) By default, **Rolling** updates is selected. You cannot change this selection.
        Databases are updated in a rolling manner, one database instance at a time, with no downtime.
      - **Schedule Preferences**: Provide the day of the week and the maintenance window that you prefer for Oracle managed database software updates.
      - **Version Preference**: Choose your preferred Release update (RU) for the Oracle-managed database updates. By default, Stay on the latest available Release Update is selected. Options are as follows:
        - Stay on the latest available Release Update
        - Stay one Release Update behind the latest
        - Stay two Release Updates behind the latest

> **Note**
>
> If you select options other than **Stay on the latest available Release Update**, the Confirm selection dialog opens. Read the acknowledgment, check the MOS Note box after reading the MOS Note, and click **Confirm**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> * Enabling Oracle Managed Software Updates for a database automatically creates a Delegation Control for the VM cluster hosting the database. Oracle uses delegated access control to the database VMs to investigate and resolve eligible issues without requiring your intervention. Use the Notifications service to create your subscription topic and add it to the Delegation Control for the VM cluster and subscribe to the notification topic to receive Delegate Access Control related notifications.
> * To enable these notifications, use the Notifications service to create a subscription topic and add the required subscribers. Then, configure an OCI Events rule for Oracle Managed Database Software Updates events related to databases and database homes.
> * When a new Release Update (RU) is available, Oracle automatically schedules your maintenance based on your scheduling preferences.
> * Once you subscribe to Oracle Managed Database Software Updates related events, you will get notified at least one week in advance of scheduled maintenance.

For more information, see About Release Updates.

> **Note**
>
> * When Oracle Managed Database Software Updates option is selected, then Oracle will create, manage, and delete Oracle managed Database homes based on the requirement.
> * You cannot perform any operations on Oracle managed Database homes. These operations will be managed automatically.
> * Oracle automatically applies the latest RU for the Database release version.

- **Database Home**: *Applicable only if you have selected Customer managed software updates.* The Oracle Database home for the database. Choose the applicable options:
  - **Select an existing Database Home**: *Applicable only if you have selected Customer managed software updates.* The **Database Home display name** field allows you to choose the database home from the existing homes for the Oracle Database release that you specified. If no database home with that release exists, then you must create a new one.
  - **Create a new Database Home**: *Applicable only if you have selected Customer managed software updates.* Provide a Display name for the new database home, and either choose the default latest database image or click **Change database image** to select the database image from the available database images.
- **Database Storage**: Select the storage for your database.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> This section applies only to Oracle Database 19c container databases (CDBs) running on VM Clusters configured with Exascale Block Storage. There is no need to allocate any storage for Oracle AI Database CDBs running on VM Clusters configured with Exascale Smart Storage.

Storage vault: (read-only) Displays the Storage vault associated with the VM cluster.

Available vault storage capacity (GB): (read-only) The available GB of storage capacity in the vault is displayed

> **Note**
>
> If you don't have enough available vault storage capacity for database storage scale up operation, you need to first scale up the storage capacity of the storage vault as described in "Scaling Exascale Database Storage Vault" in Manage Exascale Database Vaults on Oracle Exadata Database Service on Exascale Infrastructure

Allocated storage capacity: If you are deploying Oracle Database 19c, then allocate storage:

- **Data storage (GB)** (`DATA`): Allocate storage size for data. The default size is 150 GB.

  The minimum size for Data storage is 56 GB.
- **Recovery area storage (GB)** (`RECO`): Allocate recovery file storage. The default is 50 GB.

  > **Note**
  >
  > The minimum size for the Recovery area is 44 GB for a 2-node cluster.
- **Log storage (GB)**: (read-only) The log storage size is calculated automatically based on the number of nodes in the VM cluster.

  > **Note**
  >
  > For a 2-node VM cluster, the Log storage (`LOG`) size is 33 GB.
- **Total allocated storage capacity (GB)**: (read-only) This field displays the total allocated storage for data files, the recovery area storage, and log storage.

- **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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 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, then the TDE wallet password is not applicable to your database 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, then 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.

- **Database Backups** Specify the settings for backing up the database to Autonomous Recovery Service or Object Storage:
  - **Enable automatic backup**: Select this option to enable automatic incremental backups for this database. Automatic backup is the preferred method for backing up Oracle Cloud databases. If you are creating a database in a security zone compartment, then you must enable automatic backups.
  - **Backup Destination**: Your choices are **Autonomous Recovery Service** or **Object Storage**.

    > **Note**
    >
    > Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

  - **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 data files 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**:
      - **Incremental backup scheduling time (UTC)**: Specify the time window when the incremental backups start when the automatic backup capability is selected.
    - **Autonomous Recovery Service**:

Copyright © 2023, 2026, Oracle and/or its affiliates.

    - **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 data files and the control file that constitute an Oracle Database. A full backup should also include any parameter files 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.

      > **⚠ Caution**
      >
      > If you choose to defer the first full backup, then your database may not be recoverable in the event of a database failure.

  - **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 according to the protection policy 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.
  - Scheduled time for daily backup
  - **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
  - **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.

- **Security**:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **Key Management**: Choose an encryption option: Oracle Wallet or OCI Vault. The default option is Oracle Wallet.
  - **Oracle Wallet**:
    Use the administrator password for the TDE wallet: When this option is checked, the password entered for the SYS user is also used for the TDE wallet. To set the TDE wallet password manually, uncheck this option and enter the TDE wallet password and confirm it.
  - **OCI Vault**: Select the compartment, and then select the Vault. Then, select the compartment where the Master encryption key is stored and choose the specific **Master encryption key**.
    Choose the key version: Enable this option to select the required key version.
- **Enable Data Safe**:

> **Note**
>
> Ensure that the required permissions are configured before enabling Data Safe. For more information, see General Steps for Creating an IAM Policy for Oracle Data Safe.

If you have the required permissions, the Enable Data Safe toggle is available. Turn on the toggle to enable Data Safe for the target database. If you do not have the required permissions, the toggle is disabled. You can skip enabling Data Safe while creating the database and enable it later. After the database is provisioned, go to the Database Information page. In the **Security** section:

- To enable Data Safe, click **Enable** next to **Data Safe**.
- To disable Data Safe, click the Actions icon (three dots), and then select **Disable**.
- To view the Data Safe configuration for the target database, click the Actions icon (three dots), and then select **View**.
  You are redirected to the target database details page in the **Oracle Data Safe** console.
- **Connectivity options to enable Data Safe**:
  Choose one of the following connectivity options to connect Oracle Data Safe to the target container database and its pluggable databases. For more information, see Connectivity Options for Target Databases.
  - **Data Safe Private Endpoint**:
    Use a Data Safe private endpoint to connect Oracle Data Safe to the target container database and its pluggable databases. You can associate only one private endpoint with the database at a time.
    - **Use an existing private endpoint**
      a. Select the compartment.
      b. Choose a private endpoint from the list of endpoints in the selected compartment.
    - **Create a new private endpoint**
      a. Click **Create Private Endpoint**.
      b. On the Create Private Endpoint page:

Copyright © 2023, 2026, Oracle and/or its affiliates.

i. Enter a descriptive name for the private endpoint.

ii. Select a compartment.

iii. In the Network Information section, select the **VCN** and **subnet**.

iv. Optionally, add tags.

v. Click **Create**.

* **Data Safe On-Premises Connector**:
Use one or more Data Safe on-premises connectors to connect Oracle Data Safe to the target container database and its pluggable databases.

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`.

  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**. When you select this option, Data is encrypted with an encryption key that Oracle maintains. You can also select **Use customer-managed keys**. Note that when you choose customer-managed keys, this option disables the TDE wallet password.

  To configure the database with encryption based on encryption keys you manage, complete the following:

  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**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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**.

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. After you successfully create the database, you can view the status of your update preferences in the Database Version section.

> **Note**
>
> You can always change your software management preference by clicking **More actions** and selecting **Manage database software update**. In the resulting page, select your software management preference and click **Save**.

> **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

To scale database storage capacity of a cluster database in the VM cluster

This topic describes how to scale storage capacity of container databases in the VM cluster configured with Exascale Block Storage.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Click the name of the VM cluster where the specific container database is located.
3. Click the **Container databases** tab.
4. On the Container Databases page, select the container database where you want to add storage.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5. Click the **Database information** tab.
6. On the **Database Information** pane, click the **Actions** drop-down list and select **Scale database storage**.

> **Note**
>
> This section applies only to Oracle Database 19c container databases (CDBs) running on VM Clusters configured with Exascale Block Storage. There is no need to allocate any storage for Oracle AI Database 26ai CDBs running on VM Clusters configured with Exascale Smart Storage.

7. In the **Scale database storage** dialog, provide the following:
   - **Storage vault**: The name of the storage vault associated with the VM cluster is displayed
   - **Available vault storage capacity (GB)**: (read-only) The available GB of storage capacity in the vault is displayed

     > **Note**
     >
     > If you don't have enough available vault storage capacity for database storage scale up operation, you need to first scale up the storage capacity of the storage vault as described in "Scaling Exascale Database Storage Vault" in Manage Exascale Database Vaults on Oracle Exadata Database Service on Exascale Infrastructure

   - **Data storage (GB)** (`DATA`): The existing data storage size is displayed. You can scale up or scale down the storage that you want to allocate to the database.

     > **Note**
     >
     > The minimum size for Data storage is 56 GB. It cannot be less than the usable space of the primary database.

   - **Recovery area storage size (GB)** The existing recovery area storage size is displayed. You can scale up or scale down the storage as needed.

     > **Note**
     >
     > The minimum size for the Recovery area is 44 GB for a 2-node cluster.

   - **Log storage (GB)** (read-only) The log storage size is calculated automatically based on the number of nodes in the VM cluster. For a 2-node VM cluster, the Log storage (LOG) size is 33 GB.
   - **Total Allocated storage capacity (GB)** (read-only) This field displays the total allocated storage for data files, the recovery area storage, and log storage.
8. Click **Scale**.

The Scale database storage dialog closes. On the container database page, the **Updating** status indicator shows that the storage capacity database is being scaled to the new allocation.

Copyright © 2023, 2026, Oracle and/or its affiliates.

When that process is complete, the **Available** status indicator is displayed.

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. Under **Oracle AI Database**, 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 AI Database**, then click **Oracle 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 Oracle Exadata Database Service on Exascale Infrastructure.

Before you begin, note the following:

- When you create a database from a backup, the availability domain can be the same as where the backup is hosted, or a different one across regions.
- The Oracle Database software version you specify must be the same or later version as that of the backed-up database.
- If you are creating a database from an automatic backup, then you can choose any level 0 weekly backup, or a level 1 incremental backup created after the most recent level 0 backup. For more information on automatic backups, see 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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to a backup.
   - *Standalone backups:* Click **Standalone backups** under **Resources**.
   - *Automatic backups:* Navigate to the Database Details page of the database associated with the backup:
     - Under **Oracle Exadata Database Service on Exascale Infrastructure**, click **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 (three dots) for the backup you chose.
5. Click **Create Database**.
6. On the **Create Database from Backup** page, configure the database as follows:
   - **PDB**: You can include all PDBs in the backup or choose specific PDBs to restore.
     - **All PDBs**: Select this option to include all PDBs in the backup.
     - **Choose the PDBs to restore**: Specify the PDBs to restore. You can select one or more PDBs. If you are selcting multiple PDBs, provide a comma-delimited list.
   - **Region**: The target region where you want to create the database.
   - **Availability domain**: It could be the same as the availability domain that hosts the backup or a different one within the same region.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **Select a service**: From the list, choose **Exadata Database Service on Exascale Infrastructure** to host the database.
- **VM Cluster**:
  - **VM Cluster**: Choose a cloud VM cluster to run the database from the chosen compartment.
  - **Software Management Preference**: Choose one of these options:
    - **Customer managed software updates**: This option is selected by default for all existing databases.
    - **Oracle managed software updates**: Select this option to enable Oracle-managed database home and software version updates.
      If you choose **Oracle managed software updates**:
      - **Maintenance method**: (read-only) By default, **Rolling** updates is selected. You cannot change this selection.
        Databases are updated in a rolling manner, one database instance at a time, with no downtime.
      - **Schedule Preferences**: Provide the day of the week and the maintenance window that you prefer for Oracle-managed database software updates.
      - **Version Preference**: Choose your preferred Release update (RU) for the Oracle-managed database updates. By default, Stay on the latest available Release Update is selected. Options are as follows:
        - Stay on the latest available Release Update
        - Stay one Release Update behind the latest
        - Stay two Release Updates behind the latest

> **Note**
>
> If you select options other than **Stay on the latest available Release Update**, the Confirm selection dialog opens. Read the acknowledgment, check the MOS Note box after reading the MOS Note, and click **Confirm**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> * Enabling Oracle Managed Software Updates for a database automatically creates a Delegation Control for the VM cluster hosting the database. Oracle uses delegated access control to the database VMs to investigate and resolve eligible issues without requiring your intervention. Use the Notifications service to create your subscription topic and add it to the Delegation Control for the VM cluster and subscribe to the notification topic to receive Delegate Access Control related notifications.
> * To enable these notifications, use the Notifications service to create a subscription topic and add the required subscribers. Then, configure an OCI Events rule for Oracle Managed Database Software Updates events related to databases and database homes.
> * When a new Release Update (RU) is available, Oracle automatically schedules your maintenance based on your scheduling preferences.
> * Once you subscribe to Oracle Managed Database Software Updates related events, you will get notified at least one week in advance of scheduled maintenance.

– **Database Home**: *Applicable only if you have selected Customer managed software updates*. Select an existing Database Home or create one as applicable. Note that this field is not available when you create a Database from the Database Home details page.

* **Select an existing Database Home**: *Applicable only if you have selected Customer managed software updates*. If you choose this option, make a selection from the **Select a Database Home** drop-down list. There is no version restriction when restoring to an existing Database Home. If the selected Database Home is running a Database Release Update earlier than the latest approved Release Update, the Console displays a message about the risks associated with using an older version.

  > **Note**
  >
  > You cannot create a database from backup in the same Database home where the source database exists.

* **Create a new Database home**: *Applicable only if you have selected Customer managed software updates* 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**.

- **Database Home**:

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> You cannot modify the `db_name`, `db_unique_name`, and SID prefix after creating the database.

- In the Database name field, name the database or accept the default name. The database name must meet the requirements:
  - Maximum of 8 characters
  - Contain only alphanumeric characters
  - Begin with an alphabetic character
  - Cannot be part of first 8 characters of a different database's db_unique_name on the VM cluster
  - DO NOT use the following reserved names: grid, ASM
- **Database unique name**: Specify a value for the `DB_UNIQUE_NAME` database parameter. The unique name must meet the following requirements:
  - Maximum of 30 characters
  - Contain only alphanumeric or underscore (_) characters
  - Begin with an alphabetic character
  - Unique across the VM cluster. Recommended to be unique across the tenancy.

  If not specified, the system automatically generates a unique name value, as follows:

  ```
  <db_name>_<3_chars_unique_string>_<region-name>
  ```

- **Administrator username**: This read-only field displays the username for the administrator, "sys".
- In the **Password** and **Confirm password** fields, enter and re-enter a password. A strong password for SYS administrator must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be _, #, or -. The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reverse order and regardless of casing.
- In the **Enter the source database's TDE wallet or RMAN password** field, enter a password that matches either the Transparent Data Encryption (TDE) wallet password or RMAN password for the source database.

7. 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 following requirements:
     - Maximum of 12 characters
     - Contain only alphanumeric characters

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Begin with an alphabetic character
- Unique in the VM cluster
- Must not use the following reserved names: grid, ASM

In the **Tags** tab, you can add tags to the database. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see Resource Tags. If you are not sure if you should apply tags, then skip this option (you can apply tags later) or ask your administrator.

8. Click **Create**.

To navigate to a list of backups for a particular database:

1. Click the Exadata cloud VM cluster name that contains the specific database to display the details page.
2. From the list of databases, click the database name associated with the backup that you want to use. Access the list of backups for the database by clicking the **Backups** tab.

To navigate to the list of backups for your current compartment:

1. Click **Standalone Backups** under **Resources**.
2. In the list of standalone backups, find the backup you want to use to create the database.

To Create a Database from the Last Backup

Use this procedure to create a database from the last 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 can be the same as where the backup is hosted, or a different one across regions.
- The Oracle Database software version you specify must be the same or later version as that of the backed-up database.
- If the backup being used to create a database is in a security zone compartment, the database cannot be created in a compartment that is not in a security zone. See the Security Zone Policies topic for a full list of policies that affect Database service resources.

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

   The Exadata VM Clusters page is selected by default.
2. Choose your **Compartment**.
3. In the list of VM clusters, navigate to the VM cluster that contains the source database you are using to create the new database:

   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 the **Database homes** tab.
5. Find the Database home that contains the source database and click its highlighted name.
6. Click the **Databases** tab and click the name of the database you are using as the source for the new database.
7. From the Actions menu, select **Create database from backup**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

8. On the resulting Create database from backup dialog, select **Create database from last backup** or **Create database from specified timestamp**.

   Select the **Create database from specified timestamp** option if you want to restore and recover the database to the specified timestamp. Enter the timestamp date and time details.

9. Click **Create**.

   The Create Database from Backup dialog opens.

10. On the Create Database from Backup page, configure the database as follows:

    - **PDB**: You can include all PDBs in the backup or choose specific PDBs to restore.
      - **All PDBs**: Select this option to include all PDBs in the backup.
      - **Choose the PDBs to restore**: Specify the PDBs to restore. You can select one or more PDBs. If you are selcting multiple PDBs, provide a comma-delimited list.
    - **Region**: The target region where you want to create the database.
    - **Availability domain**: It could be the same as the availability domain that hosts the backup or a different one within the same region.
    - **Select a service**: From the list, choose **Exadata Database Service on Exascale Infrastructure** to host the database.
    - **VM Cluster**:
      - **VM Cluster**: Choose a cloud VM cluster to run the database from the chosen compartment.
      - **Software Management Preference**: Choose one of these options:
        * **Customer managed software updates**: This option is selected by default for all existing databases.
        * **Oracle managed software updates**: Select this option to enable Oracle-managed database home and software version updates.
          If you choose **Oracle managed software updates**:
          * **Maintenance method**: (read-only) By default, **Rolling** updates is selected. You cannot change this selection.
            Databases are updated in a rolling manner, one database instance at a time, with no downtime.
          * **Schedule Preferences**: Provide the day of the week and the maintenance window that you prefer for Oracle-managed database software updates.
          * **Version Preference**: Choose your preferred Release update (RU) for the Oracle-managed database updates. By default, Stay on the latest available Release Update is selected. Options are as follows:
            * Stay on the latest available Release Update
            * Stay one Release Update behind the latest
            * Stay two Release Updates behind the latest

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you select options other than **Stay on the latest available Release Update**, the Confirm selection dialog opens. Read the acknowledgment, check the MOS Note box after reading the MOS Note, and click **Confirm**.

> **Note**
>
> * Enabling Oracle Managed Software Updates for a database automatically creates a Delegation Control for the VM cluster hosting the database. Oracle uses delegated access control to the database VMs to investigate and resolve eligible issues without requiring your intervention. Use the Notifications service to create your subscription topic and add it to the Delegation Control for the VM cluster and subscribe to the notification topic to receive Delegate Access Control related notifications.
> * To enable these notifications, use the Notifications service to create a subscription topic and add the required subscribers. Then, configure an OCI Events rule for Oracle Managed Database Software Updates events related to databases and database homes.
> * When a new Release Update (RU) is available, Oracle automatically schedules your maintenance based on your scheduling preferences.
> * Once you subscribe to Oracle Managed Database Software Updates related events, you will get notified at least one week in advance of scheduled maintenance.

- **Database Home**: *Applicable only if you have selected Customer managed software updates* Select an existing Database Home or create one as applicable. Note that this field is not available when you create a Database from the Database Home details page.
  - **Select an existing Database Home**: *Applicable only if you have selected Customer managed software updates* If one or more Database homes already exist for the database version you have selected, then this option is selected by default. You will be presented with a list of Database Homes. Select a Database Home from the list. There is no version restriction when restoring to an existing Database Home. If the selected Database Home is running a Database Release Update earlier than the latest approved Release Update, the Console displays a message about the risks associated with using an older version.
  - **Create a new Database Home**: *Applicable only if you have selected Customer managed software updates* If no Database homes exist for the database version you have selected, then this option is selected by default.
    - Enter **Database Home display name**.
    - Click **Change Database Image** to select your software version for the new Database home.
      Select a Database Software Image window is displayed.
    - Select **Image Type**, **Oracle Provided Database Software Images**, or **Custom Database Software Images**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **Database Home**:

  > **Note**
  >
  > You cannot modify the `db_name`, `db_unique_name`, and SID prefix after creating the database.

  - **Database name**: The name for the database. The database name must meet these 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 these 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 user name for the administrator, `sys`.
- In the **Password** and **Confirm password** fields, enter and re-enter a password.
  A strong password for SYS administrator must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be _, #, or -. The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reverse order and regardless of casing.
- In the **Enter the source database's TDE wallet or RMAN password** field, enter a password that matches either the Transparent Data Encryption (TDE) wallet password or RMAN password for the source database.

11. Click **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 these requirements:
  - Maximum of 12 characters

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Contain only alphanumeric characters
- Begin with an alphabetic character
- Unique in the VM cluster
- DO NOT use the following reserved names: grid, ASM

Tags: Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

12. Click **Create**.

Enable Oracle Managed Database Software Updates for an Existing Database

Learn how you can modify an existing database to support Oracle Managed Database Software Updates for Oracle Exadata Database Service on Exascale Infrastructure.

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

   The Exadata VM Clusters page is selected by default.
2. Choose your **Compartment** that contains the VM cluster that hosts the database.
3. Click the name of the VM cluster that contains the database.
4. On the Cloud VM Cluster details page, click the **Container Databases** tab.
5. Click the name of the container database that you want to modify.

   The Database Details page displays information about the selected database.
6. Click **More Actions** and select **Manage database software update**.
7. In the resulting window, choose **Oracle managed software updates**.
   a. **Maintenance method**: (read-only) By default, Rolling updates is selected. You cannot change this selection.

      Exadata Infrastructure is updated in a rolling manner, one database instance at a time, with no downtime.
   b. **Schedule Preferences**: Provide the day of the week and the maintenance window that you prefer for Oracle Managed Database Software Updates.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> - Enabling Oracle Managed Software Updates for a database automatically creates a Delegation Control for the VM cluster hosting the database. Oracle uses delegated access control to the database VMs to investigate and resolve eligible issues without requiring your intervention. Use the Notifications service to create your subscription topic and add it to the Delegation Control for the VM cluster and subscribe to the notification topic to receive Delegate Access Control related notifications.
> - To enable these notifications, use the Notifications service to create a subscription topic and add the required subscribers. Then, configure an OCI Events rule for Oracle Managed Database Software Updates events related to databases and database homes.
> - When a new Release Update (RU) is available, Oracle automatically schedules your maintenance based on your scheduling preferences.
> - Once you subscribe to Oracle Managed Database Software Updates related events, you will get notified at least one week in advance of scheduled maintenance.

For more information, see About Release Updates

c. **Version Preference**: Choose your preferred Release update (RU) for the Oracle-managed database updates. By default, Stay on the latest available Release Update is selected. Options are as follows:

- Stay on the latest available Release Update
- Stay one Release Update behind the latest
- Stay two Release Updates behind the latest

> **Note**
>
> If you select options other than **Stay on the latest available Release Update**, the Confirm selection dialog opens. Read the acknowledgment, check the MOS Note box after reading the MOS Note, and click **Confirm**.

d. Click **Save**.

On the next scheduled maintenance date, your database will be updated and moved to an Oracle managed Database home. After the scheduled update, you can view and access the new Database home in the **Database Version** section.

> **Note**
>
> If you disable Oracle managed software updates for a database, the modified database continues to be hosted by Oracle managed Database home and is unaffected by the scheduled maintenance.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle 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 wan 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 Databse**. 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 Reschedule Oracle Managed Database Software Updates for a Database

Learn how to reschedule the Oracle managed database software updates for Oracle Exadata Database Service on Exascale Infrastructure.

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

   The Exadata VM Clusters page is selected by default.
2. Choose your **Compartment** that contains the VM cluster that hosts the database.
3. Click the name of the VM cluster that contains the database.
4. On the Cloud VM Cluster details page, click the **Container Databases** tab.
5. Click the name of the container database that you want to modify.

   The Database Details page displays information about the selected database.
6. In the **Database Version** section, locate the Next Scheduled Maintenance field, and click **Reschedule**.
7. In the **Reschedule maintenance** window, review the read-only **Next Scheduled Maintenance** and **Target Database Software Version** fields.
8. Review the warning message.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> You can only reschedule the upcoming maintenance instance and not modify the maintenance day and time. Oracle will determine the next feasible maintenance date and time based on the schedule preference that you chose for the database. To view and edit the schedule preference of a database, see *Enable Oracle Managed Database Software Updates for an Existing Database*.

9. Click **Confirm**.
10. To view the maintenance reschedule request, click the **Work requests** tab.

Related Topics

- Enable Oracle Managed Database Software Updates for an Existing Database
  Learn how you can modify an existing database to support Oracle Managed Database Software Updates for Oracle Exadata Database Service on Exascale Infrastructure.

To View and Resolve Oracle Managed Database Software Updates Issues

Learn how to view and resolve the Oracle managed database software updates issues for Oracle Exadata Database Service on Exascale Infrastructure.

If a scheduled maintenance for a container database fails the readiness checks, a banner at the top of the Database panel provides a message indicating the issues that require your action.

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

   The Exadata VM Clusters page is selected by default.
2. Choose your **Compartment** that contains the VM cluster that hosts the database.
3. Click the name of the VM cluster that contains the database.
4. On the Cloud VM Cluster details page, click the **Container Databases** tab.
5. Click the name of the container database.

   The Database Details page displays information about the selected database.
6. In the Database panel, click **View details** found on the banner.
7. Address the issues by following the Fix Recommendation comments.

   To know more about these issues and their resolutions, see *Troubleshooting Oracle Managed Software Updates Issues*.
8. Optionally, click **Run Check** after resolving the listed issues to run the readiness checks manually. Confirm the action in the subsequent window.
9. Click **Close**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> - If you have not resolved the listed issues at least two hours before the scheduled maintenance time, the current schedule will be canceled and no further schedules will be created for that database until the issues are addressed. In the Database Details page, the status of the Next Scheduled Maintenance will be set as **Awaiting Customer Action**.
> - If an issue occurs during an Oracle managed software update, you will be notified with the status "Awaiting Resolution" and the message: "Oracle managed maintenance to apply <*version number*> database software version failed to complete successfully. Oracle operations is reviewing the root cause to address the failure and reschedule maintenance if required."
> - If the operations-related issues are still unresolved, the maintenance schedule will be temporarily canceled. You will be notified with the status "Awaiting Resolution" and the message: "Oracle is unable to schedule your next maintenance. Oracle operations is reviewing root cause to address the failure and reschedule maintenance."

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 AI Database**, then click **Oracle 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

> **Caution**
>
> You must not use the console to delete a database marked as "Failed" if you previously attempted to delete the database using the Database as a service command-line interface API (`dbaascli`), and afterward reuse the diskgroups in that database to create another database using the `dbaascli` API.
>
> If you delete the Failed status database from the Console after attempting to delete it with the command-line API, then this can result in detaching and deleting the underlying reused Exascale Direct Volumes (EDV), which can cause the failure of the newly created disk groups.

To administer Vault encryption keys

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

After you provision a database in cloud VM cluster, you can rotate the Vault encryption key or change the encryption management configuration for that database.

> **Note**
>
> - To ensure that your Exadata database uses the most current version of the Vault encryption key, rotate the key from the database details page on the Oracle Cloud Infrastructure Console. Do not use the Vault service.
> - You can rotate Vault encryption keys only on databases that are configured with customer-managed keys.
> - You can change encryption key management from Oracle-managed keys to customer-managed keys but you cannot change from customer-managed keys to Oracle-managed keys.
> - Oracle supports administering encryption keys on databases after Oracle Database 11*g* release 2 (11.2.0.4).

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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.

To Create and List Container Database Snapshot

Using the console, you can create and list the snapshots of the container database (CDB) in a VM Cluster.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM Clusters page is selected by default.
2. Choose your **Compartment**.
3. In the list of VM clusters, find the VM cluster that contains the specific CDB and click its highlighted name to view the details page for the cluster.
4. On the VM Cluster details page, click the **Container Databases** tab.
5. In the list of databases, find the CDB you want to access. Click the CDB name to display details about it.
6. To create a new snapshot, click **Actions** and select **Create Exascale CDB snapshot**.
   Enter a name for the CDB snapshot. Avoid entering confidential information.
7. Click **Create**.
   When the CDB snapshot creation is complete, the CDB state changes from Updating to Available.
8. To view the list of CDB snapshots:
   On the Container Database details page, select the **Exascale CDB snapshots** tab to view the list of snapshots created for the CDB.

To Create a Clone Database from a Snapshot

Using the console, you can create a clone database from a container database (CDB) snapshot.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM Clusters page is selected by default.

Copyright © 2023, 2026, Oracle and/or its affiliates.

2. Choose your **Compartment**.
3. In the list of VM clusters, find the VM cluster that contains the specific CDB and click its highlighted name to view the details page for the cluster.
4. On the VM Cluster details page, click the **Container Databases** tab.
5. In the list of databases, find the CDB you want to access. Click the CDB name to display details about it.
6. Select the **Exascale CDB snapshots** tab to view the list of CDB snapshots and locate the CDB snapshot you want to access.
7. Click the Action menu of the specific CDB snapshot.
8. Select the **Clone from snapshot** option.
   The Clone Database dialog opens.
9. On the Clone Database page, configure the database as follows:
   a. **Destination**:
      i. **Clone type**:
         i. **Local clone**: The CDB clone is created within the same VM Cluster.

> **Note**
>
> You cannot change the Compartment and VM Cluster values with this selection.

         ii. **Remote clone**: The CDB clone is created on a different VM Cluster.
            i. Modify Compartment and VM Clusters as needed.

> **Note**
>
> - You can choose either a thick or a thin clone if you select a VM Cluster from the same vault.
> - You can only create a thick clone if you select a VM Cluster from a different vault.

      ii. **Target Database name**: Enter the name for the database. The database name must meet the requirements:
         - Maximum of 8 characters
         - Contain only alphanumeric characters
      iii. **Unique target database name**: (Optional) Specify a unique value for the database. The value is case insensitive.
         If not specified, the system automatically generates a unique name value, as follows:

```
<db_name>_<3_chars_unique_string>_<region-name>
```

      iv. **Database version**: (Read-only) Database version for your database.
      v. **Database Home**: Select a Database home where you want the CDB snapshot to be cloned.

Copyright © 2023, 2026, Oracle and/or its affiliates.

vi. **Create administrator credentials**: A database administrator SYS user will be created with the password you supply.

   i. **Username**: SYS

   ii. **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 user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.

   iii. **Confirm password**: Re-enter the SYS password you specified.

   iv. Using a **TDE wallet password** is optional. If you are using customer-managed encryption keys stored in a Vault in your tenancy, then the TDE wallet password is not applicable to your database 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, then 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.

vii. **Database Backups**: Specify the settings for backing up the database to Autonomous Recovery Service or Object Storage:

   i. **Enable automatic backup**: Select this option to enable automatic incremental backups for this database. Automatic backup is the preferred method for backing up Oracle Cloud databases. If you are creating a database in a security zone compartment, then you must enable automatic backups.

   ii. **Backup Destination**: Your choices are **Autonomous Recovery Service** or **Object Storage**.

> **Note**
>
> Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

   - **Object Storage**:

      i. **Backup retention period**: Select one of 7, 15, 30, 45, or 60 days. Default: 30 days. The system automatically deletes your incremental backups at the end of your chosen retention period. You can edit the value after creating the clone.

      ii. **Scheduled day for full backup**: Choose the day for the backups to start.

      iii. **Scheduled time for full backup**: Choose the preferred time window.

      iv. **Scheduled time for incremental backup**: Choose the time window that matches or is later than the full backup schedule time.

   - **Autonomous Recovery Service**:

Copyright © 2023, 2026, Oracle and/or its affiliates.

i. Select your **Compartment**.

ii. **Protection policy**: These policies enable automated backup retention management for protected databases.

- **Bronze**: 14 days
- **Silver**: 35 days (selected by default)
- **Gold**: 65 days
- **Platinum**: 95 days

iii. **Backup location**: (read-only) Backups are stored in OCI.

iv. **Retention lock**: (read-only) Disabled.

v. **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.

vi. **Scheduled Time for Daily Backup**: Select the time window for the daily backup.

iii. **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 files 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 when a database fails.

iv. **Deletion options after database termination**:

- **Retain backups according to the protection policy 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.

b. **Source Database**:

i. **CDB name**: (read-only) Name of the database that you clone.

ii. **Source TDE wallet password**: Enter the TDE wallet password.

iii. **Enable thin clone**: Select this option to create a database clone using copy-on-write technology.

> **Note**
>
> This option is only available when you choose a VM Cluster from the same vault.

c. Expand **Advanced Options** to specify advanced options for the database:

i. **Management**:

i. **Oracle SID prefix**: This option is in the Management tab. The Oracle Database instance number is automatically added to the SID prefix to create

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 following requirements:

– Maximum of 12 characters

– Contain only alphanumeric characters

ii. **Character set**: The character set for the database. The default is AL32UTF8.

iii. **National character set**: The national character set for the database. The default is AL16UTF16.

ii. **Encryption**: By default, the database is configured using **Oracle-managed encryption keys**. You can also select **Use customer-managed keys**. Note that when you choose customer-managed keys, this option disables the TDE wallet password.

> **Note**
>
> If you are using your own encryption keys to protect the database, then encryption is mandatory for the clone database.

To configure the database with encryption based on encryption keys you manage, complete the following:

i. **Key management**: The field is auto-selected and read-only if you have already encrypted the source database using your own encryption keys.

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

ii. Choose the **Compartment** and **Vault**.

iii. Select a Master encryption key.

iv. 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.

iii. **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.

10. Click **Clone**.
A new clone database appears in the VM Cluster Databases resource list

To Create a Clone of an Existing Container Database

Using the console, you can create a clone of an existing container database.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
The VM Clusters page is selected by default.

2. Choose your **Compartment**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

3. In the list of VM clusters, find the VM cluster that contains the specific CDB and click its highlighted name to view the details page for the cluster.
4. On the VM Cluster details page, click the **Container Databases** tab.
5. In the list of databases, find the CDB you want to access. Click the Action menu of the specific CDB.
6. Select **Clone**.
   The Clone Database dialog opens.
7. On the Clone Database page, configure the database as follows:
   a. **Destination**:
      i. **Clone type**:
         i. **Local clone**: The CDB clone is created within the same VM Cluster.

> **Note**
>
> You cannot change the Compartment and VM Cluster values with this selection.

         ii. **Remote clone**: The CDB clone is created on a different VM Cluster.
            i. Modify **Compartment** and **VM Clusters** as needed.

> **Note**
>
> - You can choose either a thick or a thin clone if you select a VM Cluster from the same vault.
> - You can only create a thick clone if you select a VM Cluster from a different vault.

      ii. **Target Database name**: Enter the name for the database. The database name must meet the requirements:
         - Maximum of 8 characters
         - Contain only alphanumeric characters
      iii. **Unique target database name**: (Optional) Specify a unique value for the database. The value is case insensitive.
      If not specified, the system automatically generates a unique name value, as follows:

```
<db_name>_<3_chars_unique_string>_<region-name>
```

      iv. **Database version**: (Read-only) Database version for your database.
      v. **Database Home**: Select a Database home where you want the CDB to be cloned.
      vi. **Create administrator credentials**: A database administrator SYS user will be created with the password you supply.
         i. **Username**: SYS
         ii. **Password**: Supply the password for this user. The password must meet the following criteria:

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.

iii. **Confirm password**: Re-enter the SYS password you specified.

iv. Using a **TDE wallet password** is optional. If you are using customer-managed encryption keys stored in a Vault in your tenancy, then the TDE wallet password is not applicable to your database 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, then 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.

vii. **Database Backups**: Specify the settings for backing up the database to Autonomous Recovery Service or Object Storage:

i. **Enable automatic backup**: Select this option to enable automatic incremental backups for this database. Automatic backup is the preferred method for backing up Oracle Cloud databases. If you are creating a database in a security zone compartment, then you must enable automatic backups.

ii. **Backup Destination**: Your choices are **Autonomous Recovery Service** or **Object Storage**.

> **Note**
>
> Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

- **Object Storage**:

  i. **Backup retention period**: Select one of 7, 15, 30, 45, or 60 days. Default: 30 days. The system automatically deletes your incremental backups at the end of your chosen retention period. You can edit the value after creating the clone.

  ii. **Scheduled day for full backup**: Choose the day for the backups to start.

  iii. **Scheduled time for full backup**: Choose the preferred time window.

  iv. **Scheduled time for incremental backup**: Choose the time window that matches or is later than the full backup schedule time.

- **Autonomous Recovery Service**:

  i. Select your **Compartment**.

  ii. **Protection policy**: These policies enable automated backup retention management for protected databases.

  – **Bronze**: 14 days

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **Silver**: 35 days (selected by default)
- **Gold**: 65 days
- **Platinum**: 95 days

iii. **Backup location**: (read-only) Backups are stored in OCI.

iv. **Retention lock**: (read-only) Disabled.

v. **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.

vi. **Scheduled Time for Daily Backup**: Select the time window for the daily backup.

iii. **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 files 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 when a database fails.

iv. **Deletion options after database termination**:

- **Retain backups according to the protection policy 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.

b. **Source Database**:

i. **CDB name**: (read-only) Name of the database that you clone.

ii. **Source CDB admin password**: Enter the CDB admin password

iii. **Source TDE wallet password**: Enter the TDE wallet password.

iv. **Enable thin clone**: Select this option to create a database clone using copy-on-write technology.

> **Note**
>
> This option is only available when you choose a VM Cluster from the same vault.

c. Expand **Advanced Options** to specify advanced options for the database:

i. **Management**:

i. **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 following requirements:

Copyright © 2023, 2026, Oracle and/or its affiliates.

– Maximum of 12 characters

– Contain only alphanumeric characters

ii. **Character set**: The character set for the database. The default is AL32UTF8.

iii. **National character set**: The national character set for the database. The default is AL16UTF16.

ii. **Encryption**: If you are creating a database in a 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**. When you select this option, Data is encrypted with an encryption key that Oracle maintains. You can also select **Use customer-managed keys**. Note that when you choose customer-managed keys, this option disables the TDE wallet password.

> **Note**
>
> If you are using your own encryption keys to protect the database, then encryption is mandatory for the clone database.

To configure the database with encryption based on encryption keys you manage, complete the following:

i. **Key management**: The field is auto-selected and read-only if you have already encrypted the source database using your own encryption keys.

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

ii. Choose the **Compartment** and **Vault**.

iii. Select a Master encryption key.

iv. 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.

iii. **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.

8. Click **Clone**.
   When the CDB cloning is complete, the state changes from Creating to Available.

To Delete a Container Database Snapshot

Using the console, you can delete the snapshot of a container database (CDB) in a VM Cluster.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM Clusters page is selected by default.
2. Choose your **Compartment**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

3. In the list of VM clusters, find the VM cluster that contains the specific CDB and click its highlighted name to view the details page for the cluster.
4. On the VM Cluster details page, click the **Container Databases** tab.
5. In the list of databases, find the CDB you want to access. Click the CDB name to display details about it.
6. Select the **Exascale CDB snapshot** tab to view the list of CDB snapshots.
7. In the row corresponding to the snapshot that you want to delete, click the Action menu.
8. Click **Terminate**.
9. Read the message and acknowledge that you want to delete the CDB snapshot by entering the CDB Snapshot name and clicking **Terminate**.
   The snapshot briefly displays the Terminated status in the console before eventually disappearing from the snapshot list.

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**

Copyright © 2023, 2026, Oracle and/or its affiliates.

    - **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, 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, 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, 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.The Exadata VM Clusters page is selected by default.
2. 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.
3. On the VM cluster details page, click the **Container Databases** tab, and click the name of the database to display the Database information page.
4. Click the **Pluggable Databases** tab.
5. Click **Create**.
6. In the **Create Pluggable Database** dialog, enter the following:
   - **PDB Name**: Enter a name for the PDB. The name must begin with an alphabetic character and can contain a maximum of 30 alphanumeric characters.
   - **Unlock the 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **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.

- **Advanced Options**
  Encryption: Read-only. Contains the name of the Encryption Key.

  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 **Create**.

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 AI Database**, then click **Oracle 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **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';
    ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 Convert a Physical Standby Database to Snapshot Standby Database
  A snapshot standby database is a fully updateable standby database created by converting a physical standby database into a snapshot standby database.
- To Convert a Snapshot Standby Database to Physical Standby Database
  Oracle recommends converting a snapshot standby database back to a physical standby database within 14 days.
- To create a pluggable database snapshot
  To create a PDB snapshot, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure.
- To delete a pluggable database snapshot
  To delete an existing PDB snapshot, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure.
- To create a cloned pluggable database from a snapshot
  To create a cloned PDB from a PDB snapshot, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure.

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 AI Database**, then click **Oracle 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle 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 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 AI Database**, then click **Oracle 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.

   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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5. In the list of pluggable databases, find the pluggable database (PDB) you want to delete. Click the PDB name to display details about it.
6. Click **More Actions**, then choose **Delete**.
7. In the **Delete PDB** dialog box, enter the name of the PDB that you want to delete to confirm the action, then click **Delete**.

> **Note**
>
> You can now delete a PDB when Data Guard setup is running on another container database in the same Oracle home, and vice versa.

To get connection strings for a pluggable database

> **Note**
>
> This topic explains how to get connection strings for the administrative service of a PDB. Oracle recommends that you connect applications to an application service, using strings created for the application service.

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

   On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the PDB, and then click its name to display details about it.
6. Click **PDB Connection**.
7. In the **Pluggable Database Connection** dialog, use the **Show** and **Copy** links to display and copy connection strings, as needed.
8. Click **Close** to exit the dialog.

To Convert a Physical Standby Database to Snapshot Standby Database

A snapshot standby database is a fully updateable standby database created by converting a physical standby database into a snapshot standby database.

Snapshot standby allows administrators to perform root cause analysis and test with production data in a very simple and risk-free manner, thereby drastically reducing the risk of errors and improving resiliency. Until now, customers were able to convert a standby database

Copyright © 2023, 2026, Oracle and/or its affiliates.

into snapshot mode outside of cloud automation. Customers will now be able to create a snapshot standby database through OCI Console/API/SDK, and Terraform.

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

A snapshot standby database typically diverges from its primary database over time because redo data from the primary database is not applied as it is received. Local updates to the snapshot standby database cause additional divergence. However, the data in the primary database is fully protected because a snapshot standby can be converted back into a physical standby database at any time, and the redo data received from the primary is then applied.

A snapshot standby database provides disaster recovery and data protection benefits that are similar to those of a physical standby database. Snapshot standby databases are best used in scenarios where the benefit of having a temporary, updateable snapshot of the primary database justifies increased time to recover from primary database failures.

Note

- You will have to ensure there is sufficient space in the Fast Recovery Area. Oracle recommends converting a snapshot standby database back to a physical standby database in the shortest amount of time.
- Automatic backups are suspended on a snapshot standby. Archive logs from the production database will continue to be backed up to preserve recoverability.
- If you have performed a database upgrade while in the snapshot standby mode, you must ensure you use the correct Oracle Home when converting back to a physical standby database.

Before converting a physical standby database to a snapshot standby database, ensure the following conditions are met:

- Oracle recommends that you have more than one standby database in a Data Guard Group before converting a standby database to a snapshot standby database.
- Only one physical standby can be converted to a snapshot standby at a time. To convert multiple databases, perform the conversions sequentially.

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

   The VM Clusters page is selected by default.
2. Choose your **Compartment**.
3. Navigate to the VM cluster that contains the databases you wish to view their roles in Data Guard associations.
4. On the VM Cluster details page, click the **Container Databases** tab.

   The role of each container database in this VM Cluster is indicated in the Data Guard role column.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5. Choose a Standby database and select **Convert to snapshot standby** from the **More Actions** drop-down list.
6. Review the warning message on the Convert to snapshot standby window.
7. Click **Convert**.

To Convert a Snapshot Standby Database to Physical Standby Database

Oracle recommends converting a snapshot standby database back to a physical standby database within 14 days.

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

   The VM Clusters page is selected by default.
2. Choose your **Compartment**.
3. Navigate to the VM cluster that contains the databases you wish to view their roles in Data Guard associations.
4. On the VM Cluster details page, click the **Container Databases** tab.

   The role of each container database in this VM Cluster is indicated in the Data Guard role column.
5. Choose a Snapshot standby database and select **Convert to standby** from the **More Actions** drop-down list.

   > **Note**
   >
   > If you convert your snapshot standby to physical standby database, all local updates to your snapshot standby database will be discarded and data from your primary database will be applied.

6. Click **Convert**.
   - You cannot perform a switchover operation on a snapshot standby database. You will need to convert the snapshot standby to a standby database before initiating a switchover.
   - Oracle managed automatic backups will be disabled for a snapshot standby database. Automatic backups, if configured, will resume upon conversion from snapshot standby to standby database.
   - You can take a standalone backup of a snapshot standby database.
   - While a database in the Data Guard configuration is in snapshot standby mode, there is no support to create additional PDBs on the primary or any of the standby databases.

To create a pluggable database snapshot

To create a PDB snapshot, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure.

> **Note**
>
> This feature is available starting with Oracle AI Database 26ai and later versions.

Copyright © 2023, 2026, Oracle and/or its affiliates.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The Exadata VM Clusters page is selected by default.
2. 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.
3. On the Cloud VM Cluster details page, click the **Container Databases** tab, and click the name of the database to display the Database Details page.
4. Click the **Pluggable Databases** tab.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to access. Click the PDB name to display details about it.
6. Click the **Exascale PDB snapshot** tab to view the existing pluggable database snapshots or to create a new one.
7. Click **Create Exascale PDB snapshot**.
   Enter a name for the PDB snapshot. Avoid entering confidential information.
8. Click **Create**.
   When the PDB snapshot creation is complete, the state changes from Creating to Available.

To delete a pluggable database snapshot

To delete an existing PDB snapshot, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure.

> **Note**
>
> This feature is available starting with Oracle AI Database 26ai and later versions.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The Exadata VM Clusters page is selected by default.
2. 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.
3. On the Cloud VM Cluster details page, click the **Container Databases** tab, and click the name of the database to display the Database Details page.
4. Click the **Pluggable Databases** tab.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to access. Click the PDB name to display details about it.
6. Click the **Exascale PDB snapshot** tab to view the existing pluggable database snapshots or to create a new one.
7. Click the Actions button (three dots) for the PDB snapshot that you want to delete.
8. Click **Delete**.
9. In the resulting dialog:
   - Review the message about deleting the PDB snapshot.
   - Click **Delete** to confirm the action.

Copyright © 2023, 2026, Oracle and/or its affiliates.

To create a cloned pluggable database from a snapshot

To create a cloned PDB from a PDB snapshot, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure.

> **Note**
>
> - This feature is available starting with Oracle AI Database 26ai and later versions.
> - The PDB snapshot must be available to use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The Exadata VM Clusters page is selected by default.
2. 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.
3. On the Cloud VM Cluster details page, click the **Container Databases** tab, and click the name of the database to display the Database Details page.
4. Click the **Pluggable Databases** tab.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to access. Click the PDB name to display details about it.
6. Click the **Exascale PDB snapshot** tab to view the existing pluggable database snapshots or to create a new one.
7. Click the Actions button (three dots) for the PDB snapshot that you want to use to create a PDB clone.
8. Click **Clone from snapshot**.
9. In the resulting Clone from snapshot dialog, enter these details:
   a. **Destination Exadata VM Cluster in *Compartment name***: The VM Cluster name where the PDB is located.
   b. **Destination Database**: The container database name where the PDB is located.
   c. **Source Database Admin Password**: Enter the password assigned to the container database.
   d. **New PDB**: Enter the following details:
      i. **PDB Name**: Enter a name for the PDB. Avoid entering confidential information.
      ii. **Database TDE Wallet Password**: Enter the Transparent Data Encryption (TDE) wallet password.
      iii. **Unlock the PDB Admin Account**: Enable this toggle to unlock the PDB admin account.
   e. **Take a backup of the PDB immediately after cloning it**: Enable this toggle to create the PDB backup after cloning.

   > **Note**
   >
   > This toggle is dimmed if the option is not supported.

Copyright © 2023, 2026, Oracle and/or its affiliates.

10. Click **Clone**.
    When the PDB clone creation is complete, its state changes from Provisioning to Available. The container database containing this PDB clone also changes its state from Updating to Available once the PDB clone is available.

Cloning an Exadata Pluggable Database

You can create local, remote, and refreshable clones.

A clone is an independent and complete copy of the given database as it existed at the time of the cloning operation. You can create clones of your PDB within the same CDB or a different CDB and also refresh the cloned PDB.

> **Note**
>
> When cloning a PDB from 19c to 26ai, the cloned PDB is automatically upgraded to 26ai. For example, if you use refreshable clones to clone to 26ai and then convert it to regular PDB, all necessary upgrade steps are automatically handled, converting the refreshable clone into a fully upgraded 26ai PDB.

The following types of clones are supported:

- **Local clone:** A clone of the PDB is created within the same CDB.
- **Remote clone:** A clone of the PDB is created in a different CDB.
- **Refreshable clone:** A clone of the PDB is created in a different CDB, and you will be able to refresh the cloned PDB.
- 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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 **Container Databases** table, click the name of the database to display the Database Details page.

4. Click the **Pluggable Databases** tab.
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 **Actions** and select **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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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.

- *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.

> **Note**
>
> This option is available starting with Oracle AI Database 26ai and later versions.

- **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 AI Database**, then click **Oracle 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 **Container Databases** table, click the name of the database to display the Database Details page.
4. Click the **Pluggable Databases** tab.
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 **Actions** and select **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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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.

> **Note**
>
> This option is available starting with Oracle AI Database 26ai and later versions.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **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 AI Database**, then click **Oracle 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 page, in the **Container Databases** table, click the name of the database to display the Database Details page.
4. Click the **Pluggable Databases** tab.
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 **Actions** and select **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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **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.

> **Note**
>
> This option is available starting with Oracle AI Database 26ai and later versions.

- **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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Choose your Compartment.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 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 AI Database**, then click **Oracle 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 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

For the complete list of APIs for the Database service, see Database Service API.

To download the PDB Connection Bundle

The PDB Connection Bundle is HA-aware and provides connection information for the PDB on all Data Guard peers.

To download the Connection Bundle of a specific pluggable database, complete this procedure for Oracle Exadata Database Service on Exascale Infrastructure.

> **Note**
>
> You must add the `db-connection-bundles` policies to download the `DbConnectionBundle` file.

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

   The Exadata VM Clusters page is selected by default.
2. 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.
3. On the Cloud VM Cluster details page, click the **Container Databases** tab, and click the name of the database to display the Database Details page.
4. Click the **Pluggable Databases** tab.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to access. Click the PDB name to display details about it.
6. Click **More actions** and select **PDB connection**.
7. On the Pluggable Database Connection dialog, click **Download** located in the Connection Bundle section.
8. Supply a new password and re-enter the supplied password in the **Confirm password** field.
9. Click **Download**.

   > **Note**
   >
   > `DbConnectionBundle` can be downloaded only when it is in the **Active** state.

   Possible lifecycle states are: **Creating**, **Active**, **Inactive**, **Updating**, **Deleting**, **Deleted**, or **Failed**.

Related Topics

- How Policies Work
- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 Oracle 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 Oracle 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 Oracle Exadata Database Service on Exascale Infrastructure

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle 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 database passwords:

1. To change the SYS password for a database:
   a. Log in to an Oracle Exadata Database Service on Exascale Infrastructure virtual machine as the `opc` user.

   b. Run the following command:

   ```
   sudo dbaascli database changepassword --dbname <database_name> --user SYS
   ```
2. To change the SYS password in a Data Guard environment:

Copyright © 2023, 2026, Oracle and/or its affiliates.

a. On the primary database, run:

```
dbaascli database changePassword --dbName <dbname> --user SYS --
prepareStandbyBlob true --blobLocation <location_to_create_blob_file>
```

b. Copy the generated blob file to all standby databases and change the file ownership to the `oracle` user.

c. On each standby database, run:

```
dbaascli database changePassword --dbName <dbname> --user SYS --
standbyBlobFromPrimary <location_of_copied_blob_file>
```

3. To change the TDE wallet password for a database:
a. Log in to an Oracle Exadata Database Service on Exascale Infrastructure virtual machine as the `opc` user.

b. 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.
- Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console
- Long-Term Retention Backup with Recovery Service
  Long-term retention backup (LTR) allows you to store full backups for periods up to ten years for compliance, regulatory, or other business needs with complete LTR lifecycle management and immutability.
- Default Backup Channel Allocation
  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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Recovery Service (Recommended)

A fully managed service based on the on-premises Oracle's Zero Data Loss Recovery Appliance technology which offers modern cybersecurity protection for Oracle Databases. Unique, automated capabilities protect Oracle Database changes in real time, validate backups without production database overhead, and enable fast, predictable recovery to any point in time.

If your backups are currently configured with Object Storage, you can seamlessly transition to Recovery Service to achieve advanced capabilities with the same cost.

For more information on Recovery Service, see About Oracle Database Autonomous Recovery Service.

Object Storage

A secure, scalable, on-demand storage solution for databases.

> **Note**
>
> If you previously used `dbaascli` to configure backups and then you switch to using the Console or the API for backups:
>
> - A new backup configuration is created and associated with your database. This means that you can no longer rely on your previously configured unmanaged backups to protect your database.

Related Topics

- Managing Exadata Database Backups by Using dbaascli

Managed Backup Types and Usage Information

There are two types of automatic Exadata database backups: Autonomous Recovery Service, and Oracle Object Storage.

The database and infrastructure (the VM cluster) must be in an "Available" state for a backup operation to run successfully. Oracle recommends that you avoid performing actions that could interfere with availability (such as patching operations) while a backup operation is in progress. If an automatic backup operation fails, 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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*.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console

> **Note**
>
> - This change applies only to backups configured using the OCI Console.
> - Standalone (on-demand) backups are not affected unless otherwise stated in the related documentation.

Effective November 04, 2026

When you enable automatic backups using the OCI Console, Autonomous Recovery Service is the only available backup destination if all of the following conditions are met:

- The tenancy was created on or after November 04, 2026.
- The database is deployed in one of the following OCI regions: `US Gov East (Ashburn)` **`LFI`** , `US Gov West (Phoenix)` **`LUF`** , `US DoD North (Chicago)` **`PIA`** , `US DoD East (Ashburn)` **`RIC`** , `US DoD West (Phoenix)` **`TUS`** , `UK Gov West (Newport)` **`BRS`** , `UK Gov South (London)` **`LTN`** , `EU Sovereign Central (Frankfurt)` **`STR`** , and `EU Sovereign South (Madrid)` **`VLL`** .
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage is displayed as the available backup destination.

Effective June 23, 2026

When you enable automatic backups using the OCI Console, Autonomous Recovery Service is the only available backup destination if all of the following conditions are met:

- The tenancy was created on or after June 23, 2026.
- The database is deployed in one of the following OCI regions: `Netherlands Northwest (Amsterdam)` **`AMS`** , `Colombia Central (Bogota)` **`BOG`** , `Brazil East (Sao Paulo)` **`GRU`** , `Indonesia North (Batam)` **`HSG`** , `India South (Hyderabad)` **`HYD`** , `US East (Ashburn)` **`IAD`** , `South Korea Central (Seoul)` **`ICN`** , `Malaysia West 2 (Kulai)` **`JBP`** , `Saudi Arabia West (Jeddah)` **`JED`** , `Morocco West (Casablanca)` **`LEJ`** , `UK South (London)` **`LHR`** , `Italy Northwest (Milan)` **`LIN`** , `Spain Central (Madrid)` **`MAD`** , `Israel Central (Jerusalem)` **`MTZ`** , `Italy North (Turin)` **`NRQ`** , `Spain Central (Madrid 3)` **`ORF`** , `Chile Central (Santiago)` **`SCL`** , `Chile West (Valparaiso)` **`VAP`** , `Brazil South East (Vinhedo)` **`VCP`** , `Canada Southeast (Montreal)` **`YUL`** , and `Canada Southeast (Toronto)` **`YYZ`** .
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage is displayed as the available backup destination.

Effective May 15, 2026

When you enable automatic backups using the OCI Console, Autonomous Recovery Service is the only available backup destination if all of the following conditions are met:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- The tenancy was created on or after May 15, 2026.
- The database is deployed in one of the following OCI regions: `Singapore (Singapore)` **`SIN`** , `US West (San Jose)` **`SJC`** , `India West (Mumbai)` **`BOM`** , `France Central (Paris)` **`CDG`**, `Australia Southeast (Melbourne)` **`MEL`** , `Japan Central (Osaka)` **`KIX`** , `Australia East (Sydney)` **`SYD`** , `South Africa Central (Johannesburg)` **`JNB`** , `Mexico Central (Queretaro)` **`QRO`** , `US Midwest (Chicago)` **`ORD`** , `Sweden Central (Stockholm)` **`ARN`** , `Mexico Northeast (Monterrey)` **`MTY`** , `Switzerland North (Zurich)` **`ZRH`** , `South Korea North (Chuncheon)` **`YNY`** , `France South (Marseille)` **`MRS`** , `Singapore West (Singapore)` **`XSP`** , `UK West (Newport)` **`CWL`** , `UAE East (Dubai)` **`DXB`** , `UAE Central (Abu Dhabi)` **`AUH`** , and `Saudi Arabia Central (Riyadh)` **`RUH`** .
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage is displayed as the available backup destination.

Effective August 6, 2025

When you enable automatic backups using the OCI Console, Autonomous Recovery Service is the only available backup destination if all of the following conditions are met:

- The tenancy was created on or after August 6, 2025.
- The database is deployed in one of the following OCI regions: `Germany Central (Frankfurt)` **`FRA`** , `US West (Phoenix)` **`PHX`** , and `Japan East (Tokyo)` **`NRT`** .
- The Oracle Database version is later than 19.18 or 23.4.

If any of these conditions are not met, OCI Object Storage is displayed as the available backup destination.

Long-Term Retention Backup with Recovery Service

Long-term retention backup (LTR) allows you to store full backups for periods up to ten years for compliance, regulatory, or other business needs with complete LTR lifecycle management and immutability.

For LTR with Recovery Service, the retention period must be in Days (90 - 3,650) or years (1 - 10) from when the backup was created.

To create an LTR backup with the required retention period, Recovery Service does not require creating a new full production backup but does so by utilizing already existing operational backups in the system within the defined recovery window in the policy. For more information, see To create an on-demand backup of a database.

You can change the retention period for a specific existing LTR backup within the retention period. For more information, see To change the retention period of an LTR backup with Recovery Service.

You can restore an LTR backup to create a new database within the retention period. For more information, see To Create a Database from a Backup.

Upon terminating a database, the LTR backups will be deleted as per the 'Deletion options after database termination' value.

- **Delete backups in 72 hours**: All backups, including long-term backups, will be deleted.
- **Delete based on policy**: LTR backups will be retained according to the retention policy of each LTR backup.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Note: Oracle recommends choosing '**Delete based on policy**' option while terminating a database to ensure the long-term backups are retained.

Consider the following additional factors for long-term backups:

- LTR backups will continue to exist independently of any automatic backups configured on the database.
- **LTR backups will be automatically deleted after the specified retention period ends.**
- In-place restore is not supported for LTR.
- For databases in a Data Guard configuration, the long-term backup will be created only for the database where it is requested.
- The database must be in the AVAILABLE state to create a LTR.
- LTR is supported for databases with file-based TDE or KMS-based keystores.
- Encryption keys will be maintained for the entire retention period of the LTR.
- An LTR backup can be canceled while it is in the 'creating' state.
- An LTR backup can be deleted at any time after it is created.
- During restore:
  - If the backup is of a supported DBHome major version, it will be restored to the latest RU of that version.
  - If the backup is of an unsupported DBHome major version, it will be restored to a supported major version, after which the database must be upgraded to any of the supported major versions.

Default Backup Channel Allocation

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 `bckup cfg` file, and setting the parameter `bkup_channels_node` to the number of channels per node desired.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Task	More Information	Required or Optional
Create IAM policies	Permissions Required for Oracle Databases in OCI to Use Recovery Service	Required
Configure network resources and register a Recovery Service subnet	Create a Recovery Service Subnet in the Database VCN	Required
Create protection policies	Creating a Protection Policy	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*
  - Review 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>
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

For more information about adding a user to a group, see Managing Groups. For more information about policies, see Getting Started with Policies.

Related Topics

- Auth Token

Using the Console to Manage Backups

You can use the Console to enable automatic incremental backups, create full backups on demand, and view the list of managed backups for a database. You can also use the Console to delete manual (on-demand) backups.

> **Note**
>
> - All backups are encrypted with the same master key used for Transparent Data Encryption (TDE) wallet encryption.
> - Backups for a particular database are listed on the details page for that database. The Encryption Key column displays either Oracle-Managed Key or a key name if you are using your own encryption keys to protect the database. See Backing Up Vaults and Keys for more information.

> **Note**
>
> Do not delete any necessary encryption keys from the vault because this causes databases and backups protected by the key to become unavailable.

- To configure automatic backups for a database
- To create an on-demand backup of a database
- To view backup status
- To cancel a backup
- To delete full backups from Object Storage
- To delete standalone backups from Object Storage
- To change the retention period of an LTR backup with Recovery Service

To configure automatic backups for a database

> **Note**
>
> Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> Databases in a *security zone compartment* must have automatic backups enabled. See the *Security Zone Policies* topic for a full list of policies that affect Database service resources.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database that 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 that you want to access and click its highlighted name to view the details page for the cluster.
4. In the list of databases, find the database for which you want to create an on-demand full backup and click its name to display database details.
5. Under **Resources**, click **Backups**.

   A list of backups is displayed.
6. Click **Create Backup**.
7. On the resulting Create backup window, do the following:
   - **Name**: Provide a descriptive name for the backup.
   - Select a **Backup retention** option:
     - **Retain backups per backup retention period**: Select this option to use the protection policy retention period for this backup.
     - **Specify long-term backup retention period**: Select this option to specify an LTR period with Autonomous Recovery Service. The retention period must be entered in Days (90 - 3,650) or Years (1 - 10) from when the backup was created.
   - Click **Create**.

To view backup status

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

4. Click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
5. In the list of databases, find the database you are interested in and click its name to display database details.
6. Under **Resources**, click **Backups**.
   A list of backups is displayed. The state column displays the status of the backup: **Active**, **Creating**, **Canceled**, **Canceling**, or **Failed**.

To cancel a backup

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on 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**.
7. A backup in the Creating state may be canceled by clicking the Actions icon (three dots) on the right of the backup row and clicking **Cancel Backup**.
   A Cancel Backup confirmation dialog will appear.
8. Enter the name of the backup, and click **Cancel Backup**.
   The state changes to **Canceling**.

   The Cancel backup Work request can be viewed, by clicking **Work requests** under **Resources**.

If the Cancel backup fails:

- In the Work requests pane under Resources, you will see a line item called "**Cancel Database Backup**" with a state of "**Failed**". There will also be a work request for the backup "**Create Database Backup**" that will reflect the state of the Backup operation.

To delete full backups from Object Storage

> **Note**
>
> You cannot explicitly delete automatic backups. Unless you terminate the database, automatic backups remain in Recovery Service and Object Storage for the number of days specified by the user, after which time they are automatically deleted.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Standalone Backups** under **Resources**.
2. In the list of standalone backups, find the backup you want to use to delete.
3. Click the Actions menu for the backup you are interested in, and then click **Delete**.
4. In the **Delete** dialog, click **Delete** to confirm the backup deletion.

To change the retention period of an LTR backup with Recovery Service

1. Open the navigation menu. Select **Oracle AI Database,** then select Oracle Exadata Database Service on Exascale Infrastructure.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database you want to change the backup retention period:
   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, click the name of the database for which you want to change the retention period.
5. Under **Resources**, click **Backups**.
   A list of backups is displayed.
6. In the list of backups, click the **Actions** menu for the backup with type **Long-term backup** for which you want to change the retention period.
7. Click **Change retention period**.
8. In the resulting **Change retention period**, change the retention period.

> **Note**
>
> The retention period must be entered in Days (90 - 3,650) or Years (1 - 10) from when the backup was created.

Copyright © 2023, 2026, Oracle and/or its affiliates.

9. Click **Save**.

To designate Autonomous Recovery Service as a Backup Destination for an Existing Database

To designate Autonomous Recovery Service as a Backup Destination for an existing database, use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the database:
   Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model): Under **Exadata on Oracle Public Cloud**, click **Exadata VM Clusters**.

   In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

   On the cloud **VM cluster** details page, in the Databases table, click the name of the database to display the **Database Details** page.
4. Click **Configure automatic backups**.
5. In the resulting window, provide the following details:
   - **Enable automatic backup**: Check the check box to enable automatic incremental backups for this database. If you are creating a database in a security zone compartment, you must enable automatic backups.
   - **Backup Destination**: Select **Autonomous Recovery Service**.
   - **Backup Scheduling**: If you enable automatic backups, you can choose a two-hour scheduling window to control when backup operations begin. If you do not specify a window, then a six-hour default window of 00:00 to 06:00 (in the time zone of the VM cluster's region) is used for your database.
   - **Protection Policy**: If you choose to enable automatic backups, you can choose a policy with one of the following preset retention periods, or a Custom policy.

     Object Storage Backup retention period: 7, 15, 30, 45, 60. Default: 30. The system automatically deletes your incremental backups at the end of your chosen retention period.

     Autonomous Recovery Service protection policy:

     – **Bronze:** 14 days
     – **Silver:** 35 days
     – **Gold:** 65 days
     – **Platinum:** 95 days
     – Custom defined by you
     – **Default:** Silver - 35 days
   - **Enable Real-Time Data Protection**: Real-time protection is the continuous transfer of redo changes from a protected database to **Autonomous Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.
6. Click **Save**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Recovering an Exadata Database from Backup Destination

This topic explains how to recover an Exadata database from a backup stored in either Object Storage or Autonomous Recovery Service by using the Console or the API.

- Object Storage service is a secure, scalable, on-demand storage solution in Exadata Cloud Infrastructure.
- OracleDatabase Autonomous Recovery Service is a centralized, fully managed, and standalone backup solution for Oracle Cloud Infrastructure (OCI) databases.

For more information about backing up your databases to Object Storage, see *Managing Exadata Database Backups*.

- Using the Console to restore a database
  You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.

Related Topics

- Managing Exadata Database Backups
  Automatic Exadata database backups are managed by Oracle Cloud Infrastructure. You configure this by using the Console or the API.

Using the Console to restore a database

You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.

> **Note**
>
> LTR backups represent a single point in time for the database, so the following options are not supported when restoring.

You can restore to:

- **Restore to latest**: Restores the database to the last known good state with the least possible data loss.
- **Restore to a timestamp**: Restores the database to the timestamp specified.
- **Restore to SCN**: Restores the database using the SCN specified. This SCN must be valid.

> **Note**
>
> You can determine the SCN number to use either by accessing and querying your database host, or by accessing any online or archived logs.

> **Note**
>
> The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `dbaascli`).

Copyright © 2023, 2026, Oracle and/or its affiliates.

- To restore a database
  Use this procedure to restore your Oracle Exadata Database Service on Exascale Infrastructure Oracle Database.

To restore a database

Use this procedure to restore your Oracle Exadata Database Service on Exascale Infrastructure Oracle Database.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
2. Click **Exadata VM Clusters**.
3. In the list of VM clusters, find the VM cluster containg the database that you want to restore, and click its highlighted name to view the details page for the cluster.
4. In the list of databases, find the database that 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.
- 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 of a database deployment on the Oracle Exadata Database Service on Exascale Infrastructure instance, use the `bkup_api` utility.
- 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. 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. Use SSH to connect to one of the database configured nodes in the VM cluster.

   ```
   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 -
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

> **Note**
>
> Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.
>
> About Using Recovery Service to Backup and Recover Oracle Cloud Databases

Parameter	Description
`bkup_disk=[yes|no]`	Whether to back up locally to disk (Fast Recovery Area).
`bkup_oss=[yes|no]`	Whether to back up to Object Storage. If yes, you must also provide the parameters `bkup_oss_url`, `bkup_oss_user`, `bkup_oss_passwd`, and `bkup_oss_recovery_window`.
`bkup_oss_url=<swift_url>`	Required if `bkup_oss=yes`. The Object Storage URL including the tenant and bucket you want to use. The URL is: `https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<tenant>/<bucket>` Where: • `<tenant>` - lowercase tenant name (even if it contains uppercase characters) that you specify when signing in to the Console • `<bucket>` - name of the existing bucket you want to use for backups.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Parameter	Description
`bkup_oss_user=<oci_user_name>`	Required if `bkup_oss=yes`. The user name for the Oracle Cloud Infrastructure user account. This is the user name you use to sign in to the Oracle Cloud Infrastructure Console. For example, **jsmith@example.com** for a local user or `<identity_provider>/jsmith@example.com` for a federated user. To determine which type of user you have, see the following topics: • Managing Users (for information on local users) • Federating with Identity Providers (for information on federated users) Note that the user must be a member of the Administrators group.
`bkup_oss_passwd=<auth_token>`	Required if `bkup_oss=yes`. The **auth token** generated by using the Console or IAM API, as described in **Prerequisites**. This is not the password for the Oracle Cloud Infrastructure user.
`bkup_oss_recovery_window=n`	Required if `bkup_oss=yes`. The number of days for which backups and archived redo logs are maintained in the Object Storage bucket. Specify 7 to 90 days.
`bkup_daily_time=hh:mm`	The time at which the daily backup is scheduled, specified in hours and minutes (`hh:mm`), in 24-hour format.
`bkup_archlog_cron_entry=[yes|no]`	When no backups are configured using `dbaastools`, setting `bkup_archlog_cron_entry=no` will remove the archive log clean up job from **crontab**. The default value is "**yes**".

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>
```

6. Use the `dbaascli database backup --status` to check the status of UUID generated for this command.

```
# dbaascli database backup --status --uuid <uuid> --dbname <database_name>
```

> **Note**
>
> A backup configuration file can contain the credentials to access the Object Storage bucket. For this reason, you might want to remove the file after successfully configuring the backup.

The following parameters can be modified to customize the backup configuration:

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> `Compatible with Console Automatic Backups=Yes` indicates the parameter is safe to change, even when using console-based automatic backups. If using parameters with `Compatible with Console Automatic Backups=No`, then do not enable backups through the console.

Table 5-6 Backup Configuration Parameters - Schedule Parameters to dbaascli

Parameter	Description	Compatible with Console Automatic Backups*
Old name: `bkup_cron_entry` New name: `scheduleBackups`	Enables the automatic backup configuration. Valid values are `yes` and `no`.	No
Old name: `bkup_archlog_cron_entry` New name: `manageArchivelogs`	Enables automatic backups of archived database log files. Valid values are `yes` and `no`. Setting `manageArchivelogs` to `no` disables automatic archive log clean-up jobs. This setting is valid only when the associated database has no automatic database backups configured.	No
Old name: `bkup_l0_day` New name: `L0BackupDay`	This parameter controls the Level 0 day of the week. Day of the week when a level 0 backup is taken. Valid values are `mon`, `tue`, `wed`, `thu`, `fri`, `sat`, and `sun`. Longer formats, for example, `Monday`, `Tuesday` are also supported. Default: `sun`.	No

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-7 Backup Configuration Parameters - General RMAN Configuration Parameters (valid for all backup destinations except Local Storage (FRA))

Parameter	Description	Compatible with Console Automatic Backups*
Old name: `bkup_rman_compression` New name: `compressionLevel`	Level of compression applied to automatic backups. Valid values for Object Storage are `NONE`, `basic`, `low`, `medium`, and `high`. Valid values for Recovery Service are `NONE` and `low`. Default value is `low`. A value of `NONE` disables RMAN compression. If RMAN compression is enabled, then any TDE encrypted datafile will be decrypted, compressed, and RMAN encrypted.	Yes
Old name: `bkup_section_size` New name: `sectionSize`	RMAN section size that is used for automatic backups. Default value is 64G.	Yes
Old name: `bkup_channels_node` New name: `channelsPerNode`	Number of RMAN channels per node used for automatic backups. Valid values are between 1 and 32. Default value is 2.	Yes
Old name: `bkup_daily_time` New name: `autoBackupTime`	Start time of the automatic daily backup expressed in 24-hour time as `hh:mm`.	Yes
Old name: `bkup_archlog_frequency` New Name: `backupFrequencyAL`	Interval in minutes between automatic backups of archived database log files. Valid values are 15, 20, 30, 60, 120 through 1440 in one-hour intervals expressed in minutes. Default value is 30 for ExaDB-D.	Yes
Old name: `bkup_type` New name: `backupDestination`	The type of the location where the backup resides. Specify OSS as the backup destination, which is the default and only option.	Yes
Old name: `bkup_filesperset_regular` New name: `filesPerSet`	Specifies the maximum number of data files that can be included in a backup set for Regular/ Archival backups.	Yes
Old name: `bkup_filesperset_al` New name: `filesPerSetAL`	Specifies the maximum number of archive log files that can be included in a backup set for Archivelog Backups.	Yes

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-7 (Cont.) Backup Configuration Parameters - General RMAN Configuration Parameters (valid for all backup destinations except Local Storage (FRA))

Parameter	Description	Compatible with Console Automatic Backups*
Old name: `bkup_encryption` New name: `encryption`	Encryption specifies whether backups should be encrypted or not. By default, encryption is enabled for OSS and Recovery Service, and this setting cannot be changed.	Yes
Old name: `rmanBackupOptimization` New name: `optimization`	Optimization is a feature in that reduces the amount of data that needs to be backed up, transferred, and restored. Recommended value is ON.	Yes
Old name: `rmanFraCleanupChannels` New name: `numberOfChannelsForFraCleanup`	Specifies the number of channels used for FRA Cleanup job.	Yes
Old name: `Compress_Archive_Logs` New name: `compressionAL`	Specifies whether to compress the archive log backups are not. Not applicable to Recovery Service.	Yes
Old name: `bkup_archlog_fra_retention` New name: `archivelogRetentionDays`	Specifies the number of days archive log to be retained in FRA.	Yes

Table 5-8 Backup Configuration Parameters - Object Storage Service (OSS) Parameters

Parameter	Description	Compatible with Console Automatic Backups*
`backupDestination=oss`	Enables backups to cloud storage. Valid values are `yes` and `no`.	No
Old name: `bkup_oss_recovery_window` New name: `ossRecoveryWindow`	Retention period for backups to cloud storage, expressed as a number of days up to 90. Applicable only when `bkup_oss` is set to `yes` or `backupdestination` is set to `OSS`. Default value is 30.	No

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-8 (Cont.) Backup Configuration Parameters - Object Storage Service (OSS) Parameters

Parameter	Description	Compatible with Console Automatic Backups*
Old name: `bkup_oss_url` New name: `ossURL`	Location of the storage container that is used for backup to cloud storage. Applicable only when `bkup_oss` is set to `yes` or `backupdestination` is set to `OSS`.	No
Old name: `bkup_oss_user` New name: `ossUserName`	User name of the Oracle Cloud user having write privileges on the cloud storage container specified in `bkup_oss_url`. Applicable only when `bkup_oss` is set to `yes` or `backupdestination` is set to `OSS`.	No
Old name: `bkup_oss_passwd` New name: `ossAuthToken`	Password of the Oracle Cloud user having write privileges on the cloud storage container specified in `bkup_oss_url`. Applicable only when `bkup_oss` is set to `yes` or `backupdestination` is set to `OSS`.	No

Table 5-9 Backup Configuration Parameters - RMAN Catalog Support Parameters

Parameter	Description	Compatible with Console Automatic Backups*
Old name: `bkup_use_rcat` New name: `useCatalog`	Enables the use of an existing RMAN recovery catalog. Valid values are `yes` and `no`.	Yes
Old name: `bkup_rcat_user` New name: `catalogUserName`	Recovery catalog user name. Applicable only when `bkup_use_rcat` is set to `yes`.	Yes
Old name: `bkup_rcat_passwd` New name: `catalogPassword`	Password for recovery catalog user specified in `bkup_rcat_user` . Applicable only when `bkup_use_rcat` is set to `yes`.	Yes
Old name: `bkup_rcat_conn` New name: `catalogConnectionString`	Connection string for the RMAN recovery catalog. Applicable only when `bkup_use_rcat` is set to `yes`.	Yes

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> Only the above parameters noted with `Compatible with Console Automatic Backups = Yes` are safe to alter in conjunction with console-based automatic backups. If any other parameters are to be altered, then do not enable backups through the console.

To create an on-demand backup

You can use the `dbaascli` to create an on-demand backup of a database.

1. SSH to one of the database configured nodes in the VM cluster resource.

   ```
   ssh -i <private_key_path> opc@<node_1_ip_address>
   ```

   To determine the first compute node, connect to any compute node as the `grid` user and execute the following command:

   ```
   $ $ORACLE_HOME/bin/olsnodes -n
   ```

   The first node has the number 1 listed beside the node name.

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

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

3. You can let the backup follow the current retention policy, or you can create a long-term backup that persists until you delete it:

   - To create a backup that follows the current retention policy, enter the following command:

     ```
     # dbaascli database backup --start --dbname <database_name>
     ```

   - To create a long-term backup, enter the following command:

     ```
     # dbaascli database backup --start --archival --dbname --tag
     <archival_tag>
     ```

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

   ```
   # exit
   $ exit
   ```

5. Use the `dbaascli database backup --status` to check the status of UUID generated for the backup command

   ```
   # dbaascli database backup --status --uuid <uuid> --dbname <database_name>
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

Related Topics

- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To remove the backup configuration

1. SSH to one of the database configured nodes in the VM cluster resource.
2. Log in as `opc` and then `sudo` to the `root` user.
3. Create a `temp` file with following parameters:
   - `bkup_oss=no`
   - `bkup_cron_entry=no`
   - `bkup_archlog_cron_entry=no`
4. Use the above file with `dbaascli database backup --configure` to remove the backup configuration for a database.

   ```
   # dbaascli database backup --configure --configFile <file_name> --dbname
   <database_name>
   ```

5. Use the `dbaascli database backup --status` to check the status of UUID generated for this command.

   ```
   # dbaascli database backup --status --uuid <uuid> --dbname <database_name>
   ```

This will disable all automatic backups.

Related Topics

- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To delete a local backup

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

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

1. Connect to the first compute node in your Exadata VM cluster 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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*.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 maintanance.
- 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 maintanance.

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.
- Applying Out-of-Cycle Kernel and Individual Package Updates

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 a 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 Oracle 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 InfrastructureOracle Exadata Database Service on Exascale Infrastructure VM clusters apply patches, and monitor the status of patch operations.
- 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 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 patch updates the Oracle Grid Infrastructure (GI) on the resource.

By default, you can update a Grid Infrastructure home to the latest approved Grid Infrastructure Release Update. You can also select either of the two preceding Release Updates after acknowledging the associated security risks. To update to an earlier Release Update, submit a Service Request.

> **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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

By default, you can patch a Database Home to the latest approved Database Release Update (N). You can also patch to either of the two preceding Release Updates (N-1 or N-2) after acknowledging the associated security risks. To patch to a Release Update older than N-2, submit a Service Request.

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 before you patch the Databases Homes and databases on that resource.
- Before you apply any patch, run the precheck operation to ensure your VM cluster or Database Home meets the requirements for that patch.
- To patch a database to a version other than the database version of the current home, move the database to a Database Home running the target version. This technique requires less downtime, and enables you to easily roll back the database to the previous version by moving it back to the old Database Home. See To Move a Database to Another Database Home.
- For Oracle Database and Oracle Grid Infrastructure major version releases available in Oracle Cloud Infrastructure, the latest approved Release Update is available by default. You can also select either of the two preceding Release Updates after acknowledging the associated security risks. To use an earlier Release Update, submit a Service Request.
- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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]
[--skipPdbs | --pdbs]
[--executePrereqs]
[--patchList]
[--skipClosedPdbs]
[--rollback]
```

Where:

- `--dbname` specifies the name of the database
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--skipPdbs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: *pdb1*,*pdb2*...
- `--pdbs` 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*...
- `--skipClosedPdbs` 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 `--skipPdbs` option followed by a comma-delimited list of PDB names to skip. For example:

```
--skipPdbs pdb1,pdb2
```

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

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

```
--pdbs 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 `--skipClosedPdbs` 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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

```
--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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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] [--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
- `--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
```

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 --dbbname dbname --targetHome Target Oracle home
location --executePrereqs
```

Prerequisites for Patching and Updating a 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Required Network Setup for Patching

Review the security rules to use with your ExaDB-XS system. Security rules control the types of traffic allowed for the client network and backup network of the Exadata's compute nodes.

Using the Console to Patch and Update Oracle 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 InfrastructureOracle 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 Oracle 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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**
2. Choose your **Compartment**.
3. Click **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 Grid Infrastructure operation.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle 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 wan to move.
4. Click **More Actions**, and then click **Move to Another Home**.
5. Select the target Database Home.
6. Click **Move Database**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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) versoins. 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Viewing Patch History of Oracle 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 AI Database**, then click **Oracle 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. 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 AI Database**, then click **Oracle 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 **Oracle 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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**.

Applying Out-of-Cycle Kernel and Individual Package Updates

Oracle recommends applying all kernel and package updates to your Exadata Database Service guest VM as part of a full Exadata software release. Exadata software is released monthly, usually on the third-Tuesday of the month. Because of the monthly cadence and time to package into a tested software release, there may be a delay from the time a critical update such as a high CVSS score security fix is made publicly available and included in the monthly Exadata software release. In the event a critical update becomes available, customers can apply it manually to their guest VMs rather than waiting for the monthly release by following the directions in the Exadata documentation linked below:

- Updating Individual Packages on Exadata Database Servers
- Installing Kernel Updates on Exadata Database Servers
- Updating Individual Packages on Exadata Storage Servers
- Installing Kernel Updates on Oracle Exadata Storage Servers

Interim Software Updates

For authorized environments, learn how to download interim software updates.

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

> **Note**
>
> To download interim software updates, you should at least have Oracle Exadata Database Service on Exascale 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Create an Interim Software Update
- Download an Interim Software Update
- Delete an Interim Software Update
- Move an Interim Software Update Resource to Another Compartment
- Using the API to Manage Interim Software Updates
  The Oracle Cloud Infrastructure APIs are typical REST APIs that use HTTPS requests and responses, and enable you ot perform interim software updates.

Create an Interim Software Update

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Under **Resources**, click **Interim software updates**.

   Interim software update page is displayed.
3. Click **Create interim software update**.

   Create interim software update panel is displayed.
4. Enter the following details in the panel:
   a. **Name**: Descriptive name for the patch download path.
   b. **Compartment**: Select a compartment where you want to create the patch resource.
   c. **Database version**: Choose the Database version for your image.
   d. **Release Update**: Choose any supported Oracle Database release update (RU).
   e. **Interim software update number**: Optionally, enter an interim patch number.
   f. **Tag**: Apply a tag.
5. Click **Create**.

Download an Interim Software Update

The patch download path is valid for four days. Download the patch within the specified timeframe.

1. On the Interim Software Updates page, click the Actions menu (three dots) of the Interim Software Update you want to download and then select **Download**.

   The system starts downloading the patch.
2. You can also download a patch from the Interim Software Updates details page.

   - Click the Actions button (three dots) for the patch you're interested in, and select **Download**.

> **Note**
>
> You can only download the patches that are in **Available** state.

Interim Software Updates Lifecycle States:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **Available**: Patch has been created successfully and the time-to-live (TTL) has not expired.
- **Creating**: The patch creation process is in progress.
- **Expired**: The lifetime of the patch download link has expired, which means you cannot download it.
- **Failed**: The patch create failed due to some error.
- **Terminating**: The patch deletion process is in progress.
- **Terminated**: The patch has been deleted.

Delete an Interim Software Update

Be discrete in deleting interim software updates. However, you can delete the interim software updates that have expired to free up space in the Object Store.

1. On the Interim Software Updates page, click the Actions menu (three dots) of the Interim Software Update you want to delete and then select **Delete**.
2. In the resulting dialog, enter the name of the patch to confirm and then click **Delete**.
3. You can also delete a patch from the Interim Software Updates details page.
   - Click the Actions button (three dots) for the patch you're interested in, and select **Delete**.

Move an Interim Software Update Resource to Another Compartment

1. On the Interim Software Updates page, click the Actions menu (three dots) of the Interim Software Update you want to move and then select **Move resource**.
2. In the resulting dialog, choose a new compartment, and click **Move Resource**.
3. You can also move a patch from the Interim Software Updates details page.
   - Click the Actions button (three dots) for the patch you're interested in, and select **Move Resource**.

Using the API to Manage Interim Software Updates

The Oracle Cloud Infrastructure APIs are typical REST APIs that use HTTPS requests and responses, and enable you ot perform 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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `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.
- Enhanced Data Guard Health Status Reporting
  The enhanced Data Guard health status reporting provides comprehensive insights into protection mode, switchover and failover readiness, and data loss exposure across primary and standby databases.
- Using the Console to Manage 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 `libkmstdepkcs11_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 it's 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.
- Reinstate
  The reinstate 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Enhanced Data Guard Health Status Reporting

The enhanced Data Guard health status reporting provides comprehensive insights into protection mode, switchover and failover readiness, and data loss exposure across primary and standby databases.

With clear visual indicators (green, yellow, red, gray), you can quickly assess the readiness of your databases for role transitions and failover events. Additionally, detailed redo transport lag metrics help you evaluate potential data loss scenarios, enabling proactive disaster recovery planning and improved operational reliability.

Protection Mode

Data Guard associations provide two protection modes:

- Maximum Availability – Enables zero data loss failover and uses synchronous redo transport to at least one standby database.
- Maximum Performance – Enables near-zero data loss failover with no performance impact by using asynchronous transport to all standby targets.

Switchover Readiness

- Primary Database: status indicators – Green, Yellow, Red, Gray
  - Green (Healthy): All standby databases have passed all switchover readiness checks.
  - Yellow (Warning): Applicable only if the primary database has multiple standby databases. A subset of standby databases is switchover-ready. Address the failed checks.
  - Red (Critical): None of the standby databases are switchover-ready. Address failed checks to avoid extended downtime.
  - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Standby Database: status indicators – Green, Red, Gray
  - Green (Healthy): The standby database has passed all switchover readiness checks.
  - Red (Critical): This standby database is not switchover-ready. Address failed checks to avoid extended downtime during switchover attempts.
  - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Disabled Standby Database and Snapshot Standby Database: status indicator – Gray
  - Gray (Unknown): Disabled standby databases always appear as Unknown because their health status cannot be determined.

Failover Readiness

- Primary Database: status indicators – Green, Yellow, Red, Gray
  - Green (Healthy): All standby databases have passed all failover readiness checks.
  - Yellow (Warning): Applicable only if the primary database has multiple standby databases. A subset of standby databases is failover-ready. Address the failed checks.
  - Red (Critical): None of the standby databases are failover-ready. Address failed checks to avoid extended downtime.
  - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Standby Database and Snapshot Standby: status indicators – Green, Red, Gray

Copyright © 2023, 2026, Oracle and/or its affiliates.

  - Green (Healthy): The standby database has passed all failover readiness checks.
  - Red (Critical): This standby database is not failover-ready. Address failed checks to avoid extended downtime during failover.
  - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Disabled Standby Database: status indicator – Gray
  - Gray (Unknown): Disabled standby databases always appear as Unknown because their health status cannot be determined.

> **Note**
>
> When the health status cannot be determined for any database in a Data Guard group at a given point in time, the status appears as null in the SDK/CLI, and Terraform, and as UNKNOWN (gray indicator) in the Console.

Data Loss Exposure

- Primary Database: Data loss exposure is defined as the redo transport lag between the primary and standby databases. "Last Computed" if displayed represents the last time the metric could be retrieved and calculated.
- Standby Database: Data loss exposure is defined as the redo transport lag between the primary and standby databases. "Last Computed" if displayed represents the last time the metric could be retrieved and calculated.

Refresh Data Guard health status

Run an explicit refresh to get the latest status.

Using the Console to Manage 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 Oracle Exadata Database Service on Exascale Infrastructure
  Learn how to set up an Oracle Data Guard group 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- To Disable Automatic Failover
  Learn how to disable Automatic Failover for an Oracle Data Guard configuration.
- 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.
- Handling Standby Database Termination
  Understand the scenarios that would affect automatic failover when you terminate a standby database.

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

Learn how to set up an Oracle Data Guard group between databases.

> **Note**
>
> - When you enable Data Guard, replication of data happens only over the client network.
> - When you configure a Data Guard group, the primary and standby databases must be on the same major release version, while the standby database can be on a higher minor version.

> **Note**
>
> You can create a Data Guard configuration when the database is encrypted using OCI Vault. However, OCI Vault currently supports key replication to only one remote region. As a result, cross-region Data Guard configurations for the databases using OCI Vault are limited to two regions: the primary region and one remote region. Multiple standby databases can be created within these regions, but configurations spanning three or more regions are not supported.

As part of the latest release, we are introducing an enhanced user experience and new APIs to improve performance and provide additional Data Guard capabilities, including support for multiple standby databases with cloud automation.

- With the new API, your new Data Guard configuration will be created as a Data Guard group resource.
- If you have an existing Data Guard setup, you can continue to use current capabilities with no impact. However, if you want to create multiple standby databases, then you must migrate to the new API model, which can be done at any time.
- If you currently have automation that manages Data Guard operations using the existing Data Guard Association API, then you must update your applications to use the new API to take advantage of these new capabilities.

Copyright © 2023, 2026, Oracle and/or its affiliates.

  - ExaDB-XS supports both the existing Data Guard Association API and the new Data Guard Group API and the associated user interface for Oracle AI Database 26ai.
  - ExaDB-XS supports only the new Data Guard Group API and the associated user interface for Oracle Database 19c.

- In addition to provisioning standby databases for Oracle AI Database 26ai, you can provision an Oracle Database 19c standby database for an Oracle Database 19c primary database on block storage.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle 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 **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.
5. On the VM cluster details page, in the **Container databases** section, click the name of the database that you want to make primary.
6. On the Database Information page, select the Data Guard Group tab and click **Add Standby**.
7. On the Add Standby page, configure your Data Guard group.
   - To explicitly run a pre-check, click **Run Precheck**.

     When you click **Run Precheck**, the following actions are performed:
     - The system closes the **Add Standby** page.
     - A message appears on the Primary **Database Details** page indicating that the Data Guard precheck is in progress, and the **Add Standby** button is disabled. Click the **View Details** link to navigate to the **Work Requests** page.
       * Review the precheck status: **Succeeded** or **Failed**.
       * If the pre-check has failed, then resolve the reported errors and retry the Data Guard setup.

     > **Note**
     >
     > While the precheck is running, the primary database and the VM Cluster remain in the **AVAILABLE** state.

   - **Choose the Data Guard experience**:
     - **Use the new Data Guard Group resource**: With this option, your new Data Guard configuration will be created as a Data Guard Group resource. This option with new APIs supports adding multiple standby databases and provides other enhancements. If you currently have automation that manages Data Guard operations using the existing Data Guard Association API, you can update your applications to use the new API to take advantage of these new capabilities.
     - **Use the existing Data Guard Association resource**: Choose this option if your automation for managing Data Guard operations relies on the existing Data Guard

Copyright © 2023, 2026, Oracle and/or its affiliates.

Association API. However, you will not be able to add multiple standby databases and will not get the enhancements provided by the new API.

- **Automatic failover**: Enabling automatic failover ensures high availability of your database and is recommended for production databases running critical workloads. Oracle automatically initiates a failover to a standby database when the primary database becomes unavailable. Oracle recommends the standby database be in the same region as the primary database. If enabled, this standby database will be the failover target for the primary database.
  Automatic failover is enabled by default, but you can disable it if needed. The behavior of Data Guard protection modes varies depending on whether automatic failover is enabled or disabled.
  - When **Automatic failover** is enabled:
    - **Data Guard protection modes**:
      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 can be Synchronous or Asnychronous. See Redo Transport Services for information on these options.
      - **Maximum data loss threshold**: Available options: 0 seconds and Custom threshold.
        - **Maximum data loss threshold (0 seconds)**:

          > **Note**
          >
          > Enabling zero data loss automatic failover configures this standby database with synchronous transport to fully protect committed transactions. Please note that synchronous transport may affect performance if network latency between primary and the standby database is high.

          - **Protection mode**: Maximum Availability
          - **Transport Type**: Synchronous
        - **Maximum data loss threshold (Custom threshold)**:
          - **Protection mode**: Maximum Performance
          - **Transport Type**: Asynchronous
  - **When Automatic failover is disabled**:
    Protection Mode and Transport Type: Rules for Standby Database Creation
    - **Creating the first standby**: You cannot modify the Protection Mode or Transport Type for the first standby database during its creation. It is possible to modify it later.
      - The default settings are:
        - **Protection mode**: Maximum Performance
        - **Transport Type**: Asynchronous
    - **Creating the second to Nth standby**: You cannot modify the Protection Mode or Transport Type for any subsequent standby databases.

Copyright © 2023, 2026, Oracle and/or its affiliates.

    * The Protection Mode is unchanged when adding subsequent standby databases.
    * Subsequent standby databases will be added to the configuration with a Transport Type of asynchronous.

- In the **Select peer VM cluster** section, provide the following information for the standby database to obtain a list of available Exadata systems in which to locate the standby database:

  > **Note**
  >
  > – If automatic failover is enabled, ensure that the peer region is configured to match the primary database region.
  > – When the maximum data loss threshold is set to 0, the peer region drop down will display only one available option—the primary database region.
  > – Additionally, the availability domain (AD) must be configured to a different AD than that of the primary database to maintain proper high availability.

  – **Region**: Select a region where you want to locate the standby database. The region where the primary database is located is selected, by default. You can choose to locate the standby database in a different region. The hint text associated with this field tells you in which region the primary database is located.
  – **Availability domain**: Select an availability domain for the standby database. The hint text associated with this field tells you in which availability domain the primary database is located.
  – **Select a service**: Exadata Database Service on Dedicated Infrastructure or Exadata Database Service on Exascale Infrastructure.
  – **Select Exadata infrastructure:** Applicable only if you have selected Exadata Database Service on Dedicated Infrastructure from the **Select a service** drop-down list.
  – **Data Guard Peer resource type**: Select **VM Cluster** from the drop-down list.
  – **Choose the Data Guard Experience**:
    * **Use the new Data Guard Group Resource**: With this option, your new Data Guard configuration will be created as a Data Guard group resource. This option with new APIs supports adding multiple standby databases and provides other enhancements. If you currently have automation that manages Data Guard operations using the existing Data Guard Association API, you can update your applications to use the new API to take advantage of these new capabilities.
    * **Use the existing Data Guard Association Resource**: Choose this option if your automation for managing Data Guard operations relies on the existing Data Guard Association API. However, you will not be able to add multiple standby databases and will not get the enhancements provided by the new API.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> For Oracle Database 19c, ExaDB-XS supports only the Data Guard Group resource.

- **Data Guard group details:**
  - **Data Guard Type**: Select Active Data Guard or Data Guard. Active Data Guard provides additional features including: Real-Time Query and DML Offload, Automatic Block Repair, Standby Block Change Tracking, Far Sync, Global Data Services, and Application Continuity. Note that Active Data Guard requires an Oracle Active Data Guard license. For more information on Active Data Guard, see Active Data Guard. For a complete overview of both Data Guard types, see Introduction to Oracle Data Guard.
  - **Configure Failover Target**: Designate the standby database that should be the target of a failover should the current primary database become unavailable. If multiple failover targets are configured, Oracle will select the first target standby database in the list as the active target. If that database becomes unavailable, Oracle will select the next standby database in the list. Only the current active target database will be failed over to if the current primary database becomes unavailable.

    Click **Add failover target** to configure additional failover targets.
  - **Protection mode**: This is a read-only field. The protection mode can be **Maximum Performance** or **Maximum Availability**. See Oracle Data Guard Protection Modes for information on these options.
  - **Transport type:** This is a read-only field. It indicates the redo transport type used for this Data Guard group. See Redo Transport Services for information on these options.
    - **Creating the first standby:** You cannot modify the Protection Mode or Transport Type for the first standby database during its creation. It is possible to modify it later.
      - The default settings are as follows:
        - **Protection Mode:** Maximum Performance
        - **Transport Type:** Asynchronous
    - **Creating the second to Nth standby**: You cannot modify the Protection Mode or Transport Type for any subsequent standby databases.
      - The Protection Mode is inherited from the first standby.
      - The default Transport Type is set to Asynchronous.
- In the **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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

b. Select the region containing the database software image you want to use to create the new Database home. Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.

c. Select the Oracle Database software version that the new Database home will use, then choose an image from the list of available images for your selected software version.

d. Click **Select**.

> **Note**
>
> * Oracle recommends that you apply the same list of patches to the Database homes of the primary and standby databases.
> * If you are using the new Data Guard group resource, then you must first create the database home before adding the standby database.

- In the **Configure standby database** section, provide standby database details.

> **Note**
>
> You cannot modify the `db_unique_name` and SID prefix after creating the database.

  - **Database unique name:** Optionally, specify a value for the `DB_UNIQUE_NAME` database parameter. This value must be unique across the primary and standby cloud VM clusters. The unique name must meet the requirements:
    * Maximum of 30 characters
    * Contain only alphanumeric or underscore (_) characters
    * Begin with an alphabetic character
    * Unique across the VM cluster. Oracle recommends that it should be unique across the tenancy.

    If not specified, the system automatically generates a unique name value, as follows:

    <db_name>_<3_chars_unique_string>_<region-name>

  - **Database password:** Enter the database administrator password of the primary database. Use this same database administrator password for the standby database.

    > **Note**
    >
    > The administrator password and the TDE wallet password must be identical. If the passwords are not identical, then follow the instructions in Changing the Database Passwords to ensure that they are.

  - **TDE wallet password:** Enter the TDE wallet password.

Copyright © 2023, 2026, Oracle and/or its affiliates.

8. *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.

   > **Note**
   >
   > This option is available starting with Oracle AI Database 26ai and later versions.

9. **Database Storage**. This section applies only to Oracle Database 19c container databases (CDBs) running on VM Clusters configured with Exascale Block Storage. There is no need to allocate any storage for Oracle AI Database 26ai CDBs running on VM Clusters configured with Exascale Smart Storage.

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

    - – In the **Tags** tab, you can add tags to the database. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see Resource Tags. If you are not sure if you should apply tags, then skip this option (you can apply tags later) or ask your administrator.

11. Click **Add**. When you create the association, the details for a database and its peer display their respective roles as **Primary** or **Standby**.

    A work request is issued to configure the Data Guard association. The progress of the request and the stages of provisioning can be viewed on the Work Requests page of the respective **Standby** database.
    When the association is created, the **Standbys** section under the **Data Guard Group** tab displays the following details:

    - **Member** – The name of the standby database member
    - **Database Unique Name** – The globally unique name of the database
    - **VM Cluster** – The Cloud VM cluster where the database resides
    - **Role** – Current role of the database: **Primary** or **Standby**
    - **Failover Targets** – List of configured failover targets for the database
    - **Region** – The OCI region in which the database is located
    - **Service Name** – The Exadata Database Service type, such as **ExaDB-D** or **ExaDB-XS**
    - **Switchover Ready** – Health status of the database for switchover: **Healthy**, **Warning**, **Critical**, or **Unknown**

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> Optionally, you can download the latest `DbConnectionBundle` from the standby database to have the latest connect strings. For more information, see *To download the PDB Connection Bundle*.

- View Data Guard Provisioning Progress
  View the progress of Data Guard Provisioning tasks using the Work Requests page.
- To Enable Automatic Failover
  Learn how to enable Automatic Failover for an Oracle Data Guard configuration.
- To Add Failover Targets to the Primary or Standby Database
  Learn how to add failover targets to a primary or standby database in an Oracle Data Guard configuration.

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.
- To download the PDB Connection Bundle
  The PDB Connection Bundle is HA-aware and provides connection information for the PDB on all Data Guard peers.

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 Enable Automatic Failover

Learn how to enable Automatic Failover for an Oracle Data Guard configuration.

Before you enable automatic failover, configure the network to allow ingress on port 7095 to the VM Cluster. The agent running in DomU uses this port. Ensure that the applicable `iptables` rules, firewall rules, and security lists allow the required ingress traffic.

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

   The VM Clusters page is selected by default.
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 VM cluster that contains a database you want to configure.

Copyright © 2023, 2026, Oracle and/or its affiliates.

4. On the VM cluster details page, in the **Container Databases** section, click the name of the database you want to configure.
5. Click the **Data Guard Group** tab.
6. Click **Enable automatic failover**.
7. On the Manage automatic failover page, enable the feature by sliding the **Enable automatic failover** toggle button.
8. **Maximum Data Loss Threshold**
   - **Threshold set to zero (default)**:
     If the Maximum data loss threshold is set to 0 (Zero):
     - The protection mode must be **Maximum Availability**.
     - The transport type must be **Synchronous**.
     - A warning message is displayed if there are no standby databases configured with these settings.
     - Only standby databases configured with:
       * **Protection mode:** Maximum Availability
       * **Transport Type:** Synchronous

       will be available for selection as failover targets.
   - **Custom threshold**:
     If a custom (non-zero) maximum data loss threshold is specified:
     - **Protection mode:** Maximum Performance
     - **Transport Type:** Asynchronous
       In this configuration, standby databases using asynchronous transport can be selected as failover targets.
9. **Configure failover targets**

   After setting the threshold, configure the failover targets. These targets represent the standby databases that will transition to the primary role if the current primary database becomes unavailable.

To Add Failover Targets to the Primary or Standby Database

Learn how to add failover targets to a primary or standby database in an Oracle Data Guard configuration.

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

   The VM Clusters page is selected by default.
2. Choose your **Compartment** that contains the Oracle Exadata Database Service on Exascale Infrastructure instance with the database you want to configure.
3. Navigate to the VM cluster that contains the database.
4. On the VM cluster details page, in the **Container Databases** section, click the name of the database you want to configure.
5. Click the **Data Guard Group** tab.
6. In the **Standbys** section, select the database to add a failover target.
7. Click the **Actions** menu (three dots) and select **Configure Failover Target**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

8. On the Configure Failover Target page, select the desired failover target.

> **Note**
>
> You can only select standby databases that are in the same region.

9. Click **Update** to save your changes.

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 AI Database**, click **Oracle 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 following information is displayed:
   - The role of each database in this VM Cluster is indicated in the **Data Guard role** column.
   - The service on which each database is running is indicated in the **Service name** column.

To enable automatic backups on a standby database

Learn to enable automatic backups on a standby database.

> **Note**
>
> Starting August 06, 2025, for tenancies created in the FRA, PHX, or NRT regions, Autonomous Recovery Service will be the only backup destination when you enable automatic backup on databases.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle 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 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, Select the Data Guard Groups option.
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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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**
   > - You can enable automatic backup on the primary database, or on the standby database, or on both. You can select Autonomous Recovery Service or Object Storage as the backup destination.
   > - If automatic backups are enabled on both primary and standby databases, the backup destination on both databases must be configured to either Autonomous Recovery Service or Object Storage.

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 AI Database**, then click **Oracle 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:

   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**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle 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 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 that 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 that you want to manage, click the Actions menu, and then click **Edit Protection Mode**.
6. In the **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 Disable Automatic Failover

Learn how to disable Automatic Failover for an Oracle Data Guard configuration.

Failover targets will remain assigned, but an automatic failover will not be initiated to those failover targets.

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

   The VM Clusters page is selected by default.
2. Choose your **Compartment** that contains the Oracle Exadata Database Service on Exascale Infrastructure instance with the database you want to configure.
3. Navigate to the VM cluster that contains a database you want to configure.

Copyright © 2023, 2026, Oracle and/or its affiliates.

4. On the VM cluster details page, in the **Container Databases** section, click the name of the database you want to configure.
5. Click the **Data Guard Group** tab.
6. Click **Edit**.
7. Slide the **Enable automatic failover** toggle to **off**.
8. Review the warning message and click **Update**.

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 AI Database**, then click **Oracle 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 AI Database**, then click **Oracle 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 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**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle 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**.

Handling Standby Database Termination

Understand the scenarios that would affect automatic failover when you terminate a standby database.

Scenario 1: Terminate a standby database when it is not the only failover target for the primary database

Automatic Failover remains unchanged. The primary database continues to have other failover targets assigned.

Scenario 2: Terminate a standby database when it is the only failover target for the primary database, but other standby databases exist in the Data Guard Group

A warning message is displayed:

Failover targets have not been defined for the primary database. Automatic failover is enabled, but there will not be a standby database available to failover to in the event the primary database is not reachable. You can configure failover targets by clicking the action icon.

Scenario 3: Terminate a standby database when it is the only standby in the Data Guard Group

The Data Guard Group becomes empty since there are no remaining standby databases.

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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> In February 2026, the Data Guard Association model and its associated APIs will be replaced by the new Data Guard Group model and APIs. Beginning February 2026, all new Data Guard configurations provisioned from the Oracle Cloud Infrastructure (OCI) Console will automatically use the Data Guard Group model.

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

- CreateDataGuardAssociation
- ListDataGuardAssociations
- GetDataGuardAssociation
- UpdateDataGuardAssociation
- SwitchoverDataGuardAssociation
- FailoverDataGuardAssociation
- ReinstateDataGuardAssociation
- DeleteDatabase - To terminate an 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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* for Oracle Dattabase 19c, or "Using Oracle Virtual Private Database to Control Data Access" for Oracle AI Database 26ai.

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

- Oracle AI Database 26ai: Using Oracle Virtual Private Database to Control Data Access in *Oracle AI Database Advanced Security Guide*
- Oracle Database 19c Using Transparent Data Encryption with Other Oracle Features in *Oracle AI Database Advanced Security Guide*
- Oracle AI Database 26ai ADMINISTER KEY MANAGEMENT in *Oracle AI Database SQL Language Reference*
- Oracle Database 19c ADMINISTER KEY MANAGEMENT in *Oracle AI 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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;
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 `ENCRYPTION` clause.
`CLOUD_ONLY`	Tablespaces created in an Exadata database are transparently encrypted with the AES128 algorithm unless a different algorithm is specified in the `ENCRYPTION` clause. For non-cloud databases, tablespaces are only encrypted if the `ENCRYPTION` clause is specified. `ENCRYPTION` is the default value.
`DDL`	During creation, tablespaces are not transparently encrypted by default, and are only encrypted if the `ENCRYPTION` 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)

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 Oracle 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database
  IAM users can connect to the database instance by using either an IAM database password verifier or an IAM token.
- Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication
  You can enable an Oracle Database instance to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier).
- Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication
  For IAM token access to the database, the client application or tool requests a database token from IAM for the IAM user.

About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database

IAM users can connect to the database instance by using either an IAM database password verifier or an IAM token.

Using the IAM database password verifier is similar to the database password authentication process. However, instead of the password verifier (encrypted hash of the password) being stored in the database, the verifier is instead stored as part of the OCI IAM user profile.

The second connection method, the use of an IAM token for the database, is more modern. The use of token-based access is a better fit for Cloud resources such as Oracle Databases in the Exadata Cloud Infrastructure. The token is based on the strength that the IAM endpoint can enforce. This can be multi-factor authentication, which is stronger than the use of passwords alone. Another benefit of using tokens is that the password verifier (which is considered sensitive) is never stored or available in memory.

> **Note**
>
> Oracle Database supports the Oracle DBaaS integration for Oracle Cloud Infrastructure (OCI) IAM with identity domains as well as the legacy IAM, which does not include identity domains. Both default and non-default domain users and groups are supported when using IAM with Identity Domains.
>
> Support for non-default custom domains are only available with Oracle Database Release 19c, Version 19.21 and higher (but not Oracle Database Release 21c).

Oracle Cloud Infrastructure IAM integration with Oracle Exadata Database Service on Dedicated Infrastructure supports the following:

- *Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication*
- *Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication*

For complete details about the architecture for using IAM users on Oracle Exadata Database Service on Dedicated Infrastructure, see *Authenticating and Authorizing IAM Users for Oracle DBaaS Databases* in the Oracle Database 19c Security Guide and Oracle AI Database 26ai Security Guide.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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)

Copyright © 2023, 2026, Oracle and/or its affiliates.

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*
    - *Oracle Developer Tools for Visual Studio*
- If the application or tool does not support requesting an IAM database token through the client API, the IAM user can first use the Oracle Cloud Infrastructure command line interface (CLI) to retrieve the IAM database token and save it in a file location. For example, to use SQL*Plus and other applications and tools using this connection method, you first obtain the database token using the Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI). For more information, see db-token get. If the database client is configured for IAM database tokens, when a user logs in with the slash login form, the database driver uses the IAM database token that has been saved in default or specified file location.
- Some Oracle AI Database 26ai clients can also get a token directly from OCI IAM instead of using the OCI command line interface. Please review the client documentation to see which clients support this native IAM integration.
- A client application or tool can use an Oracle Cloud Infrastructure IAM instance principal or resource principal to get an IAM database token and use the IAM database token to authenticate itself to an Oracle Database instance. For more information, see *Mapping Instance and Resource Principals*.
- IAM users and OCI applications can request a database token from IAM with several methods, including using an API key. See *Configuring a Client Connection for SQL*Plus That Uses an IAM Token* for an example. See *Authenticating and Authorizing IAM Users*

Copyright © 2023, 2026, Oracle and/or its affiliates.

for Oracle DBaaS Databases for a description of other methods such as using a delegation token within an OCI cloud shell.

> **Note**
>
> If your database is in Restricted Mode, only DBAs with the `RESTRICTED SESSION` privilege can connect to the database.

If a user enters a username/password to log in, then the database driver uses the password verifier method to access the database. If the parameter `PASSWORD_AUTH=OCI_TOKEN`, then the database driver will instead user the username and password to connect directly to IAM and request a database token.

Related Topics

- Support for IAM Token-Based Authentication
- JDBC and UCP Downloads
- Identity and Access Management (IAM) Token-Based Authentication
- db-token get
- Oracle Data access Components - .NET Downloads
- NuGet Gallery
- Oracle Developer Tools for Visual Studio
- 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`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 PKI Certificate 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 PKI Certificate 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.**

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 an 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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_fitch` can log in to a database instance.

```
sqlplus /nolog
connect peter_fitch@db_connect_string
Enter password: password
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

Some special characters will require double quotation marks around `user_name` and . For example:

```
"peter_fitch@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`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.co
m))
  (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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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"
     ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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.

Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure

Exadata Database Service on Exascale Infrastructure@Azure enables you to store your database's transparent data encryption (TDE) keys, also known as master encryption keys (MEKs) in either a file-based Oracle wallet or in the OCI Vault.

This feature enables Exadata Database Service on Exascale Infrastructure@Azure users to utilize Azure Key Vault (AKV) Managed HSM, AKV Premium and AKV Standard for managing TDE MEKs. This integration allows applications, Azure services, and databases to use a

Copyright © 2023, 2026, Oracle and/or its affiliates.

centralized key management solution for enhanced security and simplified key lifecycle management.

- Prerequisites
  You must complete these steps before you can configure Azure Key Vault as the key management for your databases.
- Network Requirements for Creating an Identity Connector and KMS Resources
  Azure Key Management Service (KMS) resources support both public and private connectivity.
- Using the Console to Manage Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure
  Learn how to manage Azure Key Vault integration for Exadata Database Service on Exascale Infrastructure@Azure.
- Using the API to Manage Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure
- Updating the Multicloud PKCS#11 Driver
- Release Notes
  Review the changes made across the different releases of the Azure driver.

Prerequisites

You must complete these steps before you can configure Azure Key Vault as the key management for your databases.

The following steps must be completed before you can configure Azure Key Vault as Key Management Service at the Exadata VM Cluster level:

1. If you plan to use a private network, then you must first complete the registration required for delegated subnets to use advanced network features mentioned in Network planning for Oracle Database@Azure, and then create an Azure Virtual Network with at least one delegated subnet in it to be used by Exadata VM cluster.
2. Provision an Exadata VM Cluster via the Azure interface. See Provisioning an Exadata VM Cluster for Azure for step-by-step instructions.
3. Review the networking requirements to determine whether the VM Cluster will connect to Azure KMS through a public network or through private connectivity. For more information, see Connected Machine agent network requirements or Network Requirements for Creating an Identity Connector and KMS Resources for specific steps to follow.
4. Ensure that the following policy is created before creating the database.

   ```
   allow any-user to manage oracle-db-azure-vaults IN tenancy where ALL
   { request.principal.type in ('exadbvmcluster') }
   ```

Network Requirements for Creating an Identity Connector and KMS Resources

Azure Key Management Service (KMS) resources support both public and private connectivity.

The following sections outline the network requirements for public network access:

Copyright © 2023, 2026, Oracle and/or its affiliates.

Configuration Using Private Network

- **Arc agent network configuration**
  Make sure the endpoints listed below are accessible from the VM cluster nodes:
  - `guestconfiguration.azure.com`
  - `his.arc.azure.com`
  - `login.microsoftonline.com`
  - `management.azure.com`
  - `pas.windows.net`
  - `guestconfiguration.azure.com`
  - `packages.microsoft.com`
  - `aka.ms`

  To create an Identity Connector over a private network, an Azure Arc Private Link Scope and a Private Endpoint must be configured through the Azure portal. Refer to the Azure documentation for detailed steps on setting up private connectivity for Azure Arc-enabled servers.

  Important:

  - The Private Endpoint must be created in a non-delegated subnet within the Azure Virtual Network (VNet) that hosts the Oracle Exadata VM Cluster. Private Endpoints are not supported in delegated subnets. By default, Exadata VM Clusters are provisioned in delegated subnets.
  - Managed HSM requires private connectivity and is supported only in Azure regions that offer Advanced Networking features. For a list of supported regions, see Network planning for Oracle Database@Azure.

  To allow communication with private agent resources over the private network, a private DNS zone and corresponding A records must be created within the VCN's DNS configuration in your Oracle Cloud Infrastructure (OCI) tenancy.

  The DNS configuration for the Private Endpoint associated with the Private Link Scope must include the necessary private agent resource addresses. For more information, see *URLs* section in Connected Machine agent network requirements.

  First, retrieve the list of required addresses from the Azure portal. Then, update the DNS zone entry in OCI to complete the configuration.

  Steps to retrieve the list of required IP addresses:

  1. Sign in to the Azure portal.
  2. Search for "Azure Arc Private link scopes".
  3. Choose any private link scope from the list.
  4. Under the **Configure** menu, click **Private endpoint connections**.
  5. Click the **Private endpoint** link.
  6. Under **Settings**, select DNS configuration to view the required addresses.

  Example: Add a Private Agent Resource (for example, `gbl.his.arc.azure.com`)

  The IP address associated with `gbl.his.arc.azure.com`, along with any other required agent resources, must be defined in the private DNS zone.

  Steps:

Copyright © 2023, 2026, Oracle and/or its affiliates.

1. Create a Private Zone
   For more information, see Creating a Private DNS Zone.
   - Zone type: Primary
   - Zone name: <Descriptive name>
   - Compartment: <Compartment name or OCID>
2. Add DNS Records
   - Navigate to the Records tab on the zone details page.
   - Click Manage Records, then Add Record:
     * Name: gbl.his.arc.azure.com
     * Type: A (IPv4 address)
     * TTL (seconds): 3600
     * RDATA mode: Basic
     * Address: <Private IP address>
3. Publish the Zone.
4. Ensure the record appears on the zone's page after publishing.
5. Verify that connectivity to Azure services from the VM cluster is routed through the private network.

   Even with Private Connectivity, the following endpoints must be routed through the Azure NAT gateway.

   Agent resources:
   - `packages.microsoft.com`
   - `login.microsoftonline.com`
   - `pas.windows.net`
   - `management.azure.com`
   - `aka.ms`

- **Azure Key Vault private endpoints configuration**
  To access endpoints of Azure Key Vaults over private connectivity, you must create a DNS zone. Additionally, an A record mapping the fully qualified domain name (FQDN) of the resource to the IP address of the corresponding private endpoint must be added in the OCI tenancy.

  To access the Managed HSM service over a Private Endpoint in your virtual network hosting an Exadata VM Cluster, you can establish a private link connection to Managed HSM and associate it with either the default subnet or a non-delegated subnet. Follow the steps outlined in the "Configuration Using Private Network" section of the "Network Requirements for Creating an Identity Connector and KMS Resources" topic. For more information, see Integrate Managed HSM with Azure Private Link.

Configuration Using Public Network

Create a NAT gateway in Azure portal and associate it with the delegated subnet of the Exadata VM Cluster. For more information, see Create a NAT gateway and associate it with an existing subnet.

Alternatively, you can also use Azure network firewall to restrict traffic. For more information, see Deploy and configure Azure Firewall using the Azure portal.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Using the Console to Manage Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure

Learn how to manage Azure Key Vault integration for Exadata Database Service on Exascale Infrastructure@Azure.

- Create an Identity Connector from the OCI Console
  Creating an Identity Connector installs the Azure Arc agent on the Exadata VM Cluster VMs, registering them as Azure Arc-enabled virtual machines.
- View the details of an Identity Connector
  To view the details of an identity connector, use this procedure.
- Enable or Disable the Azure Key Management
  This step installs the required library on the VM cluster to support Azure Key Vault integration. Ensure that an identity connector is created before enabling Azure Key Management on the Exadata VM Cluster.
- Create Azure Key Vault (Managed HSM, Premium, and Standard) and Assign Required Permissions
  Create Azure Key Vault Managed HSM, Azure Key Vault Premium, or Azure Key Vault Standard, then assign the permission.
- Register Azure Key Vaults in the OCI console
  This is an alternative way to register your Azure key vaults from the OCI console. If you have already registered your vault during the creation of a database in your existing VM Cluster, you can skip this step.
- Create a Database and Use Azure Key Vault as the Key Management Solution
  This topic describes only the steps for creating a database and using Azure Key Vault as the key management solution.
- Change the Key Management from Oracle Wallet to Azure Key Vault
  Learn to change encryption keys between different encryption methods.
- Rotate the Keys Managed by Azure Key Vault for a Container Database
  To rotate the Azure key vault encryption key of a container database (CDB), use this procedure.
- Rotate the Keys Managed by Azure Key Vault for a Pluggable Database
  To rotate the Azure key vault encryption key of a pluggable database (PDB), use this procedure.

Create an Identity Connector from the OCI Console

Creating an Identity Connector installs the Azure Arc agent on the Exadata VM Cluster VMs, registering them as Azure Arc-enabled virtual machines.

This enables secure communication with the Azure Key Management Service (KMS) using the Azure identity generated by the Arc agent. The Azure Arc agent can communicate with Azure services over either a public network or a private connectivity setup. Learn more about Azure Arc.

Each VM cluster must have an identity connector enabled to access Azure resources. The identity connector establishes either a public or private connection between the VM cluster and Azure Key Management resources, depending on the roles assigned. You need roles to access the Azure KMS resources.

Copyright © 2023, 2026, Oracle and/or its affiliates.

To generate an **access token** for your current Azure account, see az account get-access-token.

You can create an identity connector in one of two ways—using the Oracle Exadata Database Service on Exascale Infrastructure interface or the Database Multicloud Integrations interface.

Oracle Exadata Database Service on Exascale Infrastructure

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM clusters page is selected by default.
2. From the list of **VM Clusters**, select the cluster you are using.
3. In the **VM Cluster information** tab, navigate to **Identity connector** located under **Multicloud information**. Click the **Create** button.

   > **Note**
   >
   > If an identity connector has not been created previously, it is displayed as **None**.

4. The **Identity connector name**, **ExaDB-XS VM cluster**, **Azure subscription id**, and **Azure resource group name** are read-only fields and will be populated with values.
5. Enter your **Azure tenant id**, and **Access token**.
6. Expand the **Show advanced options** section.
   The **Private connectivity information** and **Tags** sections populate.

   To enable a private endpoint connection, enter the **Azure arc private link scope** name if it is in the same subscription of Arc agent. Enter the **Full Azure Resource ID** if the **Azure arc private link scope** is in a different subscription.

   Enter the full Azure Resource ID in the following format:

   ```
   /subscriptions/<subscription-id>/resourceGroups/<resource-group>/providers/
   Microsoft.HybridCompute/privateLinkScopes/<azure-arc-private-link-scope>
   ```

7. To add tags for your resources, click **Add tag**, and then enter required values.
8. Review your selections, and then click **Create** to create the identity connector.

Database Multicloud Integrations

1. Open the navigation menu. Click **Oracle AI Database**, click **Database Multicloud Integrations**, and then click **Microsoft Azure Integration**.
2. Select **Identity Connectors** from the left navigation menu.
3. From the **Compartment** drop-down list, select your compartment that you are using.
4. Click **Create Identity Connector**.
5. In the Identity connector information section, provide the following information:
   a. **Identity mechanism**: (read-only) Name of identity mechanism
   b. **Resource Type**: Select **ExaDB-XS VM cluster** from the drop-down options
   c. **ExaDB-XS VM cluster compartment**: Select the compartment from the drop-down options
   d. **ExaDB-XS VM cluster**: Select the VM Cluster from the drop-down options

Copyright © 2023, 2026, Oracle and/or its affiliates.

e. **Identity connector name**: Once you select your ExaDB-XS VM cluster, the **Identity connector name** automatically populates a name.

6. In the **Azure information** section, provide the following information:
   a. Enter your **Azure tenant id**.
   b. The **Azure subscription id** and **Azure resource group name** fields populate values based on your ExaDB-XS VM Cluster selection.
   c. Enter an **Access token**.
7. Expand the **Show advanced options** section. The **Private connectivity information** and **Tags** sections populate. These fields are optional.
   To enable a private endpoint connection, enter the **Azure arc private link scope** name if it is in the same subscription of Arc agent. Enter the **Full Azure Resource ID** if the **Azure arc private link scope** is in a different subscription.

   Enter the full Azure Resource ID in the following format:

   ```
   /subscriptions/<subscription-id>/resourceGroups/<resource-group>/providers/
   Microsoft.HybridCompute/privateLinkScopes/<azure-arc-private-link-scope>
   ```

8. To add tags for your resources, click **Add tag**, and then enter required values.
9. Review your selections, and then click **Create**.
   The newly created identity connector will be displayed in the **Identity Connectors** page.

> **Note**
>
> The **Identity Connectors** page displays the list of identity connectors based on the resource type and can be filtered using the **Search and Filter** functionality.

View the details of an Identity Connector

To view the details of an identity connector, use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM clusters page is selected by default.
2. Click the name of the VM cluster of your choice.
3. On the resulting **VM Cluster Details** page, in the **Multicloud Information** section, confirm that the Identity connector field displays the identity connector created previously.
4. Click the name of the **Identity Connector** to view its details. You can also view the associated ExaDB-XS VM Cluster.

Enable or Disable the Azure Key Management

This step installs the required library on the VM cluster to support Azure Key Vault integration. Ensure that an identity connector is created before enabling Azure Key Management on the Exadata VM Cluster.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM clusters page is selected by default.
2. Click the name of the VM cluster of your choice.

Copyright © 2023, 2026, Oracle and/or its affiliates.

3. On the resulting **VM Cluster Details** page, in the **Multicloud Information** section, click the **Enable** link next to **Azure key store**.
4. On the resulting **Enable Azure key management** dialog, click **Enable** to confirm the operation.
   Confirming the action will install a library on your VM Cluster.

   The status of Azure key store changes from **Disabled** to **Enabled**.
5. To disable Azure key store, click the **Disable** link.
6. On the resulting **Disable Azure key management** dialog, click **Disable** to confirm the operation.
   Disabling the Azure key management removes the library installed during enablement, which will impact the availability of databases configured to use it.

> **Note**
>
> Azure key management is configured at the VM cluster level, requiring all databases in the cluster to use the same key management solution. However, databases that use Oracle Wallet can coexist alongside those that use Azure Key Vault within the same cluster.

Create Azure Key Vault (Managed HSM, Premium, and Standard) and Assign Required Permissions

Create Azure Key Vault Managed HSM, Azure Key Vault Premium, or Azure Key Vault Standard, then assign the permission.

For more information, see Create a key vault using the Azure portal.

> **Note**
>
> There are specific roles that must be assigned to the group to grant the necessary permissions for accessing and managing **Azure Key Vault Managed HSM**, **Azure Key Vault Premium**, and **Azure Key Vault Standard** resources.

1. Create a group and add members.
   Azure groups allow you to manage users by assigning them the same access and permissions to resources.
   - To manage group in Azure, you are required to have **User Administrator** or **Groups Administrator** role. For more information, see Manage Microsoft Entra groups and group membership.
   - You must create a security group and add members from the Microsoft Azure portal. You must select the **Security** option as your **Group type**. For more information, see Create a basic group and add members.
   - You can add Arc-enabled VMs to the group. The roles mentioned in the subsequent step can be assigned to a group containing multiple VMs, instead of assigning roles to individual VMs.
2. Assign the following roles based on the type of Azure Key Vault:
   - **For Managed HSM:**

Copyright © 2023, 2026, Oracle and/or its affiliates.

  - IAM: Reader
  - Local RBAC: Managed HSM crypto officer + Managed HSM crypto user
- **For Key Vault Premium and Standard:**
  - IAM: Reader + Key Vault Crypto Officer

For detailed steps, refer to Assign Azure roles using the Azure portal.

Register Azure Key Vaults in the OCI console

This is an alternative way to register your Azure key vaults from the OCI console. If you have already registered your vault during the creation of a database in your existing VM Cluster, you can skip this step.

1. From the OCI console, navigate to **Database Multicloud Integrations**, and then select **Microsoft Azure Integration**. From the **Microsoft Azure Integration** section, select Azure Key Vaults.

   > **Note**
   >
   > At least one key must be created in the vault on the Azure portal for the registration to succeed.

2. Select the **Register Azure key vaults** button.
3. From the drop-down list, select your **Compartment**.
4. To discover Azure Key Vault or Managed HSM within the same subscription in an Azure tenancy (Identity Connector and Azure Key Vault or Managed HSM in the same subscription) and register keys:
   a. From the **Azure key vaults** section, select an **identity connector** from the **Discover Azure key vaults using connector** list.
   b. Click **Discover**.
   The list of **Vault names** is displayed.
   c. Select the check box located next to **Vault name**.
5. To discover a single Azure Key Vault or Managed HSM within the same subscription in an Azure tenancy, you can provide either the short key vault name or the full resource ID. Provide the full resource ID in one of the following formats:
   For Azure Key Vault:

   ```
   /subscriptions/<subscription-id>/resourceGroups/<resource-groups>/
   providers/Microsoft.KeyVault/vaults/<key-vault-name>
   ```

   For Azure Managed HSM:

   ```
   /subscriptions/<subscription-id>/resourceGroups/<resource-groups>/
   providers/Microsoft.KeyVault/managedHSMs/<hsm-name>
   ```

   > **Note**
   >
   > Cross subscription only supports the full resource ID.

Copyright © 2023, 2026, Oracle and/or its affiliates.

6. To discover Azure Key Vault or Managed HSM across subscriptions in an Azure tenancy (Identity Connector and Azure Key Vault or Managed HSM in different subscriptions) and register keys:
   a. From the **Azure key vaults** section, select an **identity connector** from the **Discover Azure key vaults using connector** list.
   b. Provide the full resource ID of the Azure Key Vault in the following format.

      ```
      /subscriptions/<subscription-id>/resourceGroups/<resource-groups>/
      providers/Microsoft.KeyVault/vaults/<key-vault-name>
      ```

   c. Click **Discover**.
7. If you want to add tags for your resources, expand the **Advanced options** section, and then click **Add tag**.
8. Click **Register** to register the vaults locally in the OCI.
9. Once you register the vault, you can view the **Display name**, **State**, **Type**, **Azure resource group**, and **Created** information of the vaults in the list.
10. Select the vault you are using, and then click the **Identity connector associations** tab, which lists identity connector associations in the current compartment.

    > **Note**
    >
    > A default association is automatically created between the vault and the Identity Connector used during the vault registration process. This allows the vault to be used on the VM cluster associated with that specific Identity Connector.

    If you want to use the same vault in other clusters that are registered with different Identity Connectors (that is, not the one used during vault discovery), you must explicitly create an association between the vault and those additional Identity Connectors.
11. If you expand the **Advanced options** section, you can add **Tags** for organizing your resources.
12. Review your selections, and then click **Create**.

Create a Database and Use Azure Key Vault as the Key Management Solution

This topic describes only the steps for creating a database and using Azure Key Vault as the key management solution.

For the generic database creation procedure, see To Create a Database in an Existing VM Cluster.

Prerequisites

Before creating your first database and selecting Azure Key Vault for key management, ensure the following prerequisites are met:

- All network prerequisites outlined in the Network Requirements for Creating an Identity Connector and KMS Resources section are fulfilled
- The identity connector is created and available for use
- Azure key management is enabled at the VM cluster level
- The VM cluster has the necessary permissions to access the vaults

Copyright © 2023, 2026, Oracle and/or its affiliates.

- The vaults are registered as OCI resources

Limitations

- **Virtual Machines Restriction:** Scaling out a VM cluster does not automatically extend databases that use Azure Key Vault to the newly added virtual machine. To complete the extension, you must update the existing Identity Connector for the VM Cluster by supplying the Azure access token. After updating the Identity Connector, run the dbaascli database addInstance command to add the database instance to the new VM.
- **Data Guard Restrictions:**
  - When creating a standby database for a primary that uses Azure Key Vault, ensure that the target VM cluster has an active Identity Connector, Azure key management is enabled, and the required association between the Identity Connector and the Key Vault is properly configured.
  - Cross-region Data Guard and database restore operations are not supported for databases that use Azure Key Vault for key management.
- **PDB Operations Restriction:** Remote PDB operations—such as clone, refresh, and relocate—are supported only if both the source and destination databases use the same Transparent Data Encryption (TDE) key.

Steps

> **Note**
>
> If Azure key management is enabled at the VM cluster level, you will have two key management options: **Oracle Wallet** and **Azure Key Vault**.

1. In the **Encryption** section, choose **Azure Key Vault**.
2. Select a registered **Vault** available in your compartment.

> **Note**
>
> - The Vault list populates only registered vaults.
> Click the **Register new vaults** link to register your vault. From the Register Azure key vaults page, select your vault, and then click **Register**.
> - At least one key must be registered in your vaults.

3. Select the Key available in your compartment.

Change the Key Management from Oracle Wallet to Azure Key Vault

Learn to change encryption keys between different encryption methods.

1. Navigate to your existing VM Cluster in the OCI console. Select the **Container Databases** tab. Then, select the database resource that you are using.
2. Select the **Database information tab**, and then scroll down to **Key management** section.
3. In the **Encryption** section, verify that **Key management** is set to **Oracle Wallet**, and then select the **Change** link.

Copyright © 2023, 2026, Oracle and/or its affiliates.

4. Enter the following information on the **Change key management** page:
   a. Select your **Key management** as **Azure key vault** from the drop-down list.
   b. Select your **Vault** compartment that you are using, and then select your **Vault** that is available in the compartment.
   c. Select the **Key** compartment that you are using, and then select your **Key** from the drop-down list.
   d. Click **Save changes**.

> **Note**
>
> Changing key management from Azure Key Vault to Oracle Wallet cannot be performed using the API or OCI Console—it is only supported through the dbaascli tde fileToHsm command. Additionally, switching between Azure Key Vault and OCI Vault or Oracle Key Vault (OKV) are not supported.

Rotate the Keys Managed by Azure Key Vault for a Container Database

To rotate the Azure key vault encryption key of a container database (CDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on 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 that you want to rotate encryption keys.
4. Click **Container Databases**.
5. Click the name of the database that you want to rotate encryption keys.
   The Database Details page displays information about the selected database.
6. In the **Encryption** section, verify that the **Key Management** is set to **Azure Key Vault**, and then click the **Rotate** link.
7. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

> **Note**
>
> Key rotation must be performed through the OCI interface. Rotating the key directly from the Azure interface has no effect on the database.

Rotate the Keys Managed by Azure Key Vault for a Pluggable Database

To rotate the Azure key vault encryption key of a pluggable database (PDB), use this procedure.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

A list of VM Clusters is displayed for the chosen Compartment.

3. In the list of VM clusters, click the name of the VM cluster that contains the PDB you want to start, and then click its name to display the details page.
4. Under **Container Databases**, find the database containing the PDB you want to rotate encryption keys.
5. Click the name of the database to view the Database Details page.
6. Click Pluggable Databases in the Resources section of the page.
   A list of existing PDBs in this database is displayed.
7. Click the name of the PDB that you want to rotate encryption keys.
   The pluggable details page is displayed.
8. In the **Encryption** section displays that the Key management is set as Azure Key Vault.
9. Click the **Rotate** link.
10. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

Using the API to Manage Azure Key Vault Integration for Exadata Database Service on Exascale Infrastructure@Azure

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

Use these API operations to manage Azure Key Vault integration for Exadata Database Service on Exascale Infrastructure@Azure.

Table 5-10 API operation to manage Azure Key Vault integration for Exadata Database Service on Exascale Infrastructure@Azure

API	Description
`createOracleDbAzureConnector`	Captures Azure-specific details from the customer and automates the installation of the Arc Agent on the VM Cluster.
`deleteOracleDbAzureConnector`	Deletes the Azure Connector resource and uninstalls the Arc Agent from the VM Cluster.
`getOracleDbAzureConnector`	Fetches the details of a specific Azure Connector resource.
`listOracleDbAzureConnectors`	Lists Azure Connector resources based on the specified filters.
`CreateMultiCloudResourceDiscovery`	Creates a new multi-cloud resource discovery resource.
`GetMultiCloudResourceDiscovery`	Retrieves details of a specific multi-cloud resource discovery resource.
`ListMultiCloudResourceDiscoveries`	Retrieves a list of all multi-cloud resource discovery resources.
`CreateOracleDbAzureVaultAssociation`	Creates a new association between an Oracle DB and an Azure vault.
`GetOracleDbAzureVaultAssociation`	Retrieves details of a specific Oracle DB Azure vault association.
`ListOracleDbAzureVaultAssociations`	Retrieves a list of all Oracle DB Azure vault associations.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-10 (Cont.) API operation to manage Azure Key Vault integration for Exadata Database Service on Exascale Infrastructure@Azure

API	Description
`createExadbVMCluster`	Creates a VM cluster.
`getExadbVmCluster`	Gets information about the specified VM cluster.
`listExadbVmClusters`	Gets a list of the VM clusters in the specified compartment.
`deleteExadbVMCluster`	Deletes the specified VM cluster.
`CreateDatabase`	Creates a new database in the specified Database Home. If the database version is provided, it must match the version of the Database Home.
`CreateDatabaseFromBackup`	Details for creating a database by restoring from a database backup. **Warning:** Oracle recommends that you avoid using any confidential information when you supply string values using the API.
`MigrateVaultKey`	Changes encryption key management from customer-managed, using the Vault service, to Oracle-managed.
`RotateVaultKey`	Creates a new version of an existing Vault service key.
`RestoreDatabase`	Restores a Database based on the request parameters you provide.

Updating the Multicloud PKCS#11 Driver

> **Note**
>
> The PKCS#11 driver for Azure is installed automatically when Azure key management is enabled on a VM cluster. To verify that the driver is installed and check its version, run the following command:
>
> ```
> rpm -qa | grep pkcs
> ```
>
> This command lists the PKCS package installed on the VM cluster. Only one PKCS#11 driver can be active at a time, corresponding to either Azure, Google Cloud, or AWS.
>
> Example output:
>
> ```
> pkcs-multicloud-driver-maz-0.2-250814.1708.x86_64
> ```
>
> This confirms that the `pkcs-multicloud-driver-maz` package is installed, along with its version and architecture.

The PKCS#11 driver can be updated using either a rolling or non-rolling mechanism, depending on customer preference.

Copyright © 2023, 2026, Oracle and/or its affiliates.

In rolling mode, the update is applied to one VM at a time, and the databases on that VM are restarted before proceeding to the next. In non-rolling mode, all databases in the cluster are shut down, the update is applied across all nodes, and the databases are then brought back online. Users must be aware of the `dbaastools` version in use, as subsequent steps, specifically step 3 in both scenarios, require different actions depending on the version, particularly when updating the PKCS#11 driver.

The version of `dbaastools` can be determined by running the following command:

```
/var/opt/oracle/dbaascli/dbaascli admin showLatestStackVersion
```

Rolling update

On each Guest VM, apply the PKCS#11 driver update as follows (one node at a time):

1. Shutdown the database instances that are using the Azure KMS.
   All databases that use the Azure KMS for master encryption keys (MEK) must be stopped before updating the driver. Use:

   ```
   srvctl stop instance -d <db_unique_name> -i <instance_name> [-o
   <stop_option>]
   ```

   - `-d <db_unique_name>`: The unique name of the database (same as `DB_UNIQUE_NAME` in the database).
   - `-i <instance_name>`: The name of the instance to stop, for example, `orcl1`.

   For details on additional options, refer to the SRVCTL command-line help or reference manual.
2. Verify that the databases have been shut down.
   Ensure that all database instances that are using the Azure KMS have been stopped before proceeding.
3. Update the RPM.
   As the `root` user, run the following command (for dbaastools 25.4.1 and later):

   ```
   /var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider
   AZURE --nodeList <node> --databasesStopped
   ```

   This command downloads and installs the latest RPM on the specified node.
4. Verify driver installation.
   Run the following command as the `root` user:

   ```
   /bin/rpm -qa | grep -i pkcs-multicloud-driver
   ```

5. Start the database instances.
   After the driver update, restart the database instances using:

   ```
   <oracle_home>/bin/srvctl start instance -node "<node>"
   ```

Non-Rolling update

In non-rolling mode, all databases on the cluster are stopped before applying the PKCS#11 driver update.

Copyright © 2023, 2026, Oracle and/or its affiliates.

1. Stop databases using the Azure KMS.
   Stop all databases that use the Azure KMS for master encryption keys (MEK) before updating the driver:

   ```
   dbaascli database stop [--dbname <value>]
   ```

2. Verify that the databases have been shut down.
   Ensure that all databases using the Azure KMS have been stopped before proceeding.
3. Upgrade the RPM.
   As the `root` user, run the following command (for dbaastools 25.4.1 and later):

   ```
   /var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider
   AZURE --databasesStopped
   ```

   This command downloads and installs the latest RPM.
4. Verify driver installation.
   Run the following command as the `root` user:

   ```
   /bin/rpm -qa | grep -i pkcs-multicloud-driver
   ```

5. Restart the databases.
   Once the driver has been updated, restart the databases:

   ```
   dbaascli database start [--dbname <value>]
   ```

Release Notes

Review the changes made across the different releases of the Azure driver.

- Release 2.202 (260721.2348)
- Release 2.196 (260325.2123)
- Release 2.194 (260227.1414)
- Release 0.2 (250910.0501)

Release 2.202 (260721.2348)

- Various bug fixes and stability improvements.

Release 2.196 (260325.2123)

- Various bug fixes and stability improvements.

Release 2.194 (260227.1414)

- Various bug fixes and stability improvements.

Release 0.2 (250910.0501)

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Various bug fixes and stability improvements.

Google Cloud Key Management Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud

Exadata Database Service on Exascale Infrastructure@Google Cloud now supports integration with Google Cloud Platform's Key Management Service (KMS).

This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using GCP Customer-Managed Encryption Keys (CMEKs).

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

This integration enables applications, Google Cloud services, and databases to benefit from a centralized key management solution that offers enhanced security and simplified key lifecycle management.

- Prerequisites
  Before configuring GCP Customer Managed Encryption Keys (CMEK) as the key management service for your databases, ensure the following prerequisites are met.
- Using the Console to Manage GCP KMS Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud
  Learn how to manage GCP KMS integration for Exadata Database Service on Exascale Infrastructure@Google Cloud.
- Using the API to Manage GCP KMS Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud
- Updating the Multicloud PKCS#11 Driver
- Release Notes
  Review the changes made across the different releases of the GCP driver.

Prerequisites

Before configuring GCP Customer Managed Encryption Keys (CMEK) as the key management service for your databases, ensure the following prerequisites are met.

1. Provision a VM Cluster using the Google Cloud console. See Provisioning a VM Cluster for Google Cloud for step-by-step instructions.
2. Review the Identity Connector connection to ensure it is correctly configured and active. For more information, see *Verify the Default Identity Connector Attached to the VM Cluster*.
3. Prerequisites for Configuring GCP Customer Managed Encryption Keys (CMEK) at the VM Cluster Level.
   To enable Google Cloud Platform (GCP) Customer Managed Encryption Keys (CMEK) for databases deployed with Exadata Database Service on Exascale Infrastructure@Google Cloud, you must configure CMEK as the key management option at the VM cluster level. Once CMEK is enabled, all database encryption and decryption operations will use the specified GCP-managed key.

   Before enabling CMEK, ensure that:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- The required GCP key rings and encryption keys are already created in GCP.
- These keys are mirrored as anchor resources in Oracle Cloud Infrastructure (OCI), ensuring synchronization between GCP and OCI.
- The anchor resources are in place for database provisioning and for managing the encryption key lifecycle, including key rotation, revocation, and auditing.

4. IAM Policy Requirements for Accessing GCP Key Resources.
   The database uses the cluster resource principal to securely retrieve GCP key resources. To enable this functionality, you must define the appropriate IAM policies in your OCI tenancy.

   Read-Only Access to Oracle GCP Keys:

   ```
   Allow any-user to read oracle-db-gcp-keys in compartment id <your-
   compartment-OCID>
   where all { request.principal.type = 'exadbvmcluster',}
   ```

   This policy grants read-only access to GCP key resources for the VM cluster resource principal.

5. Required Google Cloud Services Connectivity GCP CMEK Integration.
   Google Cloud VPCs typically include default routes to the services listed below. Ensure that no firewall egress rules block access to these endpoints.
   - `https://iamcredentials.googleapis.com/`
   - `https://sts.googleapis.com/`
   - `https://cloudkms.googleapis.com/`

Using the Console to Manage GCP KMS Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud

Learn how to manage GCP KMS integration for Exadata Database Service on Exascale Infrastructure@Google Cloud.

- Verify the Default Identity Connector Attached to the VM Cluster
  To view the details of an identity connector attached to a VM cluster, use this procedure.
- Create a Key Ring in Google Cloud Console
  To create a key ring, use this procedure.
- Create a Key in Google Cloud Console
  To create a raw symmetric encryption key in the specified key ring and location, use this procedure.
- Grant Permissions in Google Cloud KMS for Key Discovery by Oracle Cloud Infrastructure (OCI)
  To allow a key to be discoverable in Oracle Cloud Infrastructure (OCI), use this procedure.
- Register GCP Key Ring in Oracle Cloud Infrastructure (OCI)
  To enable Google Cloud Customer Managed Encryption Keys (CMEK) for your VM cluster, you must first register the GCP Key Ring in OCI.
- Enable or Disable Google Cloud Key Management
  To enable GCP CMEK for your VM Cluster, use this procedure.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Create a Database and Use GCP Customer-Managed Encryption Key (CMEK) as the Key Management Solution
  This topic describes only the steps for creating a database and using GCP Customer-managed encryption key (CMEK) as the key management solution.
- Change the Key Management from Oracle Wallet to GCP Customer Managed Encryption Key (CMEK)
  To change encryption keys between different encryption methods, use this procedure.
- Rotate the GCP Customer Managed Encryption Key of a Container Database (CDB)
  To rotate the GCP Customer Managed Encryption Key of a container database (CDB), use this procedure.
- Rotate the GCP Customer Managed Encryption Key of a Pluggable Database (PDB)
  To rotate the GCP Customer Managed Encryption Key of a pluggable database (PDB), use this procedure.

Verify the Default Identity Connector Attached to the VM Cluster

To view the details of an identity connector attached to a VM cluster, use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM clusters page is selected by default.
2. From the list of **VM Clusters**, select the cluster of your choice.
3. On the resulting **VM Cluster Details** page, in the **Multicloud Information** section, confirm that the Identity connector field displays the identity connector attached to this VM cluster.
4. Click the name of the **Identity Connector** to view its details.
   You will be redirected to the Database Multicloud Integrations portal.

Create a Key Ring in Google Cloud Console

To create a key ring, use this procedure.

1. Open the **Google Cloud Console**, navigate to the **Key Management** page.
2. Click **Create key ring**.
3. Provide the following details:
   - **Name**: Enter a descriptive name for the key ring.
   - **Location**: Select a location for your key ring.
     Important:
     - Key rings with the same name can exist in different locations, so you must always specify the location.
     - Choose a location close to the resources you want to protect.
     - For Customer Managed Encryption Keys, ensure the key ring is in the same location as the resources that will use it.

     Choosing a location for your Key Ring:

     When creating a key ring in Google Cloud Key Management Service (KMS), selecting the right location is crucial. Your choice affects where your cryptographic keys are stored and how they're replicated. For more information, see Cloud KMS locations.

     - **Region:**
       * Data is stored in a specific geographic region.

Copyright © 2023, 2026, Oracle and/or its affiliates.

        - Keys remain within the boundaries of this single region.
        - Ideal for:
            - Low-latency applications
            - Compliance with data residency requirements
            - Region-specific workloads
    - **Multi-region:**
        - Data is replicated across multiple regions within a larger geographical area.
        - Google manages distribution and replication automatically.
        - You cannot select individual data centers or regions.
        - Ideal for:
            - High availability
            - Resilient, fault-tolerant applications
            - Services serving a wide regional area
    - **Global:**
        - A special type of multi-region.
        - Keys are distributed across Google data centers worldwide.
        - Location selection and control are not available.
        - Ideal for:
            - Applications with global users
            - Use cases needing maximum redundancy and reach
4. Click **Create**.

Once the key ring is created, you can begin creating and managing encryption keys within it.

Create a Key in Google Cloud Console

To create a raw symmetric encryption key in the specified key ring and location, use this procedure.

1. Open the **Google Cloud Console**, navigate to the **Key Management** page.
2. Click the name of the key ring where you want to create the key.
3. Click **Create key**.
4. Provide the following details:
    - **Key name**: Enter a descriptive name for your key.
    - **Protection level**: Choose **Software** or **HSM** (Hardware Security Module).
      The protection level of a key can't be changed after the key is created. For more information, see Protection levels.
    - **Key material**: Select **Generate key** or **Import key**.
      Generate key material in Cloud KMS or import key material that is maintained outside of Google Cloud. For more information, see Customer-managed encryption keys (CMEK).
    - **Purpose and Algorithm**:
      For more information, see Key purposes and algorithms.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Set **Purpose** to **Raw encryption/decryption**.
- For **Algorithm**, select **AES-256-CBC**.

5. Click **Create**.

After creation, you can use this key for cryptographic operations that require AES-CBC encryption and decryption.

Grant Permissions in Google Cloud KMS for Key Discovery by Oracle Cloud Infrastructure (OCI)

To allow a key to be discoverable in Oracle Cloud Infrastructure (OCI), use this procedure.

1. In **Google Cloud KMS**, select the key you want to make discoverable.
2. Navigate to the **Permissions** tab and click **Add principal**.
3. In the **New principals** field, enter the service account associated with your **Workload Resource Service Agent**.

> **Note**
>
> You can find this service account on the **Identity Connector details** page, under the **GCP Information** section. Look for the **Workload resource service agent** and note its ID — this is the required service account.

4. Under **Assign roles**, add a role of your choice.

> **Note**
>
> Create a custom role with the following minimum permissions and assign it to the key ring of your choice.

These permissions together allow OCI to:

- Discover KMS resources like key rings and keys.
- Access metadata about keys and their versions.
- Use the keys for cryptographic operations (encryption/decryption).
- Create key versions.

Minimum Required Permissions:

- `cloudkms.cryptoKeyVersions.get`
  Allows retrieval of metadata for a specific key version.
- `cloudkms.cryptoKeyVersions.manageRawAesCbcKeys`
  Enables management of raw AES-CBC key material (import, rotation, and so on).
- `cloudkms.cryptoKeyVersions.create`
  Allows creation of new key versions within a key.
- `cloudkms.cryptoKeyVersions.list`
  Lists all versions of a given key.
- `cloudkms.cryptoKeyVersions.useToDecrypt`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Grants permission to use a key version for decrypting data.

- `cloudkms.cryptoKeyVersions.useToEncrypt`
  Grants permission to use a key version for encrypting data.
- `cloudkms.cryptoKeys.get`
  Allows retrieval of metadata for a key.
- `cloudkms.cryptoKeys.list`
  Lists all keys within a key ring.
- `cloudkms.keyRings.get`
  Allows retrieval of metadata for a key ring.
- `cloudkms.locations.get`
  Retrieves information about supported key locations.

5. Click **Save** to apply the changes.
6. Click **Refresh** to confirm that the updated permissions have taken effect.

Register GCP Key Ring in Oracle Cloud Infrastructure (OCI)

To enable Google Cloud Customer Managed Encryption Keys (CMEK) for your VM cluster, you must first register the GCP Key Ring in OCI.

> **Note**
>
> Before proceeding, ensure that the permissions outlined in Grant Permissions in Google Cloud KMS for Key Discovery by Oracle Cloud Infrastructure (OCI) have been granted.

1. In the **Database Multicloud Integrations** portal, navigate to: **Google Cloud Integration** > **GCP Key Rings**.
2. Click **GCP Key Ring**,
3. Click **Register GCP key rings**
4. On the resulting **Register GCP key rings** page, provide the following details:
   - **Compartment**: Select the compartment where the VM cluster resides.
   - **Identity Connector**: Choose the Identity Connector attached to the VM cluster.
   - **Key Ring**: Enter the name of the GCP key ring to register.
     To discover all available key rings through a single identity connector, you must grant the following permissions to that identity connector. These permissions should be assigned at the appropriate project or folder level to ensure the connector can access all key rings across the intended scope.
     - `cloudkms.keyRings.list`
       Allows listing all key rings within a project.
     - `cloudkms.locations.get`
       Allows retrieving metadata for a specific key ring.
5. Click **Discover** to verify if the key ring exists in GCP.
   If successful, the key ring's details will be displayed.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> Only key rings can be registered — not individual keys. All supported keys associated with a registered key ring will be available, provided the required permissions are in place.

6. Click **Register**.

Enable or Disable Google Cloud Key Management

To enable GCP CMEK for your VM Cluster, use this procedure.

> **Note**
>
> When you provision an Exadata VM Cluster, GCP CMEK is disabled by default.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
   The VM clusters page is selected by default.
2. From the list of **VM Clusters**, select the cluster you want to configure.
3. On the **VM Cluster Details** page, scroll to the **Multicloud Information** section and click **Enable** next to GCP CMEK.
4. To disable GCP CMEK, click **Disable**.

Create a Database and Use GCP Customer-Managed Encryption Key (CMEK) as the Key Management Solution

This topic describes only the steps for creating a database and using GCP Customer-managed encryption key (CMEK) as the key management solution.

For the generic database creation procedure, see To Create a Database in an Existing VM Cluster.

Prerequisites

- Enable Google Cloud Key Management at the VM cluster level.
- Register the GCP key rings in OCI.

Steps

If Google Cloud Key Management is enabled at the VM cluster, you will have two key management options: **Oracle Wallet** and **GCP Customer Managed Encryption Key**.

1. In the Encryption section, choose **GCP Customer Managed Encryption Key**.
2. Select a registered Key ring available in your compartment.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> - Only registered key rings are listed.
> - If your desired key ring is not visible, it may not have been registered yet. Click **Register Key Rings** to discover and register it.
> For detailed instructions, refer to Register GCP Key Ring in Oracle Cloud Infrastructure (OCI).

3. Select the key within the selected key ring in your compartment.

Change the Key Management from Oracle Wallet to GCP Customer Managed Encryption Key (CMEK)

To change encryption keys between different encryption methods, use this procedure.

> **Note**
>
> - You cannot migrate from GCP Customer Managed Encryption Key to Oracle Wallet.
> - Your database will experience a brief downtime while the key management configuration is being updated.

1. Navigate to your database details page in the OCI console.
2. In the **Encryption** section, verify that **Key management** is set to **Oracle Wallet**, and then click the **Change** link.
3. Enter the following information on the **Change key management** page.
   a. Select your **Key management** as **GCP Customer Managed Encryption Key** from the drop-down list.
   b. Select the compartment you are using, and then choose the Key Ring available in that compartment.
   c. Next, select the Key compartment you are using, and then choose the desired Key from the drop-down list.
   d. Click **Save changes**.

Rotate the GCP Customer Managed Encryption Key of a Container Database (CDB)

To rotate the GCP Customer Managed Encryption Key of a container database (CDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on 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 that you want to rotate encryption keys.
4. Click **Container Databases**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5. Click the name of the database that you want to rotate encryption keys.
   The Database Details page displays information about the selected database.
6. In the **Encryption** section, verify that the **Key Management** is set to **GCP Customer Managed Encryption Key**, and then click the **Rotate** link.
7. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

Rotate the GCP Customer Managed Encryption Key of a Pluggable Database (PDB)

To rotate the GCP Customer Managed Encryption Key of a pluggable database (PDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on 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 PDB you want to start, and then click its name to display the details page.
4. Under **Container Databases**, find the database containing the PDB you want to rotate encryption keys.
5. Click the name of the database to view the Database Details page.
6. Click **Pluggable Databases** in the **Resources** section of the page.
   A list of existing PDBs in this database is displayed.
7. Click the name of the PDB that you want to rotate encryption keys.
   The pluggable details page is displayed.
8. In the **Encryption** section displays that the Key management is set as GCP Customer Managed Encryption Key.
9. Click the **Rotate** link.
10. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

Using the API to Manage GCP KMS Integration for Exadata Database Service on Exascale Infrastructure@Google Cloud

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

The following resources will be made available to customers through OCI SDK, CLI, and Terraform. These APIs will be used by customers who wish to integrate Exascale Infrastructure on Exadata with Google Cloud Services.

Table 5-11 OracleDbGcpIdentityConnectors

API	Description
`ListOracleDbGcpIdentityConnectors`	Lists all GCP Identity Connector resources based on the specified filters.
`GetOracleDbGcpIdentityConnector`	Retrieves detailed information about a specific GCP Identity Connector resource.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-11 (Cont.) OracleDbGcpIdentityConnectors

API	Description
`CreateOracleDbGcpIdentityConnector`	Creates a new GCP Identity Connector resource for the specified ExaDB-D VM Cluster.
`UpdateOracleDbGcpIdentityConnector`	Updates the configuration details of an existing GCP Identity Connector resource.
`ChangeOracleDbGcpIdentityConnectorCompartment`	Moves the GCP Identity Connector resource to a different compartment.
`DeleteOracleDbGcpIdentityConnector`	Deletes the specified GCP Identity Connector resource.

Table 5-12 OracleDbGcpKeyRings

API	Description
`ListOracleDbGcpKeyRings`	Lists all GCP Key Ring resources based on the specified filters.
`CreateOracleDbGcpKeyRing`	Creates a new GCP Key Ring resource.
`ChangeOracleDbGcpKeyRingCompartment`	Moves the GCP Key Ring resource to a different compartment.
`RefreshOracleDbGcpKeyRing`	Refreshes the details of a GCP Key Ring resource.
`GetOracleDbGcpKeyRing`	Retrieves detailed information about a specific GCP Key Ring resource.
`UpdateOracleDbGcpKeyRing`	Updates the configuration details of an existing GCP Key Ring resource.
`DeleteOracleDbGcpKeyRing`	Deletes the specified GCP Key Ring resource.

Table 5-13 OracleDbGcpKeyKeys

API	Description
`ListOracleDbGcpKeys`	Lists all GCP Key Ring resources based on the specified filters.
`GetOracleDbGcpKey`	Retrieves detailed information about a specific GCP Key resource.

Updating the Multicloud PKCS#11 Driver

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> The PKCS#11 driver for Google Cloud is installed automatically when Google Cloud Key Management is enabled on a VM cluster. To verify that the driver is installed and check its version, run the following command:
>
> ```
> rpm -qa | grep pkcs
> ```
>
> This command lists the PKCS package installed on the VM cluster. Only one PKCS#11 driver can be active at a time, corresponding to either Azure, Google Cloud, or AWS.
>
> Example output:
>
> ```
> pkcs-multicloud-driver-gcp-0.1-250723.0511.x86_64
> ```
>
> This confirms that the `pkcs-multicloud-driver-gcp` package is installed, along with its version and architecture.

The PKCS#11 driver can be updated using either a rolling or non-rolling mechanism, depending on customer preference.

In rolling mode, the update is applied to one VM at a time, and the databases on that VM are restarted before proceeding to the next. In non-rolling mode, all databases in the cluster are shut down, the update is applied across all nodes, and the databases are then brought back online. Users must be aware of the `dbaastools` version in use, as subsequent steps, specifically step 3 in both scenarios, require

The version of `dbaastools` can be determined by running the following command:

```
/var/opt/oracle/dbaascli/dbaascli admin showLatestStackVersion
```

Rolling update

On each Guest VM, apply the PKCS#11 driver update as follows (one VM at a time):

1. Shutdown the database instances that are using the Google Cloud KMS.
   All databases that use the Google Cloud KMS for master encryption keys (MEK) must be stopped before updating the driver. Use:

   ```
   srvctl stop instance -d <db_unique_name> -i <instance_name> [-o
   <stop_option>]
   ```

   - `-d <db_unique_name>`: The unique name of the database (same as `DB_UNIQUE_NAME` in the database).
   - `-i <instance_name>`: The name of the instance to stop, for example, `orcl1`.

   For details on additional options, refer to the SRVCTL command-line help or reference manual.
2. Verify that the databases have been shut down.
   Ensure that all database instances that are using the Google Cloud KMS have been stopped before proceeding.

Copyright © 2023, 2026, Oracle and/or its affiliates.

3. Upgrade the RPM.
   As the `root` user, run the following command (for dbaastools 25.4.1 and later):

   ```
   /var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider
   GOOGLE --nodeList <node> --databasesStopped
   ```

   This command downloads and installs the latest RPM on the specified node.
4. Verify driver installation.
   Run the following command as the `root` user:

   ```
   /bin/rpm -qa | grep -i pkcs-multicloud-driver
   ```

5. Start the database instances.
   After the driver update, restart the database instances using:

   ```
   <oracle_home>/bin/srvctl start instance -node "<node>"
   ```

Non-Rolling update

In non-rolling mode, all databases on the cluster are stopped before applying the PKCS#11 driver update.

1. Stop databases using the Google Cloud KMS.
   Stop all databases that use the Google Cloud KMS for master encryption keys (MEK) before updating the driver:

   ```
   dbaascli database stop [--dbname <value>]
   ```

2. Verify that the databases have been shut down.
   Ensure that all databases using the Google Cloud KMS have been stopped before proceeding.
3. Update the RPM.
   As the `root` user, run the following command (for dbaastools 25.4.1 and later):

   ```
   /var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider
   GOOGLE --databasesStopped
   ```

   This command downloads and installs the latest RPM.
4. Verify driver installation.
   Run the following command as the `root` user:

   ```
   /bin/rpm -qa | grep -i pkcs-multicloud-driver
   ```

5. Restart the databases
   Once the driver has been updated, restart the databases:

   ```
   dbaascli database start [--dbname <value>]
   ```

Release Notes

Review the changes made across the different releases of the GCP driver.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Release 2.202 (260721.2348)
- Release 2.196 (260325.2123)
- Release 2.185 (251201.0551)

Release 2.202 (260721.2348)

- Various bug fixes and stability improvements.

Release 2.196 (260325.2123)

- Various bug fixes and stability improvements.

Release 2.185 (251201.0551)

- Various bug fixes and stability improvements.

AWS Key Management Service Integration for Exadata Database Service on Exascale Infrastructure@AWS

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

For Exadata Database Service on Exascale Infrastructure@AWS TDE MEKs can be stored in a file-based Oracle Wallet, Oracle Cloud Infrastructure (OCI) Vault, Oracle Key Vault (OKV), or AWS KMS, providing options to align with organization-specific security policies. Integration with AWS KMS enables applications, AWS services, and databases on VM Clusters to leverage a single centralized key management solution.

- Prerequisites
  Before using AWS KMS as the key management solution for your databases, you must complete the following steps.
- Understanding Key Pools
  Key Pools help you organize and control the AWS keys used to encrypt your Oracle Databases.
- Migrate from an Individual AWS KMS Key to a Key Pool
  Oracle AI Database@AWS supports Key Pools for AWS KMS keys. A Key Pool contains multiple AWS KMS keys and provides centralized control over the keys used for database encryption.
- Create a Database with an Individual AWS KMS Key Using dbaascli
- Using the Console to Manage AWS KMS Integration for Exadata Database Service on Exascale Infrastructure@AWS
  Learn how to manage AWS KMS integration for Exadata Database Service on Exascale Infrastructure@AWS.
- Using the API to Manage AWS KMS Integration for Exadata Database Service on Exascale Infrastructure@AWS
- Updating the Multicloud PKCS#11 Driver

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Release Notes
  Review the changes made across the different releases of the AWS driver.

Prerequisites

Before using AWS KMS as the key management solution for your databases, you must complete the following steps.

- Configure OCI Identity Domain
  OCI Identity domain is automatically configured during the Oracle AI Database@AWS onboarding process, and no action is required.
- Create or Use an Existing ODB Network
  Learn how to create or leverage an existing ODB network,
- Create an Exascale VM Cluster
- Configure Identity Provider
  Create a stack for each cluster. Use the following steps to create a CloudFormation stack. This needs to be performed for each Exascale VM Cluster.
- Associate an IAM Role to an Exascale VM Cluster
  You must attach an IAM role that you created previously to the Exascale VM Cluster to assign an identity connector that enables access to AWS resources.
- Create a Key
- Register AWS KMS Key
  To enable AWS KMS for your Exascale VM Clusters, you must first register the AWS KMS key in OCI.
- OCI IAM Policy Requirements
  Configure the required IAM policies based on the compartments containing the registered KMS resource and the Cloud VM cluster.

Configure OCI Identity Domain

OCI Identity domain is automatically configured during the Oracle AI Database@AWS onboarding process, and no action is required.

Complete the following steps only for accounts that were linked before the general availability of AWS KMS integration (August 2026).

Configure an OCI identity domain to enable AWS integration for your VM Clusters. It allows you to associate an AWS Identity and Access Management (IAM) service role with AWS integrations.

1. From the AWS console, select **Oracle AI Database@AWS**, and then select **Settings**.
2. Select the **Configure** button to configure your **OCI identity domain**.
3. Once it is complete, you can review the **Status**, **OCI identity domain ID** and **OCI identity domain URL** information from the **Settings** page.

Create or Use an Existing ODB Network

Learn how to create or leverage an existing ODB network,

1. If you do not have an existing ODB Network, you can provision one by following the ODB Network step-by-step instructions. If you have an existing ODB Network, you can modify it

Copyright © 2023, 2026, Oracle and/or its affiliates.

by selecting the **Modify** button. Follow Modify an ODB Network for step-by-step instructions.

2. From the **Configure service integrations** section:
   a. Select the **Security Token Service (STS)** option to setup the networking for AWS KMS and AWS STS access from the database.
   b. Select the **AWS KMS** checkbox to enable AWS KMS to use KMS keys in your authentication polices.

Create an Exascale VM Cluster

Exascale VM Cluster creation is only available through the AWS Console and AWS CLI. For more information, see Exascale VM Cluster.

Configure Identity Provider

Create a stack for each cluster. Use the following steps to create a CloudFormation stack. This needs to be performed for each Exascale VM Cluster.

When your CloudFormation stack creates an OIDC (OpenID Connect) Provider, it establishes a trust relationship so that an external identity source (like OCI) can authenticate to AWS. The associated IAM Role defines what that trusted identity is allowed to do.

1. Select your existing Exascale VM Cluster from the list to open its details page in the AWS Management Console.

> **Note**
>
> If you do not have an existing Exascale VM Cluster, you can provision one by following the Exascale VM Cluster step-by-step instructions.

2. Select the **IAM service roles** tab and then click the **CloudFormation** link.
3. From the Quick create stack page, review the pre-filled information.
   a. Under the **Parameters** section, parameters are defined in your template and allow you to enter custom values when you create or update a stack.
      i. Review your pre-filled **OCIIdentityDomainUrl** information.
      ii. The **OIDCProviderArn** field is optional. When you create the CloudFormation stack for the first time, it creates an OIDC Provider and an associated IAM Role. If you are executing the stack for the first time, you do not need to provide a value for the **OIDCProviderArn** field.
      iii. For any additional execution time, you must provide the ARN of the existing OIDC Provider and specify it in the **OIDCProviderArn** field. To obtain your OIDCProviderArn information, complete the following steps:
         i. From the AWS console, navigate to **IAM**, and then select **Identity providers**.
         ii. From the **Identity providers** list, you can filter your provider selecting **Type** as **OpenID Connect.**
         iii. Select the **Provider** link that was created when the CloudFormation stack was initially created.
         iv. From the **Summary** page, copy the **ARN** information.

Copyright © 2023, 2026, Oracle and/or its affiliates.

v. Paste the **ARN** information into the **OIDCProviderArn** field.

b. From the **Permissions** section, select IAM role name from the drop-down list, and then select the IAM role for CloudFormation to use for all operations performed on the stack.

c. Select the **Create stack** button to apply the changes.

4. After the CloudFormation stack deployment is complete, navigate to the **Resources** section of the stack and verify the created IdP and IAM Role resources. Copy the IAM Role ARN for future use.

Associate an IAM Role to an Exascale VM Cluster

You must attach an IAM role that you created previously to the Exascale VM Cluster to assign an identity connector that enables access to AWS resources.

1. From the AWS console, select **Oracle AI Database@AWS**.
2. From the left menu, select **Exascale VM clusters**, and then select your Exascale VM Cluster from the list.
3. Select the **IAM roles** tab, and then select the **Associate** button.
   a. The **AWS integration** field is read-only.
   b. Enter the Amazon Resource Name (ARN) of the IAM role you want to associate with the VM cluster in the **Role ARN** field. You can obtain the **ARN** information from the **Summary** section of the role that you previously created.
   c. Select the **Associate** button to attach the role.

> **Note**
>
> Once you associate an IAM role to your VM cluster, an identity connector gets attached to your Exascale VM Cluster.

Once you complete the prerequisites section, continue with the following actions:

1. From the **AWS Console**, create a customer-managed key in AWS KMS.
2. From the **OCI Console**, enable Exascale VM Clusters and databases to utilize the AWS KMS key created in the step 1.

Create a Key

1. From the AWS console, select **Key Management Service (KMS)**.
2. From the left menu, select **Customer managed keys**, and then select the **Create key** button.
3. In the **Configure key** section, enter the following information:
   a. Choose the **Symmetric** option as the **Key type**.
   b. Choose the **Encrypt and decrypt** option as the **Key usage**.
   c. Expand the **Advanced options** section. Both **AWS KMS** and **AWS CloudHSM** are supported.

Copyright © 2023, 2026, Oracle and/or its affiliates.

i. If you want to use KMS standard key store, choose the **KMS - recommended** option as **Key material origin**, and then choose either the **Single-Region key** option or the **Multi-Region key** option from the **Regionality** section.

ii. If you want to use **AWS CloudHSM**, choose the **AWS CloudHSM key store** option as **Key material origin**.

d. Select the **Next** button to continue the creation process.

4. In the **Add labels** section, enter the following information:

a. Enter a descriptive display name in the **Alias** field. Maximum 256 characters. Use alphanumeric and '_-/' characters.

> **Note**
>
> - The alias name cannot begin with `aws/`. The `aws/` prefix is reserved by AWS to represent AWS managed keys in your account.
> - An alias is a display name that you can use to identify the KMS key. We recommend that you choose an alias that indicates the type of data you plan to protect or the application you plan to use with the KMS key
> - Aliases are required when you create a KMS key in the AWS console.

b. The **Description** field is optional.

c. The **Tags** section is optional. You can use tags to categorize and identify your KMS keys and help you track your AWS costs.

d. Select the **Next** button to continue the creation process or the **Previous** button to return the previous page.

5. In the **Define key administrative permissions** section, complete the following substeps:

a. Search the role that you previously created, and then select the checkbox. Select the IAM users and roles that can administer the KMS key.

> **Note**
>
> This key policy gives the AWS account full control of this KMS key. It allows account administrators to use IAM policies to give other principals permission to manage the KMS key.

b. From the **Key deletion** section, the **Allow key administrators to delete this key** checkbox is selected by default. To prevent the selected IAM users and roles from deleting the KMS key, you can deselect the checkbox.

c. Select the **Next** button to continue the creation process or the **Previous** button to return the previous page.

6. In the **Define key usage permissions** section, complete the following substeps:

a. Search the role that you previously created, and then select the check box.

b. Select the **Next** button to continue the creation process or the **Previous** button to return the previous page.

7. In the **Edit key policy** section, complete the following substeps:

Copyright © 2023, 2026, Oracle and/or its affiliates.

a. From the **Preview** section, you can review the key policy. If you want to make a change, select the **Edit** tab.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "KMSKeyMetadata",
            "Effect": "Allow",
            "Principal": {
                "AWS": "<arn>"
            },
            "Action": [
                "kms:DescribeKey"
            ],
            "Resource": "*"
        },
        {
            "Sid": "KeyUsage",
            "Effect": "Allow",
            "Principal": {
                "AWS": "<arn>"
            },
            "Action": [
                "kms:Encrypt",
                "kms:Decrypt"
            ],
            "Resource": "*"
        }
    ]
}
```

b. Select the **Next** button to continue the creation process or the **Previous** button to return the previous page.

8. In the **Review** section, review your information and then select the Finish button.

For more information, see Create a symmetric encryption KMS key.

Register AWS KMS Key

To enable AWS KMS for your Exascale VM Clusters, you must first register the AWS KMS key in OCI.

1. From the OCI console, select **Oracle AI Database**, and then select **Database Multicloud Integrations**.
2. Select **AWS Integration**, and then select **AWS Keys**.
3. Select **Register AWS keys**, and then complete the following steps:
   a. From the **Compartment** dropdown list, select the compartment where your VM Cluster resides.
   b. Optionally, select a previously created key pool. For more information, see Create a Key Pool.
   c. If you select a key pool, ensure that its type matches the AWS key type:
      - A single-region AWS key must be added to a single-region key pool.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- A multi-region AWS key must be added to a multi-region key pool.

d. Under **AWS keys**, select your identity connector from the dropdown list.

> **Note**
>
> Ensure that the role associated with the connector has the `DescribeKey` permission on the key. This permission is required to successfully perform discovery. For more information, see step 2 in Configure Identity Provider.

e. Optionally, enter the key ARN in the **Key ARN** field.

f. Select **Discover**.

4. Review the discovered keys listed in the results, select the keys you want to register, and then select **Register**.

OCI IAM Policy Requirements

Configure the required IAM policies based on the compartments containing the registered KMS resource and the Cloud VM cluster.

- If the KMS resource and Cloud VM cluster are in different compartments, configure two policies:
  - In the compartment containing the KMS resource, allow the Cloud VM cluster resource principal to read the registered KMS resource:

    ```
    Allow any-user to read oracle-db-aws-keys in compartment id <key-compartment-ocid> where all { request.principal.type = 'exadbvmcluster' }
    ```

  - In the compartment containing the Cloud VM cluster, authorize the same resource principal for the PKCS data-object creation flow:

    ```
    Allow any-user to read oracle-db-aws-keys in compartment id <vm-cluster-compartment-ocid> where all { request.principal.type = 'exadbvmcluster' }
    ```

- If the KMS resource and Cloud VM cluster are in the same compartment, configure only one policy in that compartment:

  ```
  Allow any-user to read oracle-db-aws-keys in compartment id <vm-cluster-compartment-ocid> where all { request.principal.type = 'exadbvmcluster' }
  ```

- To create and manage Key Pools, you must have the following IAM policy:

  ```
  Allow any-user to manage oracle-db-aws-key-pools in tenancy
  ```

Understanding Key Pools

Key Pools help you organize and control the AWS keys used to encrypt your Oracle Databases.

Copyright © 2023, 2026, Oracle and/or its affiliates.

A key pool is a collection of customer-managed AWS keys used for database encryption. You create the pool in OCI, add AWS keys to it, and define whether those keys can be reused by more than one database. The pool provides a source of keys for database creation and key management operations.

Why use a Key Pool?

Without a key pool, AWS keys can be registered individually for use with databases. A key pool groups multiple AWS keys, allowing key rotation to select a key from the pool and providing centralized control over key usage.

With a key pool:

- Your organization controls which AWS keys are available.
- Different CDBs and PDBs can use separate key.
- Key selection and rotation can follow a defined reuse policy.
- You can review how keys are assigned and used.

How Key Pools work

First, create a key pool with a name, compartment, regional context, and key-reuse policy. You can create the pool before any AWS keys are added, allowing you to prepare the pool for future database deployments.

Next, add eligible AWS keys to the pool. A key can belong to only one pool, which helps ensure that key ownership and usage remain clear. You can also associate existing discovered keys with a pool when moving from the previous single-key encryption model.

When you provision a CDB, you first select a key pool and then a key from that pool. The CDB is created using the selected key for encryption.

PDB creation remains simple. You do not select a specific key when creating a PDB. The service automatically selects a key from the associated key pool for PDB creation and other supported key operations. Key selection from the key pool considers the pool's key-reuse policy, which determines whether a key can be reused by more than one database.

Key rotation follows the same pool-based approach. You initiate rotation as usual, and the service selects an appropriate key according to the policy associated with the pool. No PDB-specific key selection is required.

Managing and auditing a pool

You can view the keys in a pool, their usage information, and relevant lifecycle activity. Audit records help you track pool creation, policy changes, key additions and removals, key movement, and deletion activity.

Multi-region deployments

For supported multi-region and Data Guard configurations, Key Pools can be replicated across regions. Pool information and membership changes remain aligned so that standby environments can use the appropriate regional configuration.

Moving from the previous model

Existing single-key encryption configurations continue to operate while you prepare for migration. Guided migration capabilities help you verify prerequisites, associate existing keys with pools, and transition database environments with minimal disruption.

- Create a Key Pool
- Discover and Register Existing AWS Keys in a Key Pool

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Replicate a Key Pool Across Regions
- Add or Remove a Key from a Key Pool
- View Key Usage Count

Create a Key Pool

1. From the OCI Console, navigate to **Oracle AI Database** and then select **Database Multicloud Integrations**.
2. From the left menu, select **AWS Integration** and then select **Key Pools**.
3. Select the **Create key pool** button to create one.
   a. From the Create key pool page, complete the following steps.
   b. From the dropdown list, select the **Compartment**.
   c. Enter a descriptive name for your key pool. The Display name must contain only alphanumeric characters and hyphens (-)and underscores (_).
   d. The **Allow key reuse** toggle is enabled by default. Enable Allow key reuse to allow keys within the key pool to be reused for multiple databases. If disabled, one key can only be used for one database.
   e. The Multi-region key pool toggle is disabled by default. This is optional step. Enable Multi-region key pool to allow only Multi-region keys to be added to the key pool. If disabled, only single-region keys can be added to the key pool.
   f. The **Tags** section is an optional step. If you want to add a tag your resources, select the Add tag button.
4. Once the key pool is created, navigate back to the Key Pools page to view the State, Allow key reuse, Multi-region key pool and Created information.

Discover and Register Existing AWS Keys in a Key Pool

1. From the OCI Console, navigate to **Oracle AI Database**, and then select **Database Multicloud Integrations**.
2. From the left menu, select **AWS Integration** and then select **Key Pools**.
3. From the Key Pools list, select the key pool that you created previously.
4. Select the **Keys** tab. If you have not registered your AWS Key(s), register AWS Key(s) by selecting Register AWS keys button.
5. From the Register AWS keys page, complete the following steps:
   a. Select the **Compartment** which your AWS Key will be registered in.
      - The discovered keys will be added to the current key pool.
      - If the key pool selected is a multi-region key pool, only the multi-region keys from the discovered results are enabled and can be added to this key pool.
   b. From the Discover AWS keys using identity connector list, select an identity connector.
   c. Enter your AWS key ARN information in the **Key name** field.
   d. Select the **Discover** button.
   e. Expand the Advanced options section. This is an optional step. If you want to add a tag your resources, select the **Add tag** button.

Copyright © 2023, 2026, Oracle and/or its affiliates.

f. Review the discovered keys listed in the results, select the keys you want to add to the Key Pool, and then select **Register**.

6. Select the **Add keys to key pool** button to add your AWS Key(s) to your Key Pool.
   a. From the Add keys to key pool page, select your key(s).
   Select one or more existing AWS keys to add to the key pool. The list displays only keys that are in the current compartment and are not associated with another key pool.
   b. Select the **Add button**.

Replicate a Key Pool Across Regions

1. From the OCI Console, navigate to **Oracle AI Database**, and then select **Database Multicloud Integrations**.
2. From the left menu, select **AWS Integration** and then select **Key Pools**.
3. In the Key Pools list, select the multi-Region key pool that you want to replicate.
4. From the **Actions** menu, select **Replicate Key Pool**.
5. In the Replicate Key Pool window, select the destination regions where you want to replicate the key pool and its keys.
6. Select **Replicate**.
7. Select the **Cross-region replications** tab to view the regions where the key pool has been replicated.

Add or Remove a Key from a Key Pool

1. From the OCI Console, navigate to **Oracle AI Database**, and then select **Database Multicloud Integrations**.
2. From the left menu, select **AWS Integration** and then select **Key Pools**.
3. In the Key Pools list, select the **Key Pool** that you created earlier.
4. To add a key to the Key Pool:
   a. Select the **Keys** tab, and then select **Add keys to key pool**.
   b. On the Add keys to key pool page, select a key from the dropdown list.

   > **Note**
   > - Only keys in the current compartment that are not assigned to another Key Pool are displayed.
   > - For a single-Region Key Pool, only single-Region keys are displayed.
   > - For a multi-Region Key Pool, only multi-Region keys are displayed.

   c. Select **Save**.
5. To remove a key from the Key Pool:
   a. Select the key that you want to remove.
   b. On the Details page, scroll to the **Key Pool** field, and select **Delete**.
   c. Enter DELETE in the confirmation field, and then select **Delete**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

The key is removed from the Key Pool.

View Key Usage Count

1. From the OCI Console, navigate to **Oracle AI Database**, and then select **Database Multicloud Integrations**.
2. From the left menu, select **AWS Integration** and then select **Key Pools**.
3. In the Key Pools list, select the Key Pool that you created earlier.
4. View the Key usage count based on the type of Key Pool:
   - Single-Region Key Pool: On the Details page, scroll to the **Key usage count** field.
   - Multi-Region Key Pool: Select the **Keys** tab, select the AWS KMS key that you want to review, and then scroll to the **Key usage count** field on the Details page.

Available Keys Below Threshold

The **Key usage count** field shows the number of active database resources currently using the AWS key. This count reflects key usage and does not indicate the number of available keys in the Key Pool.

The system displays the **Available Keys Below Threshold** warning when key reuse is false and fewer than 25% of the keys in the Key Pool are available for use. If the warning appears, add more keys to the Key Pool or enable key reuse in the Key Pool policy to help ensure that keys remain available for database operations.

Migrate from an Individual AWS KMS Key to a Key Pool

Oracle AI Database@AWS supports Key Pools for AWS KMS keys. A Key Pool contains multiple AWS KMS keys and provides centralized control over the keys used for database encryption.

With a Key Pool:

- Your organization controls which AWS KMS keys are available for database encryption.
- Different CDBs and PDBs can use separate keys.
- Key selection during rotation can follow the Key Pool's reuse policy.
- You can review how keys are assigned and used.

Before You Begin

- Review the Prerequisites for required setup, permissions, and configuration.
- Review the Release Notes for driver requirements and release-specific information.
- Ensure that your VM Cluster uses AWS multicloud PKCS#11 driver release 2.203 (260927.1348), which includes Key Pool support. Update the driver if required. See Updating the Multicloud PKCS#11 Driver.
- Ensure that the VM Cluster has permission to access every AWS KMS key that you add to the Key Pool.
- Plan a maintenance window for the driver update and the database restart required by the migration command.

Migrate a Database to a Key Pool

1. Create a Key Pool.

Copyright © 2023, 2026, Oracle and/or its affiliates.

2. Add the AWS KMS keys that you want to use to the Key Pool.
3. Grant the VM Cluster permission to access every AWS KMS key in the Key Pool.
4. Connect to the VM Cluster and run the following command for the database that you want to migrate:

```
dbaascli tde setPrimaryHSMKey \
  --dbname <db_name> \
  --awsKeyPoolOCID <key_pool_ocid> \
  --bounceDatabase
```

Replace the following placeholders:

- *<db_name>* with the name of the database.
- *<key_pool_ocid>* with the OCID of the Key Pool.

The `--bounceDatabase` option restarts the database.

Create a Database with an Individual AWS KMS Key Using dbaascli

> **Note**
>
> Although you can create a database using an individual AWS KMS key, Oracle strongly recommends using a Key Pool.

If you prefer to use an individual AWS KMS key instead of a Key Pool, create the database by specifying the AWS KMS key OCID.

```
dbaascli database create \
  --dbName <db_name> \
  --oracleHome <db_home_path> \
  --oracleHomeName <db_home_name> \
  --tdeConfigMethod AWS \
  --tdeKeyOCID <key_ocid>
```

Replace the following placeholders:

- *<db_name>* with the name of the database.
- *<db_home_path>* with the path to the Oracle Home.
- *<db_home_name>* with the name of the Oracle Home.
- *<key_ocid>* with the OCID of the AWS KMS key.

This method does not require the driver upgrade for Key Pool support.

Existing databases configured with individual AWS KMS keys continue to operate while you prepare to adopt Key Pools.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Using the Console to Manage AWS KMS Integration for Exadata Database Service on Exascale Infrastructure@AWS

Learn how to manage AWS KMS integration for Exadata Database Service on Exascale Infrastructure@AWS.

- Enable or Disable AWS Key Management
  When you enable AWS key management for your database, only AWS keys that are authorized for use with VM Cluster and registered with OCI can be used.
- Create a Database and Use AWS KMS as the Key Management Solution
  This topic describes only the steps for creating a database and using AWS KMS as the key management solution.
- Steps to Change the Key Management of Existing Databases from Oracle Wallet to AWS KMS
- Rotate the AWS KMS Key of a Container Database (CDB)
  To rotate the AWS KMS Key of a container database (CDB), use this procedure.
- Rotate the AWS KMS Key of a Pluggable Database (PDB)
  To rotate the AWS KMS Key of a pluggable database (PDB), use this procedure.
- (Optional) Disable a CDB or PDB Key
  Optionally, complete the following steps to disable a specific CDB or PDB key.

Enable or Disable AWS Key Management

When you enable AWS key management for your database, only AWS keys that are authorized for use with VM Cluster and registered with OCI can be used.

1. From the OCI console, select **Oracle AI Database** and then select **Oracle Exadata Database Service on Exascale Infrastructure**.
2. From the left menu, select **VM Clusters**, and then select your VM Cluster.
3. Select the **VM Cluster information** tab, and select the **Enable** button next to **AWS Customer Managed Encryption Key**.

> **Note**
>
> Once you enable AWS Key Management for your VM Cluster, you can disable it using the **Disable** button. Disabling this feature will affect the availability of databases that use AWS Key Management Service for encryption and decryption operations. Ensure that no database is currently using AWS key management before disabling it at VM Cluster level.

Create a Database and Use AWS KMS as the Key Management Solution

This topic describes only the steps for creating a database and using AWS KMS as the key management solution.

For the generic database creation procedure, see To Create a Database in an Existing VM Cluster.

Prerequisites

Copyright © 2023, 2026, Oracle and/or its affiliates.

Refer to Prerequisites.

Steps

If AWS KMS key management is enabled at the VM cluster level, you will have two key management options: **Oracle Wallet** and **AWS key management**.

1. In the **Encryption** section, choose **AWS key management**.
2. Select the **Compartment** and the **Key pool** from the dropdown list.
3. Select the **Compartment** and the **Key** from the dropdown list.

> **Note**
>
> - Only registered keys are listed.
> - If your desired key is not visible, it may not have been registered yet. Click **Register AWS Keys** to discover and register it.
>   For detailed instructions, refer to Register AWS KMS Key.
> - You can select a key alias from the drop-down list. If no key alias is available, the drop-down list will display the key ID.

Steps to Change the Key Management of Existing Databases from Oracle Wallet to AWS KMS

If you want to change the key management of your existing databases in the VM Cluster from Oracle Wallet to AWS KMS, complete the following steps.

1. From your **VM Clusters**, navigate to **Container Databases** tab, and then select the database that you are using.
2. From the **Encryption** section, confirm that **Key management** is set to **Oracle Wallet**, and then select the **Change** link.
3. From the **Change key management** page, enter the following information.
   a. Select your **Key management** as **AWS key management** from the dropdown list.
   b. Select the key pool compartment you are using, and then select the desired key pool from the dropdown list.
   c. Select the **Save changes** button.

Rotate the AWS KMS Key of a Container Database (CDB)

To rotate the AWS KMS Key of a container database (CDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on 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 that you want to rotate encryption keys.
4. Click **Container Databases**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

5. Click the name of the database that you want to rotate encryption keys.
The Database Details page displays information about the selected database.
6. In the **Encryption** section, verify that the **Key Management** is set to **AWS key management**, and then click the **Rotate** link.
7. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

> **Note**
>
> Rotating the AWS KMS key generates a new encryption context for the same key.

Rotate the AWS KMS Key of a Pluggable Database (PDB)

To rotate the AWS KMS Key of a pluggable database (PDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on 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 PDB you want to start, and then click its name to display the details page.
4. Under **Container Databases**, find the database containing the PDB you want to rotate encryption keys.
5. Click the name of the database to view the Database Details page.
6. Click **Pluggable Databases** in the **Resources** section of the page.
A list of existing PDBs in this database is displayed.
7. Click the name of the PDB that you want to rotate encryption keys.
The pluggable details page is displayed.
8. In the **Encryption** section displays that the Key management is set as AWS key management.
9. Click the **Rotate** link.
10. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

> **Note**
>
> Rotating the AWS KMS key generates a new encryption context for the same key.

(Optional) Disable a CDB or PDB Key

Optionally, complete the following steps to disable a specific CDB or PDB key.

1. Navigate to the AWS console, select **Key Management Service (KMS)**.
2. From the left menu, select **Customer managed keys**, and then select the **Key ID** of the key that you want to edit.
3. Select the **Edit policy** button.
4. Run the following query to obtain the **EncryptionContext MKID** of the key.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Use the following command to view the current master key IDs (MKIDs) for enabling or disabling operations:

```
dbaascli tde getHSMKeys --dbname <DB-Name>
```

You can also run the following SQL query to list all key versions:

```
SELECT key_id, con_id, creation_time, key_use FROM v$encryption_keys;
```

5. Add a deny policy using the retrieved **EncryptionContext** MKID as shown below:

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt"
  ],
  "Resource": "arn:aws:kms:us-east-1:867344470629:key/7139075d-
a006-4302-92d5-48ecca31d48e",
  "Condition": {
    "StringEquals": {
      "kms:EncryptionContext:MKID":
"ORACLE.TDE.HSM.MK.06AE5EFB528D9D4F21BFEDA63C6C8738D9"
    }
  }
}
```

6. Choose **Save** to apply the policy update.

Using the API to Manage AWS KMS Integration for Exadata Database Service on Exascale Infrastructure@AWS

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

The following resources will be made available to customers through OCI SDK, CLI, and Terraform. These APIs will be used by customers who wish to integrate Oracle AI Database on Exascale Infrastructure with AWS KMS.

- OracleDbAwsIdentityConnector
- OracleDbAwsKey

OracleDbAwsIdentityConnector

Table 5-14 OracleDbAwsIdentityConnector

API	Description
`ListOracleDbAwsIdentityConnectors`	Lists all AWS Identity Connector resources based on the specified filters.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 5-14 (Cont.) OracleDbAwsIdentityConnector

API	Description
`GetOracleDbAwsIdentityConnector`	Retrieves detailed information about a specific AWS Identity Connector resource.
`CreateOracleDbAwsIdentityConnector`	Creates a new AWS Identity Connector resource for the specified ExaDB-XS VM Cluster.
`UpdateOracleDbAwsIdentityConnector`	Updates the configuration details of an existing AWS Identity Connector resource.
`ChangeOracleDbAwsIdentityConnectorCompartment`	Moves the AWS Identity Connector resource to a different compartment.
`DeleteOracleDbAwsIdentityConnector`	Deletes the specified AWS Identity Connector resource.
`RefreshOracleDbAwsIdentityConnector`	Refreshed the configuration details of the specified AWS Identity Connector resource.

OracleDbAwsKey

Table 5-15 OracleDbAwsKey

API	Description
`ListOracleDbAwsKeys`	Lists all AWS Key resources based on the specified filters.
`CreateOracleDbAwsKey`	Creates a new AWS Key resource.
`ChangeOracleDbAwsKeyCompartment`	Moves the AWS Key resource to a different compartment.
`GetOracleDbAwsKey`	Retrieves detailed information about a specific AWS Key resource.
`UpdateOracleDbAwsKey`	Updates the configuration details of an existing AWS Key resource.
`DeleteOracleDbAwsKey`	Deletes the specified AWS Key resource.
`RefreshOracleDbAwsKey`	Refreshes the configuration details of a AWS Key resource.

Updating the Multicloud PKCS#11 Driver

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> The PKCS#11 driver for AWS is installed automatically when AWS key management is enabled on a VM cluster. To verify that the driver is installed and check its version, run the following command:
>
> ```
> rpm -qa | grep pkcs
> ```
>
> This command lists the PKCS package installed on the VM cluster. Only one PKCS#11 driver can be active at a time, corresponding to either Azure, Google Cloud, or AWS.
>
> Example output:
>
> ```
> pkcs-multicloud-driver-maz-0.2-250814.1708.x86_64
> ```
>
> This confirms that the `pkcs-multicloud-driver-maz` package is installed, along with its version and architecture.

The PKCS#11 driver can be updated using either a rolling or non-rolling mechanism, depending on customer preference.

In rolling mode, the update is applied to one VM at a time, and the databases on that VM are restarted before proceeding to the next. In non-rolling mode, all databases in the cluster are shut down, the update is applied across all nodes, and the databases are then brought back online. Users must be aware of the `dbaastools` version in use, as subsequent steps, specifically step 3 in both scenarios, require different actions depending on the version, particularly when updating the PKCS#11 driver.

The version of `dbaastools` can be determined by running the following command:

```
/var/opt/oracle/dbaascli/dbaascli admin showLatestStackVersion
```

Rolling update

On each Guest VM, apply the PKCS#11 driver update as follows (one node at a time):

1. Shutdown the database instances that are using the AWS KMS.
   All databases that use the AWS KMS for master encryption keys (MEK) must be stopped before updating the driver. Use:

   ```
   srvctl stop instance -d <db_unique_name> -i <instance_name> [-o
   <stop_option>]
   ```

   - `-d <db_unique_name>`: The unique name of the database (same as `DB_UNIQUE_NAME` in the database).
   - `-i <instance_name>`: The name of the instance to stop, for example, `orcl1`.

   For details on additional options, refer to the SRVCTL command-line help or reference manual.
2. Verify that the databases have been shut down.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Ensure that all database instances that are using the AWS KMS have been stopped before proceeding.

3. Update the RPM.
   As the `root` user, run the following command (for dbaastools 25.4.1 and later):

   ```
   /var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider AWS
   --nodeList <node> --databasesStopped
   ```

   This command downloads and installs the latest RPM on the specified node.

4. Verify driver installation.
   Run the following command as the `root` user:

   ```
   /bin/rpm -qa | grep -i pkcs-multicloud-driver
   ```

5. Start the database instances.
   After the driver update, restart the database instances using:

   ```
   <oracle_home>/bin/srvctl start instance -node "<node>"
   ```

Non-Rolling update

In non-rolling mode, all databases on the cluster are stopped before applying the PKCS#11 driver update.

1. Stop databases using the AWS KMS.
   Stop all databases that use the AWS KMS for master encryption keys (MEK) before updating the driver:

   ```
   dbaascli database stop [--dbname <value>]
   ```

2. Verify that the databases have been shut down.
   Ensure that all databases using the AWS KMS have been stopped before proceeding.
3. Upgrade the RPM.
   As the `root` user, run the following command (for dbaastools 25.4.1 and later):

   ```
   /var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider AWS
   --databasesStopped
   ```

   This command downloads and installs the latest RPM.

4. Verify driver installation.
   Run the following command as the `root` user:

   ```
   /bin/rpm -qa | grep -i pkcs-multicloud-driver
   ```

5. Restart the databases.
   Once the driver has been updated, restart the databases:

   ```
   dbaascli database start [--dbname <value>]
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

Release Notes

Review the changes made across the different releases of the AWS driver.

- Release 2.203 (260927.1348)
- Release 2.202 (260721.2348)
- Release 2.196 (260325.2123)
- Release 2.185 (251201.0551)

Release 2.203 (260927.1348)

- Support for keypool.
- Various bug fixes and stability improvements.

Release 2.202 (260721.2348)

- Various bug fixes and stability improvements.

Release 2.196 (260325.2123)

- Various bug fixes and stability improvements.

Release 2.185 (251201.0551)

- Various bug fixes and stability improvements.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.
- EU Data Act Compliance
  Oracle supports the goal of fair and transparent data sharing and access within the European Union.
- Oracle Cloud Infrastructure Operations Insights
  Oracle Cloud Infrastructure Operations Insights allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases deployed in Exadata Cloud Infrastructure.
- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

EU Data Act Compliance

Oracle supports the goal of fair and transparent data sharing and access within the European Union.

To learn more about how to export data from Oracle Cloud Infrastructure (OCI) Database as a Service (DBaaS) to on-premises systems or other cloud environments, refer to our compliance page in the base database service documentation.

Related Topics

- Export Data from OCI Database Services

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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@dg11lrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
      showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file location:
      /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
AM_46941.log

############
List of Available Database Images
#############

17.IMAGE_TAG=18.17.0.0.0
   VERSION=18.17.0.0.0
   DESCRIPTION=18c JAN 2022 DB Image

18.IMAGE_TAG=19.10.0.0.0
   VERSION=19.10.0.0.0
   DESCRIPTION=19c JAN 2021 DB Image

19.IMAGE_TAG=19.11.0.0.0
   VERSION=19.11.0.0.0
   DESCRIPTION=19c APR 2021 DB Image

20.IMAGE_TAG=19.12.0.0.0
  VERSION=19.12.0.0.0
  DESCRIPTION=19c JUL 2021 DB Image

21.IMAGE_TAG=19.13.0.0.0
  VERSION=19.13.0.0.0
  DESCRIPTION=19c OCT 2021 DB Image

Images can be downloaded using their image tags. For details, see help using
'dbaascli cswlib download --help'.
dbaascli execution completed
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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 23.26.0.0.0
   ```

   > **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
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.
>
> – The dbaascli utility is only compatible with provisioning Oracle AI Database 26ai on Oracle Exadata Database Service on Exascale Infrastructure. Oracle recommends that you use the OCI console-based API rather than dbaascli for provisioning Oracle Database 19c on this service.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home Path
   ```

   Where:

   - `--dbName` specifies the name of the database
   - `--oracleHome` specifies Oracle home location

   To create a CDB, run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home Path
   ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

To create a non-CDB, run the following command:

```
dbaascli database create --dbName database name --oracleHome Oracle Home
Path --createAsCDB false
```

When prompted, enter the `sys` and `tde` passwords.

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli database create`.

- Running Prerequisite Checks Prior to Creating Oracle Database
  To run prerequisites 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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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 database passwords:

1. To change the SYS password for a database:
   a. Log in to an Exadata Database Service on Exascale Infrastructure virtual machine as the `opc` user.

   b. Run the following command:

   ```
   sudo dbaascli database changepassword --dbname <database_name> --user SYS
   ```

2. To change the SYS password in a Data Guard environment:
   a. On the primary database, run:

   ```
   dbaascli database changePassword --dbName <dbname> --user SYS --
   prepareStandbyBlob true --blobLocation <location_to_create_blob_file>
   ```

   b. Copy the generated blob file to all standby databases and change the file ownership to the `oracle` user.

   c. On each standby database, run:

   ```
   dbaascli database changePassword --dbName <dbname> --user SYS --
   standbyBlobFromPrimary <location_of_copied_blob_file>
   ```

3. To change the TDE wallet password for a database:
   a. Log in to an Exadata Database Service on Exascale Infrastructure virtual machine as the `opc` user.

   b. 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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@dg11lrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
      showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file location:
      /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
AM_46941.log

############
List of Available Database Images
#############

17.IMAGE_TAG=18.17.0.0.0
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
   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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **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 `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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Example: Log collection for all databases

```
# 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
```

Example: Log collection in background mode

```
# 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.

Example: Log collection for a specified database

```
# 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-
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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
```

Example: Log collection to a specified Guest VM directory

```
# 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=8a26cffcfdd72c261660d4f736c615981856e35
7749d90751b94f3eda19a9a70}"
}
dbaascli execution completed
```

Example: Log collection for a specified timeframe

```
# 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
-----------------
..
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
---------- 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
```

Example: Log collection for specific nodes

```
# 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}"
}
```

Example: Log collection for specific components

```
# 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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

Example: Saving collected logs to Object Storage using a PAR URL

```
# dbaascli diag collect --objectStoreBucketUri https://objectstorage.us-
phoenix-1.oraclecloud.com/p/aL-IbIKQ1j6lWNftJc2rLoLh6o9bJgbZm8z0S--
BeVuXaipSEEMISrSCfFrVEolG/n/intexadatateam/b/diag_collect_test/o/
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --objectStoreBucketUri https://
objectstorage.us-phoenix-1.oraclecloud.com/p/aL-
IbIKQ1j6lWNftJc2rLoLh6o9bJgbZm8z0S--BeVuXaipSEEMISrSCfFrVEolG/n/
intexadatateam/b/diag_collect_test/o/
Job id: 028151b7-cbc4-409a-9ec6-69affe10f3bb
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-51-36-
PM_2963.log
Loading PILOT...
Session ID of the current execution is: 20
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-51-40-
PM_3555
-----------------
..
---------- DIAG COLLECT PLUGIN RESULT ----------
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2151/
diag_cloudlogs_20231214-2151_node1.zip=71633e13ccd06de15cb26850bb0266cf0d869e2
59550515c5b1fb734c487b470}"
}
dbaascli execution completed
```

Related Topics

- dbaascli diag collect
  To collect diagnostics, use the `dbaascli diag collect` command.

Performing a Health Check Examples

Use dbaascli `dbaascli diag healthcheck` command to perform a health check on all system nodes.

See *dbaascli diag healthcheck* for the syntax details in the dbaascli Command Reference.

Example: Health check information

```
# 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
```

Example: Saving health check information to a specified domU directory

```
# dbaascli diag healthcheck --destLocation /tmp/test
DBAAS CLI version MAIN
Executing command diag healthcheck --destLocation /tmp/test
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
INFO: Starting diag healthcheck
INFO: Collected diag logs at: /tmp/test/diag_cloudlogs_20210322-2250.tar.gz
```

Example: Health check information for specific nodes

```
# 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
```

Example: Saving health check information to Object Storage using a PAR URL

```
# 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*.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 ment:

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

```
dbaascli 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

  ```
  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>
  ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascl, 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> ExaDB-XS supports patching only through the OCI Console or supported cloud-managed patching workflows. Manual patching using `dbaascli` commands is not supported.

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 recover
  To recover a database, use the `dbaascli database recover` command.
- dbaascli database createTemplate
  To create database templates (DBCA templates) that can subsequently be used to create databases, use the `dbaascli database createTemplate` command.
- dbaascli database start
  To start an Oracle Database, use the `dbaascli database start` command.
- dbaascli database status
  To check the status of an Oracle Database, use the `dbaascli database status` command.
- dbaascli database stop
  To stop an Oracle Database, use the `dbaascli database stop` command.
- dbaascli database upgrade
  To upgrade an Oracle Database, use the `dbaascli database upgrade` command.
- 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.
- dbaascli dataguard updateDGConfigAttributes
  To update Data Guard automation attributes across all the cluster nodes, use the `dbaascli dataguard updateDGConfigAttributes` command.
- dbaascli dbhome create
  To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.
- dbaascli dbHome delete
  To delete a given Oracle Database home, use the `dbaascli dbHome delete` command.
- dbaascli dbhome getDatabases
  To view information about all Oracle Databases running from a given database Oracle home, use the `dbaascli dbHome getDatabases` command. Specify either the Oracle home location or Oracle home name.
- dbaascli dbHome getDetails
  To view information about a specific Oracle home, use the `dbaascli dbHome getDetails` command. Specify either the Oracle home location or Oracle home name.
- dbaascli dbimage purge
  The `dbimage purge` command removes the specified software image from your Oracle Exadata Database Service on Exascale Infrastructure environment.
- dbaascli diag collect
  To collect diagnostics, use the `dbaascli diag collect` command.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- dbaascli diag healthCheck
  To run diagnostic health checks, use the `dbaascli diag healthCheck` command.
- dbaascli grid configureTCPS
  To configure TCPS for the existing cluster, use the `dbaascli grid configureTCPS` 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 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 getConnectString
  To display Oracle Net connect string information for a pluggable database (PDB) run the `dbaascli pdb getConnectString` command.
- dbaascli pdb create
  To create a new pluggable database (PDB), use the `dbaascli pdb create` command.
- dbaascli pdb delete
  To delete a pluggable database (PDB) run the `dbaascli pdb delete` command.
- dbaascli pdb getDetails
  To view details of a pluggable database (PDB), use the `dbaascli pdb getDetails` command.
- dbaascli pdb list
  To view the list of pluggable databases (PDB) in a container database, use the `dbaascli pdb list` command.
- dbaascli pdb localClone
  To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the `dbaascli pdb localClone` command.
- dbaascli pdb open
  To open a pluggable database (PDB), use the `dbaascli pdb open` command.
- dbaascli pdb recover
  To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.
- dbaascli pdb refresh
  To refresh a specified pluggable database (PDB), use the `dbaascli pdb refresh` command.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
############ 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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>]
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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> [--json <value>]
            | --getBackupReport
                {
                    --tag <value>
                    | --latest
                }
                --json <value>
            | --configure
                {
                    --configFile <value>
                    | --enableRTRT
                    | --disableRTRT
                    | --disableCatalog
                    | --deleteImmutableConfiguration
                }
            | --getConfig
                {
                    [--configFile <value>]
                    | [--showOldParams]
                }
            | --validate [--untilTime <value>]
            | --showHistory [--all]
            | --getSchedules
        }
```

Where:

Copyright © 2023, 2026, Oracle and/or its affiliates.

```
--dbname specifies Oracle Database name
--list returns database backup information
            [--backupType | --json]
            [--backupType specifies backupType (REGULAR-L0 | REGULAR-L1 | ARCHIVELOG |
LONGTERM). ]
            [--json specifies file Name for JSON output. ]
--start begins database backup.
            [--level0 creates a Level-0 (full) backup. ]
            [--level1 creates a Level-1 (incremental) backup. ]
            [--archival | --archivelog]
            [--archival creates an archival full backup. ]
                --tag specifies backup tag.
            [--archivelog ]
--delete deletes Archival backup.
            --backupTag specifies backup tag to delete.
--status displays the details about a backup job process.
            --uuid unique identifier of the backup operation. Input format: xxxxxxxx-
xxxx-xxxx-xxxx-xxxxxxxxxxxx.
                [--json specifies file Name for JSON output. ]
--getBackupReport returns BackupReport.
            --tag | --latest
            --tag specifies backup tag.
            --latest returns latest backup report (all types of database backup).
            --json specifies file Name for JSON output.
--configure configures database for backup.
            --configFile | --enableRTRT | --disableRTRT | --disableCatalog | --
deleteImmutableConfiguration
            --configFile specifies database backup configuration file.
            --enableRTRT enables Real Time Redo Transport.
            --disableRTRT disables Real Time Redo Transport.
            --disableCatalog disables recovery catalog.
            --deleteImmutableConfiguration
--getConfig returns database backup configuration.
            [--configFile | --showOldParams]
            [--configFile specifies database backup configuration file. ]
            [--showOldParams returns old parameter names of backup
configuration. ]
--validate validates that backups are complete and corruption-free.
            [--untilTime validates from closest Level-0 (full) backup until time
provided. Input format: DD-MON-YYYY HH24:MI:SS.]
--showHistory displays history of backup operations.
            [--all displays all backup operations. ]
--getSchedules returns all backup schedules for a given database.
```

> **Note**
>
> `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 dbaascli database backup command?

A: The `dbaascli 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 dbaascli database backup command?

A: The command must be run as the `root` user to have the necessary permissions for backup management.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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):

Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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
  ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli 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 dbaascli 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 dbaascli 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:

```
dbaascli database bounce --dbname <value> --rolling true
```

Q: Is it safe to run the dbaascli 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 `dbaascli 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli database bounce

```
dbaascli database bounce --dbname dbname
```

dbaascli database changepassword

To change the password of a specified Oracle Database user, use the `dbaascli database changePassword` command. When prompted enter the user name for which you want to change the password and then enter the password.

Prerequisites

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli database changePassword [--dbname <value>] [--user <value>]
{
  [--prepareStandbyBlob <value> [--blobLocation <value>]] | [--
standbyBlobFromPrimary <value>]
}
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the Oracle Database that you want to act on
- `--user` specifies the user name whose password change is required
- `--prepareStandbyBlob` specifies `true` to generate a blob file containing the artifacts needed to change the password in a Data Guard environment. Valid values: `true|false`
- `--blobLocation` specifies the custom path where blob file will be generated
- `--standbyBlobFromPrimary` specifies the standby blob file, which is prepared from the primary database
- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID

Frequently Asked Questions

Q: What does the dbaascli database changePassword command do?

A: The `dbaascli 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 dbaascli 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

A: Use the `--dbname` option to specify the name of the Oracle Database you want to act on. For example:

```
dbaascli 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:

```
dbaascli database changePassword --user <username>
```

Q: What is the purpose of the --prepareStandbyBlob option in the dbaascli 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:

```
dbaascli 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 dbaascli 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 dbaascli 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 dbaascli 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli database changePassword

```
dbaascli database changepassword --dbname db19
```

dbaascli database convertToPDB

To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.

Syntax

```
dbaascli database convertToPDB --dbname <value> [--cdbName <value>] [--
executePrereqs] [--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
- `--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 dbaascli database convertToPDB

To run pre-conversion prechecks:

```
dbaascli database convertToPDB --dbname ndb19 --cdbname cdb19 --
backupPrepared --executePrereqs
```

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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 dbaascli database create command do?

A: The `dbaascli 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:

```
dbaascli 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:

```
dbaascli 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

```
dbaascli 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:

```
dbaascli 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`).

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli database create

```
dbaascli database create --dbName db19 --oracleHomeName myhome19 --dbSid
db19sid --nodeList node1,node2 --createAsCDB true
```

dbaascli database delete

To delete an Oracle Database, use the `dbaascli database delete` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database delete --dbname <value>
[--deleteArchiveLogs <value>]
[--deleteBackups <value>]
[--precheckOnly <value>]
[--waitForCompletion <value>]
[--force]
[--dbSID <value>]
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the database.
- `--deleteArchiveLogs` specifies `true` or `false` to indicate deletion of database archive logs.
- `--deleteBackups` specifies `true` or `false` to indicate deletion of database backups.
- `--precheckOnly` specifies `yes` to run only the prechecks for this operation. Valid values: `yes` or `no`.
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`.
- `–-force` flag to force delete database.
- `--dbSID` specify database SID.
- `--resume` to resume the previous execution.
- `--sessionID` to resume a specific session id.

Frequently Asked Questions

Q: What is the purpose of the dbaascli database delete command?

Copyright © 2023, 2026, Oracle and/or its affiliates.

A: The `dbaascli 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:

```
dbaascli 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:

```
dbaascli 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:

```
dbaascli 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:

```
dbaascli 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:

```
dbaascli 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:

```
dbaascli 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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>]
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
[--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`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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 dbaascli database duplicate command?

A: The `dbaascli 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 dbaascli 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>
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521/
mynew.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com --oracleHome /u02/app/
oracle/product/19.0.0.0/dbhome_2 --dbName newdup --
sourceDBWalletLocation /var/opt/oracle/dbaas_acfs/tmp/prim_wallet
```

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>
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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 `oracle`user.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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*:

  ```
  dbaascli database recover --dbname myTestDb --status --uuid
  2508ea18be2911eb82d0020017075151
  ```

dbaascli database createTemplate

To create database templates (DBCA templates) that can subsequently be used to create databases, use the `dbaascli database createTemplate` command.

Prerequistes:

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 parallelsim 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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 dbaascli 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 dbaascli database start

```
dbaascli database start --dbname dbname --mode mount
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

dbaascli database status

To check the status of an Oracle Database, use the `dbaascli database status` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli database status
[--service][--dbname]
[--user]
[--password]
```

Where:

- `--service` specifies the name of the service
- `--dbname` specifies the name of the database
- `--user` specifies the user name of the service
- `--password` specifies the password of the user

Output from the command includes the open mode of the database, the software release and edition of the database, and release version of other software components.

Example 6-20 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli 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 dbaascli database stop command?

A: No, specifying a `shutdown` mode is optional. If not provided, the default immediate mode will be used.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Example 6-21 dbaascli database stop

```
dbaascli database stop --dbname db19
```

dbaascli database upgrade

To upgrade an Oracle Database, use the `dbaascli database upgrade` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database upgrade --dbname <value>
{--targetHome <value> | --targetHomeName <value>}
{ [--executePrereqs | --postUpgrade | --rollback]}
{[--standBy | --allStandbyPrepared]}
{[--upgradeOptions <value>]  | [--standBy]}
[--removeGRP]
[--increaseCompatibleParameter]
[--resume [--sessionID <value>]]
[--waitForCompletion <value>]
```

Where:

- `--dbname` (mandatory) specifies the name of the database.
- `--targetHome` specifies the target Oracle home location
- `--targetHomeName` specifies the name of the target Oracle Database home
- `--standBy` use this option to upgrade standby databases in Data Guard configurations
- `--allStandbyPrepared` required for Data Guard configured primary databases. Flags to acknowledge that all the required operations are performed on the standby databases prior to upgrading primary database
- `--removeGRP` automatically removes the Guaranteed Restore Point (GRP) backup only if the database upgrade was successful
- `--increaseCompatibleParameter` automatically increases the compatible parameter as part of the database upgrade. The parameter will get increased only if the database upgrade was successful
- `--executePrereqs` runs only the preupgrade checks
- `--postUpgrade` use this option if postupgrade fails and needs to rerun the postupgrade steps
- `--rollback` reverts an Oracle Database to its original Oracle home
- `--upgradeOptions` use this option to pass DBUA-specific arguments to perform the Oracle Database upgrade. Refer to the corresponding Oracle documentation for the supported arguments and options.
  `--standby`
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

```
dbaascli dataguard prepareStandbyBlob --dbname <value> --blobLocation <value>
```

Where:

- `--dbname` specifies the Oracle Database name
- `--blobLocation` specifies the custom directory location where the standby blob file will be generated in a Data Guard environment

Frequently Asked Questions

Q: What is the purpose of the dbaascli dataguard prepareStandbyBlob command?

A: The `dbaascli 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 dbaascli 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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

```
--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 dbaascli 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 `oracle`user.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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>]
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
[--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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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 --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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

`--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 dbaascli 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 dbaascli dbHome getDatabases with the Oracle home location?

A: Here is an example of using the command with the Oracle home location:

```
dbaascli 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 dbaascli dbHome getDatabases with the Oracle home name?

A: Here is an example of using the command with the Oracle home name:

```
dbaascli 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 dbaascli 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 dbaascli 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli dbHome getDatabases --oracleHome

```
dbaascli dbHome getDatabases --oracleHome /u02/app/mar_home/
```

dbaascli dbHome getDetails

To view information about a specific Oracle home, use the `dbaascli dbHome getDetails` command. Specify either the Oracle home location or Oracle home name.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbHome getDetails
{ --oracleHomeName value | --oracleHome value }
```

Where:

- `--oracleHomeName` specifies user-defined Oracle home name
- `--oracleHome` specifies the location of Oracle home

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

`--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 dbaascli dbHome getDetails command with an Oracle home location?

A: Here is an example:

```
dbaascli 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 dbaascli dbHome getDetails command with an Oracle home name?

A: Here is an example:

```
dbaascli 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 dbaascli 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 dbaascli 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 `dbaascli 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 dbaascli dbHome getDetails - using Oracle home location

```
dbaascli dbHome getDetails --oracleHome /u02/app/home_db19c/
```

Example 6-26 dbaascli dbHome getDetails - using Oracle home name

```
dbaascli dbHome getDetails --oracleHomeName home_db19c
```

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Syntax

```
dbaascli diag collect [--components <value>] [--startTime <value>] [--endTime
<value>] [--nodes <value>] [--dbNames <value>]
        {
            [--objectStoreBucketUri <value>]
            | [--destLocation <value>]
        }
        [--waitForCompletion <value>]
```

Where:

- `--components` specifies a list of components for log collection.
  Valid values:
  - `db`
  - `gi`
  - `os`
  - `dbaastools`
  - `all`
- `--startTime` specifies the start time for log collection. Valid date and time format: `YYYY-MM-DDTHH24:MM:SS`
- `--endTime` specifies the end time for log collection. Valid date and time format: `YYYY-MM-DDTHH24:MM:SS`
- `--nodes` specifies a comma-delimited list of nodes to collect logs
- `--dbNames` specifies the database name for which to collect logs. You can specify only one database name.
- `--objectStoreBucketURI` specifies an Object Storage service pre-authenticated request (PAR) URL used to upload collected logs. Logs are collected from Guest VM. For more information, see *Using Pre-Authenticated Requests*.
- `--destLocation` specifies the location on Guest VM to collect logs. Default: `/var/opt/oracle/dbaas_acfs`
- `--waitForCompletion` Values: `true|false`. Default `true`. Specify `false` to run in the background.

Related Topics

- Using Pre-Authenticated Requests
- Collecting Tooling Log Data Examples
  The dbaascli `dbaascli diag collect` command uses the syntax shown below to collect tooling log data:

dbaascli diag healthCheck

To run diagnostic health checks, use the `dbaascli diag healthCheck` command.

Prerequisite

Run the command as the `root` user.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Syntax

```
dbaascli diag healthCheck
[--destLocation]
[--nodes]
[--objectStoreBucketURI]
```

Where:

- `--destLocation` specifies the location on Guest VM to collect logs. Default: `/var/opt/oracle/dbaas_acfs`
- `--nodes` specifies a comma-delimited list of nodes to collect logs
- `--objectStoreBucketURI` specifies an Object Storage service pre-authenticated request (PAR) URL used to upload collected logs. Logs are collected from Guest VM. For more information, see *Using Pre-Authenticated Requests*.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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 dbaascli 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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

`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:

```
dbaascli 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 dbaascli 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 dbaascli grid configureTCPS

To configure grid using self-signed certificate:

```
dbaascli grid configureTCPS
```

To configure grid using CA-signed certificate:

```
dbaascli grid configureTCPS --cert_type CA_SIGNED_CERT --server_cert /tmp/
certs/server_cert.pem --ca_cert_chain /tmp/certs/ca.pem --private_key /tmp/
certs/encrypted_private.key --private_key_password_protected false
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli 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:

```
dbaascli 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:

```
dbaascli 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`).

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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,
  "updatedTime" : 1628095633660,
  "description" : "Service job report for operation database create",
  "appMessages" : {
    "schema" : [ ],
    "errorAction" : "SUCCEED_AND_SHOW"
  },
  "resourceList" : [ ],
  "pct_complete" : "100"
}
```

dbaascli patch db prereq

> **Note**
>
> `dbaascli patch db prereq` and `dbaascli patch db apply` commands have been deprecated in `dbaascli` release 21.2.1.2.0, and replaced with `dbaascli grid patch`, `dbaascli dbhome patch`, and `dbaascli database move` commands.
> For more information, see:
>
> - `dbaascli grid patch`
> - `dbaascli dbhome patch`
> - `dbaascli database move`
> - *Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli*

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]
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
                    | [--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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
    {
            --pdbName <value>
            | --pdbUID <value>
        }
        --dbname <value> [--openMode <value>] [--startServices <value>] [--
waitForCompletion <value>]
```

Where:

- `--pdbName` specifies the name of the PDB
- `--pdbUID` specifies the identifier of the PDB
- `--dbname` specifies the name of the container database that hosts the PDB
- `--openMode` specifies the target `OPEN MODE` of PDB
- `--startServices` specifies to start all or a list of all services corresponding to a PDB. Accepted values are `all` or a comma-delimited list of PDB services.
- `--waitForCompletion` specifies to run the operation in the foreground or background. Valid values: `true|false`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
        {
            --pdbName <value>
            | --pdbUID <value>
        }
        --dbname <value> [--waitForCompletion <value>]
```

Where:

- `--pdbname` specifies the name of the PDB that you want to close.
- `--pdbUID` specifies the identifier of the PDB
- `--dbname` specifies the name of the container database that hosts the PDB.
- `--waitForCompletion` specifies to run the operation in the foreground or background. Valid values: `true|false`.

Upon successful completion of running this command, the PDB is closed on all of the container database instances.

Example 6-33 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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

dbaascli pdb getConnectString

To display Oracle Net connect string information for a pluggable database (PDB) run the `dbaascli pdb getConnectString` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb getConnectString --dbname <value>
        {
            --pdbName <value>
            | --pdbUID <value>
        }
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbname` specifies the name of the PDB for which you want to display connect string information
- `--pdbUID` specifies the identifier of the PDB

Example 6-34 dbaascli pdb getConnectString

```
dbaascli pdb getConnectString --dbname dbname --pdbName pdbName
```

dbaascli pdb create

To create a new pluggable database (PDB), use the `dbaascli pdb create` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb create --pdbName <value> --dbName <value> [--maxCPU <value>] [--
maxSize <value>] [--pdbAdminUserName <value>] [--lockPDBAdminAccount <value>]
[--resume [--sessionID <value>]] [--executePrereqs] [--waitForCompletion
<value>]
        {
            [--blobLocation <value>]
            | [--standbyBlobFromPrimary <value>]
        }
        [--pdbTdeKeyVersionOCID <value>]
```

Where:

- `--pdbName` specifies the name of the new PDB that you want to create
- `--dbName` specifies the name of the container database that hosts the new PDB

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--maxCPU` optionally specifies the maximum number of CPUs that are available to the PDB. Setting this option is effectively the same as setting the `CPU_COUNT` parameter in the PDB
- `--maxSize` optionally specifies the maximum total size of data files and temporary files for tablespaces belonging to the PDB. Setting this option is effectively the same as setting the `MAXSIZE PDB` storage clause in the `CREATE PLUGGABLE DATABASE` SQL command. You can impose a limit by specifying an integer followed by a size unit (`K`, `M`, `G`, or `T`), or you can specify `UNLIMITED` to explicitly enforce no limit
- `--pdbAdminUserName` specifies the new PDB admin user name
- `--lockPDBAdminAccount` specifies `true` or `false` to lock the PDB admin user account. Default value is `true`.
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specifies the location of the standby blob file, which is prepared from the primary database. This is required only for standby database PDB operations.

> **Note**
>
> the parameters`blobLocation` and `standbyBlobFromPrimary` are mutually exclusive.

- `--pdbTdeKeyVersionOCID`

During the PDB creation process, you are prompted to specify the administration password for the new PDB.

Example 6-35 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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

dbaascli pdb delete

To delete a pluggable database (PDB) run the `dbaascli pdb delete` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb delete --dbName <value>
        {
            --pdbName <value>
            | --pdbUID <value>
        }
        [--executePrereqs] [--waitForCompletion <value>] [--resume [--
sessionID <value>]] [--allStandbyPrepared] [--cleanupRelocatedPDB]
```

Where:

- `--dbName` specifies the name of the container database that hosts the PDB
- `--pdbName` specifies the name of the PDB that you want to delete
- `--pdbUID` specifies the UID of the PDB that you want to delete
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID
- `--allStandbyPrepared` specifies to confirm that the operation has been successfully run on all the standby databases
- `--cleanupRelocatedPDB` specify to cleanup source database after a PDB has been relocated

Example: dbaascli pdb delete
To delete a PDB from a standard database in a non-Data Guard environment or from Standby database in Data Guard environment.

```
dbaascli pdb delete --dbName db721 --pdbName pdb1
```

To create PDB from Primary database in Data Guard environment:

```
dbaascli pdb create --dbName db721 --pdbName pdb1 --allStandbyPrepared
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

dbaascli pdb getDetails

To view details of a pluggable database (PDB), use the `dbaascli pdb getDetails` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb getDetails --dbname <value>
        {
            --pdbName <value>
            | --pdbUID <value>
        }
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbname` specifies the name of the PDB that you want to delete
- `--pdbUID` specifies the identifier of the PDB

Example 6-36 dbaascli pdb getDetails

```
dbaascli pdb getDetails--dbname cdb name --pdbName pdb name associated with
the CDB


dbaascli pdb getDetails--dbname cdb name --pdbUID con_uid of that pdb
```

dbaascli pdb list

To view the list of pluggable databases (PDB) in a container database, use the `dbaascli pdb list` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb list --dbname
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 6-37 dbaascli pdb list

```
dbaascli pdb list --dbname cdb name
```

dbaascli pdb localClone

To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the `dbaascli pdb localClone` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb localClone --pdbName <value> --dbName <value> [--targetPDBName
<value>] [--powerLimit <value>] [--maxCPU <value>] [--maxSize <value>] [--
resume [--sessionID <value>]] [--executePrereqs] [--waitForCompletion
<value>] { [--blobLocation <value>] | [--standbyBlobFromPrimary <value>] } [--
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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 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`.

The newly cloned PDB inherits administration passwords from the source PDB.

Example 6-38 dbaascli pdb localClone

```
dbaascli pdb localClone --dbName db35 --pdbName PDB35 --targetPDBName
local_clone1 --maxCPU 2 --maxSize 15
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

dbaascli pdb open

To open a pluggable database (PDB), use the `dbaascli pdb open` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli pdb open
        {
            --pdbName <value>
            | --pdbUID <value>
        }
        --dbname <value> [--openMode <value>] [--startServices <value>] [--
waitForCompletion <value>] [--setPDBRefreshModeNone [--skipPDBRefresh] [--
pdbAdminUserName <value>]] [--executePrereqs] [--resume [--sessionID
<value>]] [--blobLocation <value>]
```

Where:

- `--pdbName` specifies the name of the PDB that you want to open
- `--pdbUID` specifies the identifier of the PDB
- `--dbname` specifies the name of the container database that hosts the PDB.
- `--openMode` specifies the target OPEN MODE of PDB
- `--startServices`: specifies to start all or list all services corresponding to a PDB. Accepted values are `all` or a comma-delimited list of PDB services.
- `--waitForCompletion`: specify `false` to run the operation in the background. Valid values: `true|false`
- `--setPDBRefreshModeNone`: specifies to convert a refreshable PDB to non-refreshable PDB
  - `--skipPDBRefresh`: specifies to skip refreshable PDB refresh
  - `--pdbAdminUserName`: specifies new PDB admin user name
- `--executePrereqs` specifies to run the prerequisite checks and report the results
- `--resume` resumes the previous operation
  - `--sessionID` specifies to resume a specific session ID
- `--blobLocation` specifies the custom directory location where the standby blob file will be generated in a Data Guard environment

Upon successful completion, the PDB is opened on all of the container database instances.

Example 6-39 dbaascli pdb open

```
dbaascli pdb open --dbname cdb name --pdbName pdb name associated with the CDB
```

```
dbaascli pdb open --dbname cdb name --pdbUID con_uid of that pdb
```

Optional: `--openMode READ_WRITE/READ_ONLY`

Copyright © 2023, 2026, Oracle and/or its affiliates.

dbaascli pdb recover

To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.

Prerequisite

- Run the command as the `root` user.
- Database must be configured with backup storage destination details where backups are stored.

Syntax

```
dbaascli pdb recover --dbname <value> --pdbName <value>
        {
            --start
                {
                    --untilTime <value> [--nonUTC]
                    | --untilSCN <value>
                    | --latest
                    | --tag <value>
                }
            | --status --uuid <value>
        }
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbName` specifies the name of the PDB that you want to recover
- `--start` begins the recovery of a PDB.
  - `--untilTime` recovers PDB until time. Input format: `DD-MON-YYYY HH24:MI:SS`
  - `--untilSCN` recovers PDB until SCN
  - `--latest` recovers PDB to the last known state
  - `--tag` recovers PDB to archival tag
- `--status` displays the details about a PDB recovery job process
  - `--uuid` unique identifier of the PDB recovery operation. Input format: `xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx`

Example 6-40 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
  ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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`

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 relocate

To relocate the specified PDB from the remote database into local database, use the `dbaascli pdb relocate` command.

Prerequisite

Run the command as the `oracle` user. When prompted, you must supply the SYS user password for the source database.

Syntax

```
dbaascli pdb relocate --pdbName <value> --dbName <value> --
sourceDBConnectionString <value> [--targetPDBName <value>] [--powerLimit
<value>] [--maxCpu <value>] [--maxSize <value>] [--resume [--sessionID
<value>]] [--executePrereqs] [--sourcePDBServices <value>] [--
sourcePDBReadOnlyServices <value>] [--relocateAvailabilityType <value>] [--
waitForCompletion <value>]
        {
            [--blobLocation <value>]
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
            | [--standbyBlobFromPrimary <value>]
        }
        [--upgradePDB] [--updateDBBlockCacheSize]
        {
            [--skipOpenPDB [--enableRefreshMode --refreshMode <value> [--
refreshIntervalInMinutes <value>]]]
            | [--completePDBRelocate]
        }
        {
            [--dblinkUsername <value> [--honorCaseSensitiveUserName]]
            | [--dbLinkName <value>]
        }
```

Where:

- `--pdbName` specifies the source PDB name to relocate
- `--dbName` specifies the target database name
- `--sourceDBConnectionString` specifies the source database connection string in the format *<scan_name>*:*<scan_port>*/*<database_service_name>*
- `--targetPDBName` specifies a name for the target PDB (new relocated PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the relocate operation
- `--maxCpu` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes|no`
- `--sourcePDBServices` specifies a list of comma-delimited source PDB services
- `--sourcePDBReadOnlyServices` specifies a comma-delimited list of source PDB read-only services
- `--relocateAvailabilityType` specifies `MAX|NORMAL` availability to enable the listener connection forwarding
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`
- `--blobLocation` specifies a custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specifies the location of the standby BLOB file, which is prepared from the primary database. This is required only for standby operations.

> **Note**
>
> The parameters `--blobLocation` and mutually exclusive.

- `--upgradePDB` specify `true` specifies to upgrade the PDB as part of this operation. Valid values : `true` | `false`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--updateDBBlockCachesize` specifies to enable application to set `db block cache size` initialization parameters to support data copy with different block size
- `--skipOpenPDB` specifies to indicate that the PDB should not be opened at the end of the current operation.
  - `--enableRefreshMode` specifies to enable PDB refresh support in PDB relocate step-one.
    * `--refreshMode` specifies to refresh mode for refreshable PDB. Valid values: `AUTO|MANUAL`
      * `--refreshIntervalInMinutes` specifies to refresh interval for `refreshablePDB` in minutes
- `--completePDBRelocate` specifies to complete the PDB relocation if done as a two-step operation.

Example 6-41 dbaascli pdb relocate

```
dbaascli pdb relocate --sourceDBConnectionString test-
scan.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521/
source_cdb_service_name --pdbName source_pdb --dbName target_db
```

dbaascli pdb remoteClone

To create a new pluggable database (PDB) as a clone of an existing PDB in another container database (CDB), use the `dbaascli pdb remoteClone` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli pdb remoteClone --pdbName <value> --dbName <value> --
sourceDBConnectionString <value> [--targetPDBName <value>] [--powerLimit
<value>] [--maxCPU <value>] [--maxSize <value>] [--resume [--sessionID
<value>]] [--executePrereqs] [--waitForCompletion <value>] [--
sourcePDBExportedTDEKeyFile <value>] { [--blobLocation <value>] | [--
standbyBlobFromPrimary <value>] } [--excludeUserTablespaces <value>] [--
excludePDBData <value>] [--pdbAdminUserName <value>] [--lockPDBAdminAccount
<value>] [--sourcePDBServiceConvertList <value>] [--upgradePDB] [--
refreshablePDB --refreshMode <value> [--refreshIntervalInMinutes <value>]] [--
updateDBBlockCacheSize] { [--dblinkUsername <value> [--
honorCaseSensitiveUserName]] | [--dbLinkName <value>] } [--
keepSourceDBServices <value>]
```

Where:

- `--pdbName` specifies the name of the source PDB that you want to clone
- `--dbname` specifies the name (`DB_NAME`) of the CDB that hosts the newly cloned PDB
- `--sourceDBConnectionString` specifies the source database connection string in the format *`scan_name`*`:`*`scan_port/database_service_name`*
- `--targetPDBName` specifies the name for the target PDB (new cloned PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the clone operation. Valid value is between 1 and 128

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--maxCPU` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--sourcePDBExportedTDEKeyFile` specifies the source PDB exported key file. This variable is applicable to only 12.1 database.
- `--blobLocation` specifies the custom path where the standby blob file will be generated in a Data Guard environment
- `--standbyBlobFromPrimary` specify the location of the standby blob file, which is prepared from the primary database. This is required only for standby database PDB operations

> **Note**
>
> The parameters `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive.

- `--excludeUserTablespaces` option to skip user table spaces, example *t1*,*t2*,*t3*.
- `--excludePDBData` specify `true/yes` to skip user data from source PDB.
- `--pdbAdminUserName` specifies new PDB admin user name
- `--lockPDBAdminAccount` specify `true` or `false` to lock the PDB admin user account. Default value is `true`.
- `--sourcePDBServiceConvertList` specify a comma-delimited list of source to target service names, which need to be converted. Syntax is `source_srv1:new_srv1`, `source_srv2:new_srv2`.
- `--upgradePDB` specify to upgrade the PDB as part of this operation
- `--refreshablePDB` specifies to create refreshable PDB
  - `--refreshMode` specifies refresh mode for refreshable PDB. Valid values: `AUTO|MANUAL`
    - `--refreshIntervalInMinutes` specifies refresh interval for `refreshablePDB` in minutes
  - `--dblinkUsername` specifies common user of a remote database used for database link to connect to the remote database
    - `--honorCaseSensitiveUserName` indicates specified username is case sensitive
- `--dbLinkName` specifies an existing database link name in the target database that points to the remote database

When promoted, you must supply the SYS user password for the source PDB. The newly cloned PDB inherits administration passwords from the source PDB. The cloned PDB is named using the following format: `dbname_sourcepdbname`. This command is supported only for databases that are not in a Data Guard configuration and use Oracle Database version 12.2.0.1, or later.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Example 6-42 dbaascli pdb remoteClone

```
dbaascli pdb remoteClone --sourceDBConnectionString test-
can.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521 --pdbName source_pdb1
--dbName db9944 --targetPDBName new_pdb1 --maxsize 5 --maxcpu 2


dbaascli pdb remoteClone --sourceDBConnectionString
orcla.dbaastoolslrgsu.dbaastoolslrgvc.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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 6-43 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 dbaascli tde changePassword command?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

```
dbaascli 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:

```
dbaascli tde changePassword --dbname mydatabase --pdbName mypdb
```

Q: What are the prerequisites for running the dbaascli 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 `dbaascli 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 6-44 dbaascli tde addSecondaryHsmKey

```
dbaascli tde addSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza


dbaascli tde addSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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 6-45 dbaascli tde enableWalletRoot

```
dbaascli tde enableWalletRoot --dbname db name --dbrestart rolling|full


dbaascli tde enableWalletRoot --dbname orcl


dbaascli tde enableWalletRoot --dbname orcl--dbrestart full
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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*:

```
dbaascli 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:

```
dbaascli 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:

```
dbaascli 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:

```
dbaascli tde encryptTablespacesInPDB --pdbName mypdb --useSysdbaCredential true
```

Q: What are the prerequisites for running the dbaascli 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 6-46 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `--skipPatchCheck` skips validation check for required patches if the value passed for this argument is `true`. Valid values: `true` or `false`
- `--executePrereqs` sexecute 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 dbaascli tde fileToHsm command?

A: The `dbaascli 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 dbaascli 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli 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 6-47 dbaascli tde fileToHsm --kmsKeyOCID

```
dbaascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5uxa
invvxza
```

```
dbaascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5uxa
invvxza --executePrereqs
```

```
dbaascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5uxa
invvxza --resume
```

dbaascli tde getHsmKeys

To get TDE active key details, use the `dbaascli tde getHsmKeys` command.

Prerequisite

Run the command as the `root` user.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 6-48 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 6-49 dbaascli tde getMkidForKeyVersionOCID

```
dbaascli tde getMkidForKeyVersionOCID --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha
```

dbaascli tde getPrimaryHsmKey

To get primary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde getPrimaryHsmKey` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde getPrimaryHsmKey
[--dbname]
```

Where:

- `--dbname` specifies the name of the database

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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Run the command as the `root` user.
- Be connected to an Exadata Cloud@Customer virtual machine via SSH.

Q: Who can execute the dbaascli 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 dbaascli 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 6-50 dbaascli tde getPrimaryHsmKey

```
dbaascli tde getPrimaryHsmKey --dbname dbname
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

A: The `dbaascli 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 dbaascli 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`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli 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 6-51 dbaascli tde hsmToFile

```
dbaascli tde hsmToFile --dbname dbname


dbaascli tde hsmToFile --dbname dbname --executePrereqs


dbaascli tde hsmToFile --dbname dbname --resume
```

dbaascli tde listKeys

To list TDE master keys, use the `dbaascli tde listKeys` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde listKeys
[--dbname <value>]
[--infoFilePath <value>]
```

Where:

- `--dbname` specifies the name of the database
- `--infoFilePath` specify the absolute path of the file where the results will be saved.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 6-52 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 Oracle Exadata Database Service on Exascale Infrastructure virtual machine using SSH.

Q: What does the --force parameter do in the dbaascli tde removeSecondaryHsmKey command?

A: The `--force` parameter allows the removal of the secondary HSM key without prompting the user for confirmation. If not specified, the command will prompt the user before deleting any keys.

Q: What does the --secondaryKmsKeyOCID parameter specify?

A: The `--secondaryKmsKeyOCID` parameter specifies the OCID (Oracle Cloud Identifier) of the secondary KMS key you want to remove from the existing HSM configuration.

Q: What does the --dbname parameter do?

A: The `--dbname` parameter specifies the name of the database for which the secondary HSM key is being removed.

Q: What is the purpose of the --precheckOnly parameter?

A: The `--precheckOnly` parameter, if set to `yes`, will only run the prechecks to validate the readiness for the removal operation without actually removing the secondary HSM key. If set to `no`, the full removal operation is performed.

Q: Is the --force parameter mandatory?

A: No, the `--force` parameter is optional. If it's not specified, the system will prompt the user for confirmation before proceeding with the key removal.

Q: Is the --secondaryKmsKeyOCID parameter mandatory?

A: Yes, you must provide the `--secondaryKmsKeyOCID` to identify the specific secondary HSM key that you want to remove from the configuration.

Q: Is the --dbname parameter mandatory?

A: No, the `--dbname` parameter is optional. If not specified, the command will attempt to remove the secondary HSM key from the default database on the system.

Q: What should I do if I want to remove the secondary HSM key without any user prompts?

A: You should use the `--force` parameter to bypass the confirmation prompt and remove the secondary HSM key directly:

```
dbaascli tde removeSecondaryHsmKey --force --secondaryKmsKeyOCID <value>
```

Q: How can I test whether the system is ready to remove the secondary HSM key without actually removing it?

A: You can use the `--precheckOnly` parameter set to yes to perform a precheck:

Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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 Oracle Exadata Database Service on Exascale Infrastructure virtual machine.
- The correct `--secondaryKmsKeyOCID` and `--dbname` values are provided. Check the error messages and logs for more details on the failure.

Q: What should I do if the removal operation fails partway through?

A: If the operation fails, review the error logs and try running the command with `--precheckOnly` to ensure the system is ready for the operation. If necessary, correct any issues before retrying.

Q: Can I run the dbaascli tde removeSecondaryHsmKey command while the database is running?

A: Yes, the command can be executed while the database is running, as it does not require the database to be stopped.

Q: What is the purpose of removing a secondary HSM key?

A: Removing a secondary HSM key is typically done when the key is no longer needed or when you want to manage the encryption keys used in your TDE (Transparent Data Encryption) configuration.

Q: How do I connect to the Oracle Exadata Database Service on Exascale Infrastructure virtual machine to run the command?

A: You can connect to the virtual machine using SSH. Refer to the Oracle Exadata Database Service on Exascale Infrastructure documentation for instructions on establishing a secure connection.

Example 6-53 dbaascli tde removeSecondaryHsmKey

```
dbaascli tde removeSecondaryHsmKey --dbname dbname
```

```
dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza
```

```
dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza --executePrereqs
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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` specifes 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:
  ```
  dbaascli tde setKeyVersion --kmsKeyVersionOCID <value> --masterKeyID <keyID>
  --dbname <DB_NAME>
  ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

- After successfully running the command on the standby database, run the command on the primary database using the `--standbySuc` parameter:
  `dbaascli 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`:

`dbaascli 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 dbaascli 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 6-54 dbaascli tde setKeyVersion

```
dbaascli 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-
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

```
dbaascli 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:

```
dbaascli 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:

```
dbaascli tde setPrimaryHsmKey --primaryKmsKeyOCID <key_OCID> --dbname <DB_NAME>
```

Then, run the command on the primary database with the `--standbySuc` parameter:

```
dbaascli 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:

```
dbaascli 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?

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 6-55 dbaascli tde setPrimaryHsmKey

```
dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza


dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocid1.key.oc1.eu-
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
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 dbaascli tde status command?

A: The status of the keystore, including whether it is open or closed, is part of the output returned by the `dbaascli 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:

```
dbaascli tde status --dbname <DB_NAME>
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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 dbaascli 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 dbaascli 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 dbaascli 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 `dbaascli 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Example 6-56 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 |
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
2210002022031021461 5 | COMPLETE        |
'---------------+--------------+--------+------+-----------
+---------------------+-----------'
```

- To start the Oracle Trace File Analyzer daemon on the local node, run the `tfactl start` command as `root`:

```
# tfactl start
Starting TFA..
Waiting up to 100 seconds for TFA to be started..
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
Successfully started TFA Process..
. . . . .
TFA Started and listening for commands
```

- To stop the Oracle Trace File Analyzer daemon on the local node, run the `tfactl stop` command as `root`:

```
# tfactl stop
Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..
```

Manage sysLens

- If sysLens is running, then once every 15 minutes data is collected in the local domU to discover the events to be reported. To check if sysLens is running, run the `systemctl status syslens` command as `root` in the domU:

```
# systemctl status syslens
? syslens.service
Loaded: loaded (/etc/systemd/system/syslens.service; enabled; vendor
preset: disabled)
Active: active (running) since Wed 2025-03-19 20:23:00 UTC; 44min ago
Process: 137603 ExecStopPost=/var/opt/oracle/syslens/bin/syslens --stop
(code=exited, status=0/SUCCESS)
Main PID: 137794 (python3)
Tasks: 7 (limit: 319999)
Memory: 194.7M
CGroup: /system.slice/syslens.service
??137794 /usr/bin/python3 /var/opt/oracle/syslens/bin/syslens_main.py --
level DRIFT=0 --daemon --service

Mar 19 20:53:12 scaqar07dv0201 su[336679]: pam_unix(su-l:session): session
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
closed for user oracle
Mar 19 20:53:13 scaqar07dv0201 su[336728]: (to oracle) root on none
Mar 19 20:53:13 scaqar07dv0201 su[336728]: pam_unix(su-l:session): session
opened for user oracle by (uid=0)
Mar 19 20:53:13 scaqar07dv0201 su[336728]: pam_unix(su-l:session): session
closed for user oracle
Mar 19 20:53:14 scaqar07dv0201 su[336808]: (to oracle) root on none
Mar 19 20:53:14 scaqar07dv0201 su[336808]: pam_unix(su-l:session): session
opened for user oracle by (uid=0)
Mar 19 20:53:14 scaqar07dv0201 su[336808]: pam_unix(su-l:session): session
closed for user oracle
Mar 19 20:53:15 scaqar07dv0201 su[336869]: (to oracle) root on none
Mar 19 20:53:15 scaqar07dv0201 su[336869]: pam_unix(su-l:session): session
opened for user oracle by (uid=0)
Mar 19 20:53:15 scaqar07dv0201 su[336869]: pam_unix(su-l:session): session
closed for user oracle
[root@scaqar07dv0201 opc]#
```

- If the sysLens is enabled, when there is a reboot of the domU, then sysLens starts automatically. To validate if sysLens is enabled to collect telemetry, run the `systemctl is-enabled syslens` command as `root` in the domU:

  ```
  # systemctl is-enabled syslens
  enabled
  ```

- To validate if sysLens is able to run as daemon:

  ```
  # /var/opt/oracle/syslens/bin/applin_ctl config_file get enable --
  file /etc/oracle/syslens/config/syslens.config
  true
  ```

- To validate if sysLens is configured to notify events, run the `tfactl get customerDiagnosticsNotifications` command as `root` in the domU:

  ```
  # tfactl get customerDiagnosticsNotifications
  .-------------------------------------------------------------------------.
  |                             Hostname                                    |
  +-----------------------------------------------------------+-------+
  | Configuration Parameter                                   | Value |
  +-----------------------------------------------------------+-------+
  | Send CEF notifications ( customerDiagnosticsNotifications ) | ON    |
  '-----------------------------------------------------------+-------'
  ```

Manage Database Service Agent

View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

- To check the status of the Database Service Agent, run the `systemctl status` command:

  ```
  # systemctl status dbcsagent.service
  dbcsagent.service
  Loaded: loaded (/usr/lib/systemd/system/dbcsagent.service; enabled; vendor
  preset: disabled)
  Active: active (running) since Fri 2022-04-01 13:40:19 UTC; 6min ago
  Process: 9603 ExecStopPost=/bin/bash -c kill `ps -fu opc |grep "java.*dbcs-
  ```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
agent.*jar" |awk '{print $2}' ` (code=exited, status=0/SUCCESS)
Main PID: 10055 (sudo)
CGroup: /system.slice/dbcsagent.service
 10055 sudo -u opc /bin/bash -c umask 077; /bin/java -
Doracle.security.jps.config=/opt/oracle/...
```

- To start the agent if it is not running, run the `systemctl start` command as the `root` user:

```
systemctl start dbcsagent.service
```

Related Topics

- To Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Overview of Events
- Notifications Overview

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.
>
> Potentially one hour of metrics will be lost between network restore and the first metric posted.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.
>
> Potentially one hour of metrics will be lost between network restore and the first metric posted.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**. A list of VM clusters is displayed.
3. In the list of VM clusters, click the VM cluster that contains the database for which you want to view the metrics. Details of the VM cluster you selected are displayed.
4. In the list of databases, click the database for which you want to view the metrics.
5. In the **Resources** section, click **Metrics**.
   A chart for each metrics is displayed. By default, the metrics for the last one hour are displayed.
6. Select a namespace from the **Metric namespace** from where you wish to view metrics.

   > **Note**
   >
   > - When Database Management is enabled, you will have an option to choose from `oci_database` or `oracle_oci_database` namespace.
   > - When Database Management is disabled, then you can view metrics only from the `oci_database` namespace.

7. If you want to change the interval, select the required start time and end time. Alternatively, you can select the interval from the Quick Selects drop down menu. The metrics are refreshed immediately for the selected interval.
8. For each metric, you can choose the interval and statistic independently.
   - Interval - The time period for which the metric is calculated.
   - Statistic - The mathematical method by which the metric is calculated.
9. For each metric, you can choose the following options from the 'Options' drop down menu.
   - View Query in Metrics Explorer
   - Copy Chart URL
   - Copy Query (MQL)
   - Create an Alarm on this Query
   - Table View

For Detailed information on various options for viewing the metrics chart, see Viewing Default Metric Charts.

Copyright © 2023, 2026, Oracle and/or its affiliates.

View Metrics for a PDB

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle 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 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 PBD 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).

Copyright © 2023, 2026, Oracle and/or its affiliates.

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**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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 2022-04-0113:40:19UTC; 6min ago
Process: 9603ExecStopPost=/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

> **Note**
>
> In addition to the metrics listed in the following section, Oracle analyzes additional metrics to provide the highest level of service operations and support for ensuring high availability of services.

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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`
					`resourceId_dbnode`
					`resourceName_dbnode`

`FilesystemUtilization`	**Filesystem Utilization**	percentage	Percent utilization of provisioned filesystem	1 minute	`hostName`
					`deploymentType`
					`filesystemName`
					`resourceId_dbnode`
					`resourceName_dbnode`

`LoadAverage`	**Load Average**	number	System load average over 5 minutes	1 minute	`hostName`
					`deploymentType`
					`resourceId_dbnode`
					`resourceName_dbnode`

`MemoryUtilization`	**Memory Utilization**	percentage	Percentage of memory available for starting new applications, without swapping. The available memory can be obtained via the following command: `cat /proc/meminfo`	1 minute	`hostName`
					`deploymentType`
					`resourceId_dbnode`
					`resourceName_dbnode`

`NodeStatus`	**Node Status**	integer	Indicates whether the host is reachable	1 minute	`hostName`
					`deploymentType`
					`resourceId_dbnode`
					`resourceName_dbnode`

`SwapUtilization`	**Swap Utilization**	percentage	Percent utilization of total swap space	1 minute	`hostName`
					`deploymentType`
					`resourceId_dbnode`
					`resourceName_dbnode`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
EcpusEnabled	**ECPUs Enabled**	integer	Number of ECPUs enabled.	1 minute	deploymentType
ExascaleVaultUtilization	**Exascale Vault Utilization**	percentage	Percent utilization of vault.	1 minute	deploymentType vaultName

The metrics listed in the following table are automatically available for the database.

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions
CpuUtilization	**CPU Utilization**	percentage	The CPU utilization expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
StorageUtilization	**Storage Utilization**	percentage	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.	1 hour	deploymentType resourceId_{database|pdb} resourceName_{database|pdb}
BlockChanges	**DB Block Changes**	Changes per second	The Average number of blocks changed per second.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database|pdb} resourceName_{database|pdb}

Copyright © 2023, 2026, Oracle and/or its affiliates.

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions

`ExecuteCount`	**Execute Count**	Count	The number of user and recursive calls that executed SQL statements during the selected interval.	5 minutes	`instanceNumber`
					`instanceName`
					`hostName`
					`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

`CurrentLogons`	**Current Logons**	Count	The number of successful logons during the selected interval.	5 minutes	`instanceNumber`
					`instanceName`
					`hostName`
					`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

`TransactionCount`	**Transaction Count**	Count	The combined number of user commits and user rollbacks during the selected interval.	5 minutes	`instanceNumber`
					`instanceName`
					`hostName`
					`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions

`UserCalls`	**User Calls**	Count	The combined number of logons, parses, and execute calls during the selected interval.	5 minutes	`instanceNumber`
					`instanceName`
					`hostName`
					`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

`ParseCount`	**Parse Count**	Count	The number of hard and soft parses during the selected interval.	5 minutes	`instanceNumber`
					`instanceName`
					`hostName`
					`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

`StorageUsed`	**Storage Space Used**	GB	Total amount of storage space used by the database at the collection time.	1 hour	`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

`StorageAllocated`	**Storage Space Allocated**	GB	Total amount of storage space allocated to the database at the collection time	1 hour	`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions

`StorageUsedByTablespace`	**Storage Space Used By Tablespace**	GB	Total amount of storage space used by tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	`tablespaceName`
					`tablespaceType`
					`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

`StorageAllocatedByTablespace`	**Allocated Storage Space By Tablespace**	GB	Total amount of storage space allocated to the tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	`tablespaceName`
					`tablespaceType`
					`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

`StorageUtilizationByTablespace`	**Storage Space Utilization By Tablespace**	percentage	This indicates the percentage of storage space utilized by the tablespace at the collection time. In case of container database, this metric provides root container tablespaces..	1 hour	`tablespaceName`
					`tablespaceType`
					`deploymentType`
					`resourceId_{database|pdb}`
					`resourceName_{database|pdb}`

Oracle Exadata Database Service on Exascale Infrastructure Events

Oracle Exadata Database Service on Exascale Infrastructure resources emit events, which are structured messages that indicate changes in resources.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> In addition to the events listed below, Oracle analyzes additional events to provide highest level of service operations and support for ensuring high availability of services.

- About Event Types on 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 Oracle 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 Event Types
  Review the list of event types that Data Guard group and Data Guard Associations emit.
- Oracle Database Home Event Types
  Review the list of events emitted by Oracle Database Homes.
- Database Event Types
  These are the event types that Oracle Databases in Exadata Cloud Service instances emit.
- Pluggable Database Event Types
  These are the event types that Oracle pluggable databases in Oracle Cloud Infrastructure emit.
- Database Service Events
  The Database Service emits events, which are structured messages that indicate changes in resources.
- Application VIP Event Types
  These are the event types that Application VIPs in Oracle Cloud Infrastructure emit.
- Interim Software Updates Event Types
  These are the event types that Interim Software Updates in Oracle Cloud Infrastructure emit.
- Serial Console Connection Event Types
  Review the list of event types that serial console connection emits.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Oracle Exadata Database Service on Exascale Infrastructure Event Types

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*.

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`
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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-1 (Cont.) Resource Events and Operations for ExaDB-XS

Friendly Name	Begin Event Sample	End event Sample
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": "ocid1.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": "ocid1.tenancy.oc1<UniqueID>",
        "userAgent": "Jersey/2.38 (HttpUrlConnection 17.0.6)"
      },
      "message": "GetExadbVmCluster succeeded",
      "request": {
        "action": "GET",
        "headers": {},
        "id": "/<uniqueID>",
        "parameters": {},
        "path": "/20160918/exadbVmClusters/ocid1.<uniqueID>"
      },
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "resourceId": "ocid1.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": "ocid1.compartment.oc1<UniqueID>",
      "ingestedtime": "2024-06-13T18:21:06.462Z",
      "loggroupid": "_Audit",
      "tenantid": "ocid1.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

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **Note**
>
> Exadata systems that use the old DB system resource model are deprecated and desupported.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance Reminder (ROLLING)	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancereminder`	This is an Oracle Cloud Operations reminder notice. Oracle has scheduled a quarterly maintenance update installation for Cloud Exadata Infrastructure instance <*infra_name*>, OCID *infra_ocid* in approximately *no_of_days_left* days schedule_time. The maintenance method for this maintenance is Rolling as selected per the maintenance preferences.
Cloud Exadata Infrastructure - Maintenance Reminder (NONROLLING)	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancereminder`	This is an Oracle Cloud Operations reminder notice. Oracle has scheduled a quarterly maintenance update installation for Cloud Exadata Infrastructure instance <*infra_name*>, OCID *infra_ocid* in approximately *no_of_days_left* days on *schedule_time*. The maintenance method for this maintenance is %s as selected per the maintenance preferences. Non-rolling maintenance minimizes maintenance time but will result in full system downtime.
Cloud Exadata Infrastructure - Maintenance Begin	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.begin`	This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> part of Maintenance Run <*mr-display-name*>, OCID <*mr-ocid*>. Your maintenance update started at <*start-time*>. You will get a notification on the completion of the quarterly maintenance update.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance End	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.end`	**Success:** This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> part of Maintenance Run <*mr-display-name*>, OCID <*mr-ocid*>. Your maintenance update started at <*start-time*> and was completed successfully at <*end-time*>. You have successfully completed maintenance updates for this window. **Failed:** This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> part of Maintenance Run <*mr-display-name*>, OCID <*mr-ocid*>. Your maintenance update started at <*start-time*> and was failed to complete successfully as planned. Our operations team is evaluating the failure and will notify you of the next steps to complete the maintenance update for this quarter. **Canceled:** This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> part of Maintenance Run <*mr-display-name*>, OCID <*mr-ocid*>. Your maintenance update started at <*start-time*>. Your maintenance has been canceled as requested. And a new window will be created according to the time given. **Duration Exceeded:** This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> part of Maintenance Run <*mr-display-name*>, OCID<*mr-ocid*>. Your maintenance update started at <*start-time*>. Your window was configured for a `WINDOW_DURATION` duration. Your maintenance is taking longer than the configured window duration. This window has

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type	Event Messages
		duration enforcement enabled. Oracle automation will reschedule all updates that have not started to a future maintenance window. Please acknowledge the updates rescheduled to run in a future unplanned maintenance window created by Oracle.
Cloud Exadata Infrastructure - Maintenance Custom action time Begin (ROLLING)	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.begin`	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> for Database Server <*db-server-name*>, OCID <*db-server-ocid*>. Your custom action time started at <*start-time*>. You will get a notification on the completion of the custom action time for the Database Server.
Cloud Exadata Infrastructure - Maintenance Custom action time End (ROLLING)	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.end`	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> for Database Server <*db-server-name*>, OCID <*db-server-ocid*>. Your custom action time started at <*start-time*> ended at <*end-time*>.
Cloud Exadata Infrastructure - Maintenance Custom action time Begin (NONROLLING)	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.begin`	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> for Database Servers <*db-server-name*>, OCID <*dbserver-ocid*> | <*db-server-name*>, OCID <*dbserver-ocid*>. Your custom action time started at <*start-time*>. You will get a notification on the completion of the custom action time for the Database Servers.
Cloud Exadata Infrastructure - Maintenance Custom action time End (NONROLLING)	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.end`	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> for Database Servers <*db-server-name*>, OCID <*dbserver-ocid*> | <*db-server-name*>, OCID <*dbserver-ocid*>. Your custom action time started at <*start-time*> ended at <*end-time*>.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure – Storage Server Maintenance Begin	`com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceStorageServers.begin`	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Storage Servers of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> for Storage Server(s) count <*cell-count*>. Your maintenance update started at <*start-time*>. You will get a notification on the completion of the quarterly maintenance update for the Storage Servers.
Cloud Exadata Infrastructure – Storage Server Maintenance End	`com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceStorageServers.end`	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Storage Servers of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> for Storage Server(s) count <*cell-count*>. Your maintenance update started at <*start-time*> and was completed successfully at <*end-time*>.
Cloud Exadata Infrastructure - DB Server Maintenance Begin (ROLLING)	`com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceVM.begin`	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Database Server <db-server-name>, OCID <db-server-ocid>. Your maintenance update started at <start-time>. You will get a notification on the completion of the quarterly maintenance update for the Database Server.
Cloud Exadata Infrastructure - DB Server Maintenance End (ROLLING)	`com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceVM.end`	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> for Database Server <*db-server-name*>, OCID <*db-server-ocid*>. Your maintenance update started at <*start-time*> and was completed successfully at <end-time>.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - DB Server Maintenance Begin (NONROLLING)	`com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceVM.begin`	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> for Database Servers <*db-server-name*>, OCID <*dbserver-ocid*> | <*db-server-name*>, OCID <*dbserver-ocid*>. Your maintenance update started at <*start-time*>. You will get a notification on the completion of the quarterly maintenance update for the Database Servers.
Cloud Exadata Infrastructure - DB Server Maintenance End (NONROLLING)	`com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceVM.end`	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-D Exadata Infrastructure <*infra-name*>, OCID <*infra-ocid*> for Database Servers <*db-server-name*>, OCID <*dbserver-ocid*> | <*db-server-name*>, OCID <*dbserver-ocid*>. Your maintenance update started at <*start-time*> and was completed successfully at <*end-time*>.
Cloud Exadata Infrastructure - VM Maintenance Begin	`com.oraclecloud.databaseservice.cloudexadatainfrastructurevmmaintenancebegin`	Your maintenance update started at <*start-time*>. You will get a notification on the completion of the quarterly maintenance update for the virtual machines.
Cloud Exadata Infrastructure - VM Maintenance End	`com.oraclecloud.databaseservice.cloudexadatainfrastructurevmmaintenanceend`	Your maintenance update started at <*start-time*> and was completed successfully at <*end-time*>.
Cloud Exadata Infrastructure – Maintenance Method Change	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancemethodchange`	Oracle Cloud Operations is announcing a change related to quarterly maintenance update for Cloud Exadata Infrastructure. There's a change in maintenance method on your service instance <*infra_name*>, OCID <*infra_ocid*> to *new_patching_mode*.
Cloud Exadata Infrastructure – Maintenance Rescheduled	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancerescheduled`	Oracle Cloud Operations is announcing reschedule of a quarterly maintenance update for Cloud Infrastructure. A maintenance run has been rescheduled on your service instance <*infra_name*>, OCID <*infra_ocid*> to *new_schedule_time*.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure – Maintenance Rescheduled With Reason	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancerescheduledwithreason`	The operator entered reason for rescheduling is sent as an email.
Cloud Exadata Infrastructure – Maintenance Scheduled (ROLLING)	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancescheduled`	Oracle Cloud Operations is announcing the availability of a new quarterly maintenance update for Cloud Exadata Infrastructure. Oracle has scheduled the installation of this new update on your service instance <*infra_name*>, OCID *infra_ocid* on *schedule_time*. The maintenance method for this maintenance is %s as selected per the maintenance preferences.
Cloud Exadata Infrastructure – Maintenance Scheduled (NONROLLING)	`com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancescheduled`	Oracle Cloud Operations is announcing the availability of a new quarterly maintenance update for Cloud Exadata Infrastructure. Oracle has scheduled the installation of this new update on your service instance <*infra_name*>, OCID *infra_ocid* on *schedule_time*. The maintenance method for this maintenance is %s as selected per the maintenance preferences. Non-rolling maintenance minimizes maintenance time but will result in full system downtime.
Cloud Exadata Infrastructure - IB Switch Maintenance Begin	`com.oraclecloud.databaseservice.cloudexadatainfrastructureibswitchmaintenance.begin`	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of the network fabric switches component of your Cloud Exadata Infrastructure instance <*infra_name*>, OCID *infra_ocid* has started. A follow-up notice will be sent when the network fabric switches maintenance operation has completed.
Cloud Exadata Infrastructure - IB Switch Maintenance End	`com.oraclecloud.databaseservice.cloudexadatainfrastructureibswitchmaintenance.end`	This is an Oracle Cloud Operations notice that quarterly maintenance update of the network fabric switches component of your Cloud Exadata Infrastructure instance <*infra_name*>, OCID *infra_ocid* has completed.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

Critical Event - EventName	Description
HEALTH.INFRASTRUCTURE.CELL.SQL_QUARANTINE	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\"

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.

Example critical event for Cloud Exadata infrastructure resource

In the following example of a "critical" event, you can see within the `additionalDetails` section of the event message that this particular message concerns an SQL statement that was terminated by Oracle Database Resource Manager because it was consuming excessive CPU or I/O resources. The `eventName` and `description` fields within the `additionalDetails` section provide information regarding the critical situation:

```
 {
  "eventType" :
"com.oraclecloud.databaseservice.exadatainfrastructure.critical",
  "cloudEventsVersion" : "0.1",
  "eventTypeVersion" : "2.0",
  "source" : "Exadata Storage",
  "eventTime" : "2021-07-30T04:53:18Z",
  "contentType" : "application/json",
  "data" : {
    "compartmentId" : "ocid1.tenancy.oc1.<unique_ID>",
    "compartmentName" : "example_name",
    "resourceName" : "my_exadata_resource",
    "resourceId" : "ocid1.dbsystem.oc1.phx.<unique_ID>",
    "availabilityDomain" : "phx-ad-2",
     "additionalDetails" : {
      "serviceType" : "exacs",
      "sqlID" : "gnwfm1jgqcfuu",
      "systemId" : "ocid1.dbsystem.oc1.eu-frankfurt-1.<unique_ID>",
      "creationTime" : "2021-05-14T13:29:28+00:00",
      "dbUniqueID" : "1558836122",
      "quarantineType" : "SQLID",
      "dbUniqueName" : "AB0503_FRA1S6",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "description" : "SQL statement terminated by Oracle Database Resource
Manager due to excessive consumption of CPU and/or I/O.
                       The execution plan associated with the terminated SQL
stmt is quarantined. Please find the sql identifier in sqlId
                       field of this JSON payload. This feature protects an
Oracle database from performance degradation.
                       Please review the SQL statement. You can see the
statement using the following commands: \"set serveroutput off\",
                       \"select sql_id, sql_text from v$sqltext where sql_id
=<sqlId>\", \"set serveroutput on\"",
      "quarantineReason" : "Manual",
      "asmClusterName" : "None",
      "component" : "storage",
      "infrastructureType" : "exadata",
      "name" : "143",
      "eventName" : "HEALTH.INFRASTRUCTURE.CELL.SQL_QUARANTINE",
      "comment" : "None",
      "quarantineMode" : "\"FULL Quarantine\"",
      "rpmVersion" : "OSS_20.1.8.0.0_LINUX.X64_210317",
      "cellsrvChecksum" : "14f73eb107dc1be0bde757267e931991",
      "quarantinePlan" : "SYSTEM"
    }
  },
  "eventID" : "<unique_ID>",
  "extensions" : {
    "compartmentId" : "ocid1.tenancy.oc1.<unique_ID>"
  }
}
```

Example information event for Cloud Exadata infrastructure resource

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..aaaaaaaao3lj36x6lwxyvc4wausjouca7pwyjfwb5ebsq5emrpqlql2gj5
iq",
    "compartmentName" : "intexadatateam",
    "resourceId" :
"ocid1.dbsystem.oc1.phx.abyhqljt5y3taezn7ug445fzwlngjfszbedxlcbctw45ykkaxyzc5i
sxoula",
    "availabilityDomain" : "phx-ad-2",
    "additionalDetails" : {
      "serviceType" : "exacs",
      "component" : "storage",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "systemId" :
"ocid1.dbsystem.oc1.phx.abyhqljt5y3taezn7ug445fzwlngjfszbedxlcbctw45ykkaxyzc5i
sxoula",
      "infrastructureType" : "exadata",
      "description" : "Flash Disk Failure has been detected. This is being
investigated by Oracle Exadata team and the disk will be
                        replaced if needed. No action needed from the
customer.",
      "eventName" : "HEALTH.INFRASTRUCTURE.CELL.FLASH_DISK_FAILURE",
      "FLASH_1_1" : "S2T7NA0HC01251  failed",
      "otto-ingestion-time" : "2021-12-17T19:14:43.205Z",
      "otto-send-EventService-time" : "2021-12-17T19:14:44.198Z"
    }
  },
  "eventID" : "30130ab4-42fa-4285-93a7-47e49522c698",
  "extensions" : {
    "compartmentId" :
"ocid1.tenancy.oc1..aaaaaaaao3lj36x6lwxyvc4wausjouca7pwyjfwb5ebsq5emrpqlql2gj5
iq"
  }
}
```

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.updatecloudvmclusteriormconfig.begin
Cloud VM Cluster - Update IORM Configuration End	com.oraclecloud.databaseservice.updatecloudvmclusteriormconfig.end
Cloud VM Cluster - Add Virtual Machine Begin	com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.begin
Cloud VM Cluster - Add Virtual Machine End	com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.end

Copyright © 2023, 2026, Oracle and/or its affiliates.

This is a reference event for a cloud VM cluster resource:

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "<unique_ID>",
    "eventType":
"com.oraclecloud.databaseservice.updatecloudvmclusteriormconfig.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2022-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "<unique_ID>",
      "eventName": "UpdateCloudVmClusterIormConfig",
      "compartmentName": "example_compartment",
      "resourceName": "my_container_database",
      "resourceId": "ocid1.cloudvmcluster.oc1.<unique_ID>",
      "resourceVersion": null,
      "additionalDetails": {
        "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.<unique_ID>",
        "freeFormTags": {},
        "definedTags": {},
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "lifecycleState": "AVAILABLE",
        "giVersion": "19.0.0.0.0",
        "cpuCoreCount": 16
      }
    }
  },
  "timeCreated": "2022-06-15T16:31:31.979Z"
}
```

This is a reference event for Add Virtual Machine Begin:

```
{
  "id":
"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"
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
},
{
  "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
        "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
        "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.v87pke1z9k9u6xaqo51taf6bunf0gc2wyhrbmjzbh3h1pjwakav
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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"
      ]
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    },
    {
      "name": "dataStorageSizeInGBs",
      "type": [
        "null",
        "Integer"
      ]
    },
    {
      "name": "licenseType",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "giVersion",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "dbNodeIds",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "timeZone",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "eventType":
"com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.end",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "eventID": "ced78bb7-3903-acd8-ff78-5567aa01a912",
    "eventTime": "2023-01-06T23:18:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "csid89a04ef74ccb8b48340f56e656cf/
729c99d3e5a34d548ddc31c054810454/634F086E8618E0A660946A6862C82A68",
      "eventName": "CloudVmClusterAddVirtualMachine",
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "compartmentName": null,
      "resourceName": "my_cloud_vm_cluster",
      "resourceId": "ocid1.cloudvmcluster.oc1.....unique_id",
      "resourceVersion": null,
      "availabilityDomain": "",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
        "tagSlug": "tag_slug",
        "identity": {
          "principalName": null,
          "principalId": null,
          "authType": null,
          "callerName": null,
          "callerId": null,
          "tenantId": null,
          "ipAddress": null,
          "credentials": null,
          "authZPolicies": null,
          "userGroups": null,
          "userAgent": null,
          "consoleSessionId": null
        },
        "request": {
          "id": "07197e12-b680-475e-851e-bb89fcd8376d",
          "path": null,
          "action": null,
          "parameters": null,
          "headers": null
        },
        "response": {
          "status": null,
          "responseTime": null,
          "headers": null,
          "payload": null,
          "message": null
        },
        "stateChange": {
          "previous": null,
          "current": {
            "licenseType": "BRING_YOUR_OWN_LICENSE",
            "dataStorageSizeGb": 60,
            "lifecycleState": "AVAILABLE",
            "sshPublicKeys": "...",
            "displayName": "my_cloud_vm_cluster",
            "cpuCoreCount": 16,
            "freeTags": {},
            "definedTags": {},
            "ocpuCountFractional": 16.0
          }
        },
        "additionalDetails": {
          "timeCreated": "2023-01-06T22:18:04.000Z",
          "timeUpdated": "2023-01-06T22:20:04.000Z",
          "lifecycleState": "AVAILABLE",
          "lifecycleDetails": null,
          "cloudExadataInfrastructureId":
      "ocid1.cloudexadatainfrastructure.oc1.....unique_id",
          "cpuCoreCount": 16,
          "ocpuCountFractional": 16.0,
          "dataStorageSizeInTBs": 4,
          "dataStorageSizeInGBs": 60,
          "licenseType": "BRING_YOUR_OWN_LICENSE",
          "giVersion": "19.0.0.0.0",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
        "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.ekmz1phzp4bl1k7m7tbygulbnakmjnrsi99eqjops3zvpt337pn
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"
    },
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    {
      "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    "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": {
      "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "UpdateCloudVmCluster",
      "compartmentName": "example_compartment",
      "resourceName": "my_container_database",
      "resourceId": "ocid1.cloudvmcluster.oc1.....unique_id",
      "resourceVersion": null,
      "additionalDetails": {
        "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.....unique_id",
        "freeFormTags": {},
        "definedTags": {},
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "lifecycleState": "AVAILABLE",
        "giVersion": "19.0.0.0.0",
        "cpuCoreCount": 16
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


This is a reference event for Cloud VM Cluster - Update End:


```
{
  "id":
"ocid1.eventschema.oc1.phx.svwkildsx63clp1q6phba7d6lns1rl92yc3uyc2ea5utjprqcwu
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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    "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"
    ]
  },
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    {
      "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": {
      "eventGroupingId": "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.vmclustereraddvirtualmachine.begin

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-2 (Cont.) VM Node Subsetting Events

Friendly Name	Event Type
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
DbNode - Action Begin	com.oraclecloud.DatabaseService.DbNodeAction.begin
DbNode - Action End	com.oraclecloud.DatabaseService.DbNodeAction.end

> **Note**
>
> All virtual machine actions (start, stop, reboot, force reboot, and force stop) use the same event type name. The event payload indicates the specific virtual machine action that was performed.

Example 6-57 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": "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": {
"id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-09-03T12:00:00.000Z",
      "timeUpdated": "2019-09-03T12:30:00.000Z",
      "displayName": "testDisplayName",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "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 - 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": "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": {
"id": "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
      "dbServerIds": "[ocid1.dbserver.1, ocid1.dbserver.2,...]",
      "timeZone": "US/Pacific"
    }
  }
}
```

This is a reference event for VM Cluster - Terminate Virtual Machine Begin:

```
"exampleEvent": {
"cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.vmclusterterminatevirtualmachine.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
"compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
"compartmentId": "ocid1.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Vmcluster-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
"id": "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"
    }
  }
}
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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": {
"id": "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"
    }
  }
}
```


This is a reference event for DbNode Action:


```
{
  "datetime": 1758918737152,
  "logContent": {
    "data": {
      "additionalDetails": {
        "dbHostId": "DbHost-f57d1844-36b4-4d39-801c-f233a7390469",
        "hostName": "myhost1-hbhyq",
        "lifecycleState": "AVAILABLE",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
        "powerAction": "Restart", // <====================
        "timeCreated": "2025-09-26T06:25:42Z",
        "timeUpdated": "2025-09-26T20:32:16Z",
        "vmClusterId":
"ocid1.exadbvmcluster.oc1.iad.anuwcljrk2wwj3yaih56fosgnx7hxbnjynlke5hipdyvvkil
v4lhxkjve36q"
      },
...
```

Data Guard Event Types

Review the list of event types that Data Guard group and Data Guard Associations emit.

> **Note**
>
> To receive events related to Data Guard actions on multiple standby databases, subscribe to the Data Guard group resource events. If you have not switched to the new model, you can continue to subscribe to the Data Guard Associations resource events. However, after switching to the new model, you will need to explicitly subscribe to the new Data Guard Group resource events.

Data Guard Event Types (Data Guard Group resource)

Review the list of event types that Data Guard Group emit.

Friendly Name	Event Type
Change Protection Mode Begin	`com.oraclecloud.databaseservice.changeprotectionmode.begin`
Change Protection Mode End	`com.oraclecloud.databaseservice.changeprotectionmode.end`
Data Guard Create Standby Database - Create Begin	`com.oraclecloud.databaseservice.createstandbydatabase.begin`
Data Guard Create Standby Database - Create End	`com.oraclecloud.databaseservice.createstandbydatabase.end`
Data Guard Switchover - Begin	`com.oraclecloud.databaseservice.dataguardswitchover.begin`
Data Guard Switchover - End	`com.oraclecloud.databaseservice.dataguardswitchover.end`
Data Guard Failover - Begin	`com.oraclecloud.databaseservice.dataguardfailover.begin`
Data Guard Failover - End	`com.oraclecloud.databaseservice.dataguardfailover.end`
Data Guard Reinstate - Begin	`com.oraclecloud.databaseservice.dataguardreinstate.begin`
Data Guard Reinstate - End	`com.oraclecloud.databaseservice.dataguardreinstate.end`
Data Guard Update Config - Begin	`com.oraclecloud.databaseservice.updatedataguardconfig.begin`
Data Guard Update Config - End	`com.oraclecloud.databaseservice.updatedataguardconfig.end`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type
Refresh Data Guard Health Status - Begin	com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.begin
Refresh Data Guard Health Status - End	com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.end

Example 6-58 Data Guard Health Status

This is a reference event for Refresh Data Guard Health Status - Begin


```
"exampleEvent": {
  "eventID": "022a63a4-ff77-11e9-a0af-f45c89b1cb17",
  "eventTime": "2024-12-04T10:06:21.000Z",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.begin",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "displayName": "Refresh Data Guard Health Status - Begin",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_id",
    "compartmentName": "example_name",
    "resourceName": "my_standby_database",
    "resourceId": "ocid1.database.oc1.phx.unique_id",
    "availabilityDomain": "AD1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.database.oc1.phx.unique_id",
      "timeCreated": "2025-10-04T10:06:21.000Z",
      "timeUpdated": "2025-10-04T10:06:21.000Z",
      "lifecycleState": "AVAILABLE",
      "lifecycleMessage": null,
      "dbSystemId": "ocid1.dbsystem.oc1.phx.unique_id",
      "databaseId": "ocid1.database.oc1.phx.unique_id",
      "dbHomeId": "ocid1.dbhome.oc1.phx.unique_id",
      "lastSyncedTime": "2025-10-04T10:06:21.000Z",
      "applyLag": "0 seconds",
      "syncState": "SYNCED",
      "switchoverReadiness": "HEALTHY",
      "switchoverReadinessMessage": null,
      "failoverReadiness": "HEALTHY",
      "failoverReadinessMessage": null,
      "dataLossExposure": "0 seconds",
      "memberTimeUpdated": "2025-10-04T10:06:21.000Z",
      "lastUpdatedIdentifier": "022a6912-ff77-11e9-9e77-f45c89b1cb17",
      "currentDataGuardRole" : "STANDBY",
      "noOfStandbyDatabases" : 3,
      "primaryDatabaseId" : "ocid1.database.oc1.iad.unique_id",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "protectionMode" : "MAXIMUM_AVAILABILITY",
      "transportType" : "SYNC"
    }
  }
}
```

This is a reference event for Refresh Data Guard Health Status - End

```
"exampleEvent": {
  "eventID": "022a63a4-ff77-11e9-a0af-f45c89b1cb17",
  "eventTime": "2024-12-04T10:06:21.000Z",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.end",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "displayName": "Refresh Data Guard Health Status - End",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_id",
    "compartmentName": "example_name",
    "resourceName": "my_standby_database",
    "resourceId": "ocid1.database.oc1.phx.unique_id",
    "availabilityDomain": "AD1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.database.oc1.phx.unique_id",
      "timeCreated": "2025-10-04T10:06:21.000Z",
      "timeUpdated": "2025-10-05T10:10:21.000Z",
      "lifecycleState": "AVAILABLE",
      "lifecycleMessage": null,
      "dbSystemId": "ocid1.dbsystem.oc1.phx.unique_id",
      "databaseId": "ocid1.database.oc1.phx.unique_id",
      "dbHomeId": "ocid1.dbhome.oc1.phx.unique_id",
      "lastSyncedTime": "2025-10-04T10:06:21.000Z",
      "applyLag": "0 seconds",
      "syncState": "SYNCED",
      "switchoverReadiness": "HEALTHY",
      "switchoverReadinessMessage": null,
      "failoverReadiness": "ALERT",
      "failoverReadinessMessage": "The Oracle Data Guard broker indicates
that the member 'DGHS1DB_s9b_sea' is not ready for failover. ",
      "dataLossExposure": "12 seconds",
      "memberTimeUpdated": "2025-10-05T10:10:21.000Z",
      "lastUpdatedIdentifier": "022a6912-ff77-11e9-9e77-f45c89b1cb17",
      "currentDataGuardRole" : "STANDBY",
      "noOfStandbyDatabases" : 3,
      "primaryDatabaseId" : "ocid1.database.oc1.iad.unique_id",
      "protectionMode" : "MAXIMUM_AVAILABILITY",
      "transportType" : "SYNC"
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    }
  }
}
```

Data Guard Event Types (Data Guard Associations resource)

Review the list of event types that Data Guard Associations emit.

Table 6-3 Data Guard Associations Events

Friendly Name	Event Type
Change Protection Mode Begin	`com.oraclecloud.databaseservice.changeprotectionmode.begin`
Change Protection Mode End	`com.oraclecloud.databaseservice.changeprotectionmode.end`
Data Guard Association - Create Begin	`com.oraclecloud.databaseservice.createdataguardassociation.begin`
Data Guard Association - Create End	`com.oraclecloud.databaseservice.createdataguardassociation.end`
Data Guard Association - Failover Begin	`com.oraclecloud.databaseservice.failoverdataguardassociation.begin`
Data Guard Association - Failover End	`com.oraclecloud.databaseservice.failoverdataguardassociation.end`
Data Guard Association - Reinstate Begin	`com.oraclecloud.databaseservice.reinstatedataguardassociation.begin`
Data Guard Association - Reinstate End	`com.oraclecloud.databaseservice.reinstatedataguardassociation.end`
Data Guard Association - Switchover Begin	`com.oraclecloud.databaseservice.switchoverdataguardassociation.begin`
Data Guard Association - Switchover End	`com.oraclecloud.databaseservice.switchoverdataguardassociation.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.createdbhome.begin`
DB Home - Create End	`com.oraclecloud.databaseservice.createdbhome.end`
DB Home - Patch Begin	`com.oraclecloud.databaseservice.patchdbhome.begin`
DB Home - Patch End	`com.oraclecloud.databaseservice.patchdbhome.end`
DB Home - Terminate Begin	`com.oraclecloud.databaseservice.deletedbhome.begin`
DB Home - Terminate End	`com.oraclecloud.databaseservice.deletedbhome.end`
DB Home - Update Begin	`com.oraclecloud.databaseservice.updatedbhome.begin`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type
DB Home - Update End	com.oraclecloud.databaseservice.updatedbhome.end
Oracle managed DB Home deletion started	com.oraclecloud.DatabaseService.OracleManagedDbHomeDeletionStarted
Oracle managed DB Home deletion completed	com.oraclecloud.DatabaseService.OracleManagedDbHomeDeletionCompleted
Oracle managed DB Home deletion failed	com.oraclecloud.DatabaseService.OracleManagedDbHomeDeletionFailed
Oracle managed software update staging started	com.oraclecloud.DatabaseService.OracleManagedDbHomeStageStarted
Oracle managed software update staging failed	com.oraclecloud.DatabaseService.OracleManagedDbHomeStageFailed
Oracle managed software update staging completed	com.oraclecloud.DatabaseService.OracleManagedDbHomeStageCompleted

Oracle Database Home example event

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": "ocid1.compartment.oc1.<unique_ID>"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1.<unique_ID>",
      "compartmentName": "example_compartment",
      "resourceName": "my_dbhome",
      "resourceId": "DbHome-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.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"
      }
    }
  }
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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.automaticbackupdatabase.begin`
Database - Automatic Backup End	`com.oraclecloud.databaseservice.automaticbackupdatabase.end`
Database - Create Backup Begin	`com.oraclecloud.databaseservice.backupdatabase.begin`
Database - Create Backup End	`com.oraclecloud.databaseservice.backupdatabase.end`
Database - Critical	`com.oraclecloud.databaseservice.database.critical`
Database - Information	`com.oraclecloud.databaseservice.database.information`
Database - Delete Backup Begin	`com.oraclecloud.databaseservice.deletebackup.begin`
Database - Delete Backup End	`com.oraclecloud.databaseservice.deletebackup.end`
Database - Migrate to KMS Key Begin	`com.oraclecloud.databaseservice.migratedatabasekmskey.begin`
Database - Migrate to KMS Key End	`com.oraclecloud.databaseservice.migratedatabasekmskey.end`
Database - Move Begin	`com.oraclecloud.databaseservice.movedatabase.begin`
Database - Move End	`com.oraclecloud.databaseservice.movedatabase.end`
Database - Restore Begin	`com.oraclecloud.databaseservice.restoredatabase.begin`
Database - Restore End	`com.oraclecloud.databaseservice.restoredatabase.end`
Database - Rotate KMS Key Begin	`com.oraclecloud.databaseservice.rotatedatabasekmskey.begin`
Database - Rotate KMS Key End	`com.oraclecloud.databaseservice.rotatedatabasekmskey.end`
Database - Terminate Begin	`com.oraclecloud.databaseservice.database.terminate.begin`
Database - Terminate End	`com.oraclecloud.databaseservice.database.terminate.end`
Database - Update Begin	`com.oraclecloud.databaseservice.updatedatabase.begin`
Database - Update End	`com.oraclecloud.databaseservice.updatedatabase.end`
Database - Upgrade Begin	`com.oraclecloud.databaseservice.upgradedatabase.begin`
Database - Upgrade End	`com.oraclecloud.databaseservice.upgradedatabase.end`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type
Oracle managed database software update enabled	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateEnabled
Oracle managed database software update disabled	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateDisabled
Oracle managed database software update scheduled	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateScheduled
Oracle managed database software update readiness check successful	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateReadinessCheckSuccessful
Oracle managed database software update readiness check failed	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateReadinessCheckFailed
Oracle managed database software update rescheduled	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateRescheduled
Oracle managed database software update started	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateStarted
Oracle managed database software update completed	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateCompleted
Oracle managed database software update failed	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateFailed
Oracle managed database software update awaiting user action	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateAwaitingUserAction
Oracle managed database software update reminder	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateReminder
Oracle managed database software update target version changed	com.oraclecloud.DatabaseService.ManagedSoftwareUpdateTargetVersionChanged

Database example event

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" : "ocid1.compartment.oc1.<unique_ID>",
"compartmentName": "example_compartment_name",
"resourceName": "my_backup",
"resourceId": "ocid1.dbbckup.oc1.<unique_ID>",
"availabilityDomain": "<availability_domain>",
"additionalDetails" : {
"timeCreated" : "2020-01-08T17:31:44Z",
"lifecycleState" : "CREATING",
"dbSystemId" : "ocid1.dbsystem.oc1.<unique_ID>",
"dbHomeId" : ocid1.dbhome.oc1.<unique_ID>",
"dbUniqueName" : DB1115_iad1dv",
"dbVersion" : "11.2.0.4.190716",
"databaseEdition" : "ENTERPRISE_EDITION_HIGH_PERFORMANCE",
"autoBackupsEnabled" : "false",
"backupType" : "FULL",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
"databaseId" : "ocid1.database.oc1.<unique_ID>",
},
"definedTags" : {
      "My_example_tag_name" :
        { "Example_key" : "Example_value" }
      },
    "eventID": "<unique_ID>",
    "extensions" : {
      "compartmentId": "ocid1.compartment.oc1.<unique_ID>"
    }
}
```

Pluggable Database Event Types

These are the event types that Oracle pluggable databases in Oracle Cloud Infrastructure emit.

Friendly Name	Event Type
Pluggable Database - Create Begin	`com.oraclecloud.databaseservice.createpluggabledatabase.begin`
Pluggable Database - Create End	`com.oraclecloud.databaseservice.createpluggabledatabase.end`
Pluggable Database - Delete Begin	`com.oraclecloud.databaseservice.deletepluggabledatabase.begin`
Pluggable Database - Delete End	`com.oraclecloud.databaseservice.deletepluggabledatabase.end`
Pluggable Database - Local Clone Begin	`com.oraclecloud.databaseservice.localclonepluggabledatabase.begin`
Pluggable Database - Local Clone End	`com.oraclecloud.databaseservice.localclonepluggabledatabase.end`
Pluggable Database - Remote Clone Begin	`com.oraclecloud.databaseservice.remoteclonepluggabledatabase.begin`
Pluggable Database - Remote Clone End	`com.oraclecloud.databaseservice.remoteclonepluggabledatabase.end`
Start Pluggable Database - Begin	`com.oraclecloud.databaseservice.startpluggabledatabase.begin`
Start Pluggable Database - End	`com.oraclecloud.databaseservice.startpluggabledatabase.end`
Stop Pluggable Database - Begin	`com.oraclecloud.databaseservice.stoppluggabledatabase.begin`
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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type
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"
    },
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
        "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
  "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    "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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

> **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..
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
. . . . .
TFA Started and listening for commands
```

- To stop the Oracle Trace File Analyzer daemon on the local node, run the `tfactl stop` command as `root`:

```
# tfactl stop
Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..
```

Manage sysLens

- If sysLens is running, then once every 15 minutes data is collected in the local domU to discover the events to be reported. To check if sysLens is running, run the `systemctl status syslens` command as `root` in the domU:

```
# systemctl status syslens
? syslens.service
Loaded: loaded (/etc/systemd/system/syslens.service; enabled; vendor
preset: disabled)
Active: active (running) since Wed 2025-03-19 20:23:00 UTC; 44min ago
Process: 137603 ExecStopPost=/var/opt/oracle/syslens/bin/syslens --stop
(code=exited, status=0/SUCCESS)
Main PID: 137794 (python3)
Tasks: 7 (limit: 319999)
Memory: 194.7M
CGroup: /system.slice/syslens.service
??137794 /usr/bin/python3 /var/opt/oracle/syslens/bin/syslens_main.py --
level DRIFT=0 --daemon --service

Mar 19 20:53:12 scaqar07dv0201 su[336679]: pam_unix(su-l:session): session
closed for user oracle
Mar 19 20:53:13 scaqar07dv0201 su[336728]: (to oracle) root on none
Mar 19 20:53:13 scaqar07dv0201 su[336728]: pam_unix(su-l:session): session
opened for user oracle by (uid=0)
Mar 19 20:53:13 scaqar07dv0201 su[336728]: pam_unix(su-l:session): session
closed for user oracle
Mar 19 20:53:14 scaqar07dv0201 su[336808]: (to oracle) root on none
Mar 19 20:53:14 scaqar07dv0201 su[336808]: pam_unix(su-l:session): session
opened for user oracle by (uid=0)
Mar 19 20:53:14 scaqar07dv0201 su[336808]: pam_unix(su-l:session): session
closed for user oracle
Mar 19 20:53:15 scaqar07dv0201 su[336869]: (to oracle) root on none
Mar 19 20:53:15 scaqar07dv0201 su[336869]: pam_unix(su-l:session): session
opened for user oracle by (uid=0)
Mar 19 20:53:15 scaqar07dv0201 su[336869]: pam_unix(su-l:session): session
closed for user oracle
[root@scaqar07dv0201 opc]#
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

- If the sysLens is enabled, when there is a reboot of the domU, then sysLens starts automatically. To validate if sysLens is enabled to collect telemetry, run the `systemctl is-enabled syslens` command as `root` in the domU:

```
# systemctl is-enabled syslens
enabled
```

- To validate if sysLens is able to run as daemon:

```
# /var/opt/oracle/syslens/bin/applin_ctl config_file get enable --
file /etc/oracle/syslens/config/syslens.config
true
```

- To validate if sysLens is configured to notify events, run the `tfactl get customerDiagnosticsNotifications` command as `root` in the domU:

```
# tfactl get customerDiagnosticsNotifications
.-------------------------------------------------------------------------.
|                             Hostname                                    |
+-----------------------------------------------------------+-------+
| Configuration Parameter                                   | Value |
+-----------------------------------------------------------+-------+
| Send CEF notifications ( customerDiagnosticsNotifications ) | ON    |
'-----------------------------------------------------------+-------'
```

Manage Database Service Agent

View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

- To check the status of the Database Service Agent, run the `systemctl status` command:

```
# systemctl status dbcsagent.service
dbcsagent.service
Loaded: loaded (/usr/lib/systemd/system/dbcsagent.service; enabled; vendor
preset: disabled)
Active: active (running) since Fri 2022-04-01 13:40:19 UTC; 6min ago
Process: 9603 ExecStopPost=/bin/bash -c kill `ps -fu opc |grep "java.*dbcs-
agent.*jar" |awk '{print $2}' ` (code=exited, status=0/SUCCESS)
Main PID: 10055 (sudo)
CGroup: /system.slice/dbcsagent.service
 10055 sudo -u opc /bin/bash -c umask 077; /bin/java -
Doracle.security.jps.config=/opt/oracle/...
```

- To start the agent if it is not running, run the `systemctl start` command as the `root` user:

```
systemctl start dbcsagent.service
```

Related Topics

- To Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Overview of Events

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-4 Database Service Events

Friendly Name	Event Name	Event Type	Threshold
Resource Utilization - Disk Usage	`HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE` This event is reported when VM guest file system free space falls below 10% free, as determined by the operating system `df(1)` command, for the following file systems: • `/` • `/u01` • `/u02` • `/var` (X8M and later only) • `/tmp` (X8M and later only)	`com.oraclecloud.databaseservice.dbnode.critical`	Critical Threshold: 90%
CRS status Up/Down	`AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN`. An event of type CRITICAL is created when the Cluster Ready Service (CRS) is detected to be down.	`com.oraclecloud.databaseservice.dbnode.critical` (if .DOWN and NOT "user_action")	N/A
	`AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN_CLEARED` An event of type INFORMATION is created once it is determined that the event for CRS down has cleared.	`com.oraclecloud.databaseservice.dbnode.information` (if .DOWN_CLEARED)	N/A

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-4 (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.databaseservice.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.databaseservice.dbnode.information` (if .DOWN_CLEARED)	N/A

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-4 (Cont.) Database Service Events

Friendly Name	Event Name	Event Type	Threshold
Net Listener Up/Down	`AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN` A DOWN event is created when a client listener goes down. The event is of type INFORMATION when a client listener is shutdown due to user action, such as with the Server Control Utility (`srvctl`) or Listener Control (`lsnrctl`) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a client listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a client listener is started. There is one client listener per node, each called LISTENER.	`com.oraclecloud.databaseservice.database.critical` (if .DOWN and NOT "user_action")	N/A
	`AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN_CLEARED` An event of type INFORMATION is created once it is determined that the event for Client Listener down has cleared.	`com.oraclecloud.databaseservice.database.information` (if .DOWN_CLEARED)	N/A

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-4 (Cont.) Database Service Events

Friendly Name	Event Name	Event Type	Threshold
CDB Up/Down	`AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN` 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.	`com.oraclecloud.databaseservice.database.critical` (if .DOWN and NOT "user_action")	N/A
	`AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN_CLEARED` An event of type INFORMATION is created once it is determined that the event for the CDB down has cleared.	`com.oraclecloud.databaseservice.database.information` (if .DOWN_CLEARED)	N/A
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.	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

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-4 (Cont.) Database Service Events

Friendly Name	Event Name	Event Type	Threshold
Critical DB Errors	`HEALTH.DB_CLUSTER.CDB.CORRUPTION` 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.	`com.oraclecloud.databaseservice.database.critical`	N/A
Other DB Errors	`HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG` An event of type CRITICAL is created if a CDB is either unable to archive the active online redo log or unable to archive the active online redo log fast enough to the log archive destinations.	`com.oraclecloud.databaseservice.database.critical`	N/A
	`HEALTH.DB_CLUSTER.CDB.DATABASE_HANG` An event of type CRITICAL is created when a process/session hang is detected in the CDB.	N/A	N/A
Backup Failures	`HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE` An event of type CRITICAL is created if there is a CDB backup with a FAILED status reported in the `v$rman_status` view.	`com.oraclecloud.databaseservice.database.critical`	N/A
Disk Group Usage	`HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE` 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.databaseservice.dbsystem.critical` `com.oraclecloud.databaseservice.dbsystem.information` (if < 90%)	Critical threshold: 90%

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-4 (Cont.) Database Service Events

Friendly Name	Event Name	Event Type	Threshold
Memory Usage	`CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE` 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.	`com.oraclecloud.databaseservice.dbnode.critical`	90%
sshd Configuration	`CONFIGURATION.DB_GUEST.SSHD.INVALID` An event of type CRITICAL is created if unexpected values are set in the `/etc/ssh/sshd_config` file.	`com.oraclecloud.databaseservice.dbnode.critical`	N/A
Disk Issues	`HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION` 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: • `/` • `/u01` • `/u02`	`com.oraclecloud.databaseservice.dbnode.critical`	N/A

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-4 (Cont.) Database Service Events

Friendly Name	Event Name	Event Type	Threshold
Oracle EXAchk Reported Issues	HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT Oracle EXAchk is Exadata database platform's holistic health check that includes software, infrastructure and database configuration checks. CRITICAL check alerts should be addressed in 24 hours to maintain the maximum stability and availability of your system. This database service event alerts every 24 hours whenever there are any CRITICAL checks that are flagged in the most recent Oracle EXAchk report. The event will point to the latest Oracle EXAchk zip report.	com.oraclecloud.databaseservice.dbnode.critical	N/A
DB I/O Latency	HEALTH.DB_GUEST.PDB.HIGH_IO_LATENCY	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >35ms
DB IO Latency Distribution	HEALTH.DB_GUEST.CDB.HIGH_IO_LATENCY_HISTOGRAM	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >32ms
Unusual CPU Waits	HEALTH.DB_GUEST.PDB.HIGH_CPU_WAITS	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >2 minutes
DB Waits	HEALTH.DB_GUEST.PDB.EXCESSIVE_WAITS	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >20%
CPU Throttling	HEALTH.DB_GUEST.PDB.EXCESSIVE_CPU_THROTTLING	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >>0.25 CPU_COUNT
DB Blocker Detector	HEALTH.DB_GUEST.PDB.DB_BLOCKER_STALL_CHECK	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >= 10 number of Session waiting Threshold: >= 600 sec wait time for the chain
Exadata Cache Efficiency	HEALTH.DB_GUEST.PDB.TOTAL_EXADATA_CACHE_HIT_RATIO	com.oraclecloud.databaseservice.dbnode.warning	Threshold: < 90%
Redo Log Wait	HEALTH.DB_GUEST.CDB.HIGH_REDO_LOG_WAITS	com.oraclecloud.databaseservice.dbnode.warning	Threshold: > 50% of Overall Background Processes Time

Copyright © 2023, 2026, Oracle and/or its affiliates.

Example 6-59 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]
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

Parameters

Table 6-5 tfactl blackout Command Parameters

Parameter	Description
`add|remove|print|`	Adds, removes, or prints blackout conditions.
`targettype` *`type`*	Limits blackout only to the specified target type.
Target type: `host|crs|asm|asmdg|database|dbbackup|db_dataguard|db_tablespace|pdb_tablespace|pdb|listener|service|os`	`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. `crs`: Blackout the availability of the Oracle Clusterware resource or events in the Oracle Clusterware logs. `asm`: Blackout the availability of Oracle Automatic Storage Management (Oracle ASM) on this machine or events in the Oracle ASM alert logs. `asmdg`: Blackout an Oracle ASM disk group. `database`: Blackout the availability of an Oracle Database, Oracle Database backup, tablespace, and so on, or events in the Oracle Database alert logs. `dbbackup`: Blackout Oracle Database backup events (such as CDB or archive backups). `db_dataguard`: Blackout Oracle Data Guard events. `db_tablespace`: Blackout Oracle Database tablespace events (container database). `pdb_tablespace`: Blackout Oracle Pluggable Database tablespace events (Pluggable database). `pdb`: Blackout Oracle Pluggable Database events. `listener`: Blackout the availability of a listener. `service`: Blackout the availability of a service. `os`: Blackout one or more operating system records.
`target all|`*`name`*	Specify the target for blackout. You can specify a comma-delimited list of targets. By default, the target is set to `all`.
`container` *`name`*	Specify the database container name (`db_unique_name`) where the blackout will take effect (for PDB, `DB_TABLESPACE`, and `PDB_TABLESPACE`).
`pdb` *`pdb_name`*	Specify the PDB where the blackout will take effect (for `PDB_TABLESPACE` only).
`events all|"`*`str1,str2`*`"`	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. `all`: Blackout everything for the target specified. *string*: Blackout for incidents where any part of the line contains the strings specified. Specify a comma-delimited list of strings.
`timeout` *`nh`*`|`*`nd`*`|none`	Specify the duration for blackout in number of hours or days before timing out. By default, the timeout is set to 24 hours (24h).
`c|local`	Specify if blackout should be set to cluster-wide or local. By default, blackout is set to `local`.
`reason` *`comment`*	Specify a descriptive reason for the blackout.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-5 (Cont.) tfactl blackout Command Parameters

Parameter	Description
`docollection`	Use this option to do an automatic diagnostic collection even if a blackout is set for this target.

Example 6-60 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
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type
Application Virtual IP (VIP) - Create End	com.oraclecloud.databaseservice.createapplicationvip.end
Application Virtual IP (VIP) - Delete Begin	com.oraclecloud.databaseservice.deleteapplicationvip.begin
Application Virtual IP (VIP) - Delete End	com.oraclecloud.databaseservice.deleteapplicationvip.end

Application VIP Event Types Examples:

This is a reference event for Application Virtual IP (VIP) - Create Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.5ur5er8bddumnu9r84rtt2c3282s5no31vsthibyqvvsisotnwp
csg9idv6q",
  "serviceName": "Database",
  "displayName": "Application Virtual IP (VIP) - Create Begin",
  "eventType": "com.oraclecloud.databaseservice.createapplicationvip.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-12-15T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "definedTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
{
  "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"
  ]
}
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
  ],
  "exampleEvent": {
    "eventType": "com.oraclecloud.databaseservice.createapplicationvip.begin",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "contentType": "application/json",
    "eventID": "ab2ac219-b435-1045-aaf3-13cd909ec106",
    "eventTime": "2022-12-16T21:16:04.000Z",
    "data": {
      "resourceId": "ocid1.applicationvip.oc1.....unique_id",
      "resourceName": "my_application_vip",
      "tagSlug": null,
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "request": {
        "id": "4260c9fd-d36b-4bc8-866e-c2dd53f34b2f",
        "path": null,
        "action": null,
        "parameters": null,
        "headers": null
      },
      "response": {
        "status": null,
        "responseTime": null,
        "headers": null,
        "payload": null,
        "message": ""
      },
      "stateChange": {
        "previous": null,
        "current": {
          "lifecycleState": "PROVISIONING",
          "hostnameLabel": "my_application_vip",
          "freeTags": {},
          "definedTags": {}
        }
      },
      "eventGroupingId": "csid74237ee84398b60cf1b834c81602/
f43a881dc99542318d46fa9285bdf2c5/6AC9F7641E1A5AD5C27D1650CB17E822",
      "eventName": "CreateApplicationVip",
      "availabilityDomain": "",
      "resourceVersion": null,
      "additionalDetails": {
        "id": "ocid1.applicationvip.oc1.....unique_id",
        "freeformTags": {},
        "definedTags": {},
        "timeCreated": "2022-12-15T21:17:59.000Z",
        "timeUpdated": "2022-12-15T21:18:04.389Z",
        "lifecycleState": "PROVISIONING",
        "lifecycleDetails": "",
        "hostnameLabel": "my_application_vip",
        "cloudVmClusterId": "ocid1.cloudvmcluster.oc1.....unique_id",
        "compartmentId": "ocid1.compartment.oc1.....unique_id",
        "vcnIpId": "ocid1.privateip.oc1.....unique_id",
        "ipAddress": "10.0.0.0",
        "subnetId": "ocid1.subnet.oc1.....unique_id",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
        "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.c1ok1948lwge4il6m85ta4jdlbnh1yjrjltrabujyv52calb0el
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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
            "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    "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"
      }
    }
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
  },
  "timeCreated": "2022-12-15T16:31:31.979Z"
}
```


This is a reference event for Application Virtual IP (VIP) - Delete Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.m2gheil6f1nfzb9ggpkkv17wdomdks8zin9nntqlghui6bckh17
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"
      ]
    },
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      {
        "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    "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"
}
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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"
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
        ]
      },
      {
        "name": "cloudVmClusterId",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "compartmentId",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "vcnIpId",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "ipAddress",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "subnetId",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "networkType",
        "type": [
          "null",
          "string"
        ]
      }
    ],
    "exampleEvent": {
      "eventType": "com.oraclecloud.databaseservice.deleteapplicationvip.end",
      "cloudEventsVersion": "0.1",
      "eventTypeVersion": "2.0",
      "source": "databaseservice",
      "contentType": "application/json",
      "eventID": "17619ca1-07ae-4e2d-a818-5b5f1fcd4f70",
      "eventTime": "2022-12-16T21:16:04.000Z",
      "data": {
        "resourceId": "ocid1.applicationvip.oc1.....unique_id",
        "resourceName": "my_application_vip",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Friendly Name	Event Type
Oneoff Patch - Create Begin	com.oraclecloud.databaseservice.createoneoffpatch.begin
Oneoff Patch - Create End	com.oraclecloud.databaseservice.createoneoffpatch.end
Oneoff Patch - Delete Begin	com.oraclecloud.databaseservice.deleteoneoffpatch.begin
Oneoff Patch - Delete End	com.oraclecloud.databaseservice.deleteoneoffpatch.end
Oneoff Patch - Download Begin	com.oraclecloud.databaseservice.downloadoneoffpatch.begin
Oneoff Patch - Download End	com.oraclecloud.databaseservice.downloadoneoffpatch.end

Interim Software Updates Event Types Examples:

This is a reference event for This is a reference event for Oneoff Patch - Create Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrsllp7rfneajgq2knxbqopwux24za7qzoe3mfj2bzfxt
nwqcxpbcq",
  "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",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
  "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.abyhqljrj4vvuph4qvj5eateeel6axblhkq3caqndgmjvwl3sld
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" },
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    { "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.abyhqljrdripga5rryplwmv4ws6hqzr3pjyl7wfvoaqutvg2ey2
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" },
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
    { "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.abyhqljrgwk2gvx5lmx6fiwotgdy32mdmrnkyzznz37dpb4mmeh
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"
}
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

This is a reference event for Oneoff Patch - Download Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljr3vkb7klt5hkbsnqzjaxmszsqomanlbqmr2tsrcq7xaf
cv2c74l2q",
  "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": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Download 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":
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
"ocid1.eventschema.oc1.phx.abyhqljrn2lruez55ah56kqksi5qfg6m7igvven7o2qkahlk5tk
wrj5ll3oa",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Download End",
  "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.end",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```


Serial Console Connection Event Types

Review the list of event types that serial console connection emits.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-6 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
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 6-61 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": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "CREATING",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    }
  }
}
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

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": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "ACTIVE",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.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": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "DELETING",
      "timeCreated": "2019-08-29T12:00:00.000Z",
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.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": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "DELETED",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.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": "ocid1.compartment.oc1..unique_ID"
  },
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "ACTIVE",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
      "tenantId": "ocid1.tenant.oc1..unique_ID",
      "compartmentId": "ocid1.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": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_dbnode",
    "resourceId": "DbNode-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-08-26T12:00:00.000Z",
      "timeUpdated": "2019-08-26T12:30:00.000Z",
      "dbSystemId": "ocid1.dbsystem.oc1.phx.unique_ID",
      "lifecycleDetails": "detail message",
      "vmClusterId": "VmCluster-unique_ID",
      "dbHostId": "dbHost-unique_ID",
      "nodeNumber": 2,
      "powerAction": "HardReset",
      "hostName": "testHostName"
    }
  }
}
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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 6-62 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": {
    "eventGroupingId": "2367d62d-cff8-11ed-bfd3-02001714f979",
    "eventName": "VmConsoleConnected",
    "compartmentId": "ocid1.compartment.oc1..<TRUNCATED>aaaaxxxxx",
    "compartmentName": "exacc-dev",
    "resourceName": "",
    "resourceId":
"ocid1.dbnodeconsoleconnection.oc1.iad.<TRUNCATED>aaaaaaxxxxx",
    "availabilityDomain": null,
    "freeformTags": null,
    "definedTags": null,

    "identity": {
      "principalName": "dsaes",
      "principalId": "ocid1.user.oc1..<TRUNCATED>aaaaaaaaaxxxxxxxxxx",
      "authType": "Native",
      "callerName": null,
      "callerId": null,
      "tenantId": "ocid1.tenancy.oc1..<TRUNCATED>aaaaaaxxxxx",
      "ipAddress": null,
```


Copyright © 2023, 2026, Oracle and/or its affiliates.

```
      "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": "ocid1.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.
- Enable Database Management
  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.
- Enable Database Management
  To enable Database Management for a pluggable database (PDB), use this procedure.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.
- Edit Database Management
  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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed.

Copyright © 2023, 2026, Oracle and/or its affiliates.

3. In the list of VM Clusters, click the VM Cluster that contains the database for which you want to enable Database Management.

   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

To enable Database Management for a container database (CDB), use this procedure.

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.
   - **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 a 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 a 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.
   - Private endpoint for Oracle RAC Databases.

   If a Database Management private endpoint is not available, then you must create one.

   For information on how to create a private endpoint, see *Create a Database Management Private Endpoint*.

5. In the **Management options** section, choose between the following options:
   - **Full management**: This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **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 AI Database**, then click **Oracle 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 that contains the pluggable database for which you want to enable Database Management.

   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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

To enable Database Management for a pluggable database (PDB), use this procedure.

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.
   - **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 a 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 a 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.
   - Private endpoint for Oracle RAC Databases.

   If a Database Management private endpoint is not available, then you must create one.

   For information on how to create a private endpoint, see *Create a Database Management Private Endpoint*.

5. In the **Management options** section, choose between the following options:
   - **Full management**: This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.
   - **Basic management**: This includes basic monitoring metrics and the ASH Analytics and SQL Monitoring features in Performance Hub for container databases. For more information on management options, see *About Management Options*.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle 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 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**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

To edit the Database Management configuration for a pluggable database (PDB), use this procedure.

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.
   - **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 a 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 a 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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.
   - Private endpoint for Oracle RAC Databases.

   If a Database Management private endpoint is not available, then you must create one.

   For information on how to create a private endpoint, see *Create a Database Management Private Endpoint*.

5. In the **Management options** section, choose between the following options:
   - **Full management**: This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.
   - **Basic management**: This includes basic monitoring metrics and the ASH Analytics and SQL Monitoring features in Performance Hub for container databases. For more information on management options, see *About Management Options*.

6. Click **Enable Database Management**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to disable Database Management.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database that contains the pluggable database for which you want to disable Database Management.

   Database Details page is displayed.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database for which you want to disable Database Management.

   Pluggable Database Details page is displayed.
7. In the **Database Information** section, under the **Associated Services**, check the status of Database Management.
8. If the Database Management is displayed as **Enabled**, perform the following steps to disable Database Management:
   a. Click **Disable**.
   b. A confirmation message with a link to the **Work requests** section on the **Database information** page is displayed. Click the link to monitor the progress of the work request.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle Exadata Database Service on Exascale Infrastructure**.
2. Choose your **Compartment**.

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the database for which you want to view Performance Hub.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database for which you want to view Performance Hub.

   Database Details page is displayed.
5. Click **Performance Hub**.

With Basic Management, Performance Hub provides **ASH Analytics** and **SQL Monitoring**. Advanced Management will additionally provide **ADDM**, **Workload**, and **Blocking Sessions**.

Performance Hub allows you to download reports for your managed databases. For more information about downloading reports, see *Automatic Workload Repository (AWR) Report*, *Active Sessions History (ASH) Report*, and *Performance Hub Report*.

Related Topics

- Automatic Workload Repository (AWR) Report
- Active Sessions History (ASH) Report
- Performance Hub Report

Using the Console to View Performance Hub for a Pluggable Database (PDB)

To view Performance Hub for a pluggable database (PDB), use this procedure. You must first enable Database Management to view the performance report.

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

   A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to view Performance Hub.

   Exadata VM Cluster Details page is displayed.

   Under **Resources**, **Databases** is selected by default.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 AI Database**, then click **Oracle 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 that contains the database for which you want to view the metrics.

   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 AI Database**, then click **Oracle 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 that contains the pluggable database for which you want to view the metrics.

   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**.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.

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`
- `db-multi-cloud-family`

Individual Resource-Types

- `database-family`
  - `db-nodes`
  - `db-node-snapshots`
  - `db-homes`
  - `databases`
  - `pluggable-databases`
  - `db-backups`
  - `dbnode-console-connection`
  - `db-connection-bundles`
- `db-multi-cloud-family`
  - `oracle-db-azure-connectors`
  - `oracle-db-azure-blob-mounts`
  - `oracle-db-azure-blob-containers`
  - `oracle-db-mci-work-requests`
  - `multi-cloud-resource-discoveries`
  - `oracle-db-azure-vaults`
  - `oracle-db-azure-keys`
  - `oracle-db-azure-vault-associations`
  - `oracle-db-gcp-connectors`
  - `oracle-db-gcp-keyrings`
  - `oracle-db-gcp-keys`
  - `oracle-db-aws-connectors`

Copyright © 2023, 2026, Oracle and/or its affiliates.

- `oracle-db-aws-keys`
- `oracle-db-aws-key-pools`

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
- Soft Deprecation of Certain Database Service APIs and Fields
- Permissions and API operation details for DB Backups
- Permissions and API operation details for Databases (CDBs)
- Permissions and API operation details for CDB Snapshot
- Permissions and API operation details for Data Guard Association
- Permissions and API operation details for DB Nodes
- Permissions and API operation details for DB Node Snapshots
- Permissions and API operation details for DB Node Console Connection
- Permissions and API operation details for DB Homes
- Permissions and API operation details for Database Software Image
- Permissions and API operation details for VM Clusters
  Review the list of permissions and API operations for the `exadb-vm-clusters` resource-type.
- Permissions and API operation details for 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 Data Guard Group
- Permissions and API operation details for Key Stores
- Permissions and API operation details for DB Connection Bundles
- Permissions Required for Each API Operation
  The following tables list the API operations for Oracle Exadata Database Service on Exascale Infrastructure instances in a logical order, grouped by resource type.
- Permissions and API operation details for Pluggable Databases (PDBs)
- Permissions and API operation details for Oracle DB Azure Connectors
- Permissions and API operation details for Oracle DB Azure Blob Mounts

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Permissions and API operation details for Oracle DB Azure Blob Containers
- Permissions and API operation details for Oracle DB MCI Work Requests
- Permissions and API operation details for Multicloud Resource Discoveries
- Permissions and API operation details for Oracle DB Azure Vaults
- Permissions and API operation details for Oracle DB Azure Keys
- Permissions and API operation details for Oracle DB Azure Vault Associations
- Permissions and API operation details for Oracle DB GCP Connectors
- Permissions and API operation details for Oracle DB GCP Keyrings
- Permissions and API operation details for Oracle DB GCP Keys
- Permissions and API operation details for Oracle DB AWS Identity Connectors
- Permissions and API operation details for Oracle DB AWS Key
- Permissions and API operation details for Oracle DB AWS Key Pool

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.

Soft Deprecation of Certain Database Service APIs and Fields

Some Database Service (DBaaS) APIs, response fields, and request parameters have been soft-deprecated. Soft-deprecated items continue to function but are no longer recommended for new integrations. Oracle recommends migrating to the replacement APIs and avoiding reliance on deprecated fields to ensure long-term compatibility.

Exadata Database Service on Cloud@Customer (ExaDB-C@C)

- **Soft-Deprecated APIs**

  The following patch-related APIs are soft-deprecated and replaced by update-based APIs:

  - ListVmClusterPatches
    Replacement: ListVmClusterUpdates
  - GetVmClusterPatch
    Replacement: GetVmClusterUpdate
  - ListVmClusterPatchHistoryEntries
    Replacement: ListVmClusterUpdateHistoryEntries
  - GetVmClusterPatchHistoryEntry

Copyright © 2023, 2026, Oracle and/or its affiliates.

Replacement: GetVmClusterUpdateHistoryEntry

- **Soft-Deprecated Fields**

  The following APIs contain a soft-deprecated response field:

  – GetDatabase
  Soft-deprecated field: `keyStoreWalletName` (response field)

  – ListDatabases
  Soft-deprecated field: `keyStoreWalletName` (response field)

  – GetBackup
  Soft-deprecated field: `keyStoreWalletName` (response field)

  – ListBackups
  Soft-deprecated field: `keyStoreWalletName` (response field)

  The following API has a soft-deprecated request parameter:

  – ListDbVersions
  Soft-deprecated parameter: `dbSystemShape` (query parameter)

All Database Services (ExaDB-C@C, ExaDB-D, ExaDB-XS, and Base Database Service)

- **Soft-Deprecated APIs**

  The following pluggable database clone APIs are soft-deprecated and replaced by a unified create operation:

  – LocalClonePluggableDatabase
  Replacement: CreatePluggableDatabase

  Use `pdbCreationTypeDetails.creationType = LOCAL_CLONE_PDB`.

  – RemoteClonePluggableDatabase
  Replacement: CreatePluggableDatabase

  Use `pdbCreationTypeDetails.creationType = REMOTE_CLONE_PDB`.

Additional Notes

- Soft-deprecated APIs, fields, and parameters remain supported but may be removed in a future release.
- New development should use the recommended replacement APIs and avoid deprecated fields.
- Existing integrations should plan migration to maintain forward compatibility.

Permissions and API operation details for DB Backups

The table below lists permissions and API operations for `db-backups`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	`DB_BACKUP_INSPECT`	`GetBackup` `ListBackups`	`ChangeCloudVmClusterCompartment` (also needs `use cloud-vmclusters`, `use db-homes`, and `use databases`)

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
read	*INSPECT* + `DB_BACKUP_CONTENT_READ`	*none*	`RestoreDatabase` (also needs `use databases`)
use	*no extra*	*no extra*	*none*
manage	*USE* + `DB_BACKUP_CREATE` `DB_BACKUP_DELETE`	`DeleteBackup`	`CreateBackup` (also needs `read databases`)

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	*INSPECT*+ `DATABASE_CONTENT_READ`	*no extra*	*no extra*
use	*READ* + `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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

manage	*USE* +
	`DATABASE_CREATE`
	`DATABASE_DELETE`
		no extra	`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 CDB Snapshot

The table below lists permissions and API operations for `container-database-snapshot`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	`DATABASE_SNAPSHOT_INSPECT`	`ListDatabaseSnapshot` `GetDatabaseSnapshot`	*none*
read	*INSPECT* + `DATABASE_SNAPSHOT_CONTENT_READ`	`GetDatabaseSnapshot` `ListDatabaseSnapshot`	*none*
use	*READ* + *no extra*	*none*	*none*

manage	*USE* +
	`DATABASE_SNAPSHOT_CREATE`
	`DATABASE_SNAPSHOT_DELETE`
	`DATABASE_SNAPSHOT_UPDATE`
		`CreateDatabaseSnapshot`
		`DeleteDatabaseSnapshot`
		`UpdateDatabaseSnapshot`
			none

Permissions and API operation details for Data Guard Association

The table below lists permissions and API operations for `data-guard-association`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	DATABASE_INSPECT	ListDataGuardAssociations, GetDataGuardAssociation	CreateDataGuardAssociation
READ	*no extra*	*no extra*	*none*
USE	*READ* + VM_CLUSTER_UPDATE + DB_HOME_UPDATE DATABASE_UPDATE	DeleteDatabase SwitchoverDataGuardAssociation, FailoverDataGuardAssociation, ReinstateDataGuardAssociation	CreateDataGuardAssociation
MANAGE	*USE* + DATABASE_DELETE	DeleteDatabase	*none*

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	*none*
read	*no extra*	*no extra*	*none*
use	DB_NODE_UPDATE DB_NODE_CONTENT_READ	UpdateDbNode	*none*
manage	*USE* + DB_NODE_POWER_ACTIONS DB_NODE_CONTENT_READ	DbNodeAction	*none*

Permissions and API operation details for DB Node Snapshots

The table below lists permissions and API operations for `db-node-snapshots`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_NODE_SNAPSHOT_INSPECT	listDbNodeSnapshot getDbNodeSnapshot	*none*
read	DB_NODE_SNAPSHOT_INSPECT	listDbNodeSnapshot getDbNodeSnapshot	*none*
use	*INSPECT* + DB_NODE_SNAPSHOT_CONTENT_READ	createDbNodeSnapshot mountDbNodeSnapshot, unmountDbNodeSnapshot	*none*
manage	*USE* + DB_NODE_SNAPSHOT_CREATE DB_NODE_SNAPSHOT_DELETE	createDbNodeSnapshot deleteDbNodeSnapshot	*none*

Permissions and API operation details for DB Node Console Connection

The table below lists permissions and API operations for `dbnode-console-connection`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DBNODE_CONSOLE_CONNECTION_INSPECT	GetDbNodeConsoleConnection ListDbNodeConsoleConnections	*none*
read	*no extra*	*no extra*	*none*
use	*READ* + DBNODE_CONSOLE_CONNECTION_UPDATE PLUGGABLE_DATABASE_UPDATE	UpdateDbNodeConsoleConnection	*none*
manage	*USE* + DBNODE_CONSOLE_CONNECTION_CREATE DBNODE_CONSOLE_CONNECTION_DELETE	CreateDbNodeConsoleConnection DeleteDbNodeConsoleConnection	*none*

Permissions and API operation details for DB Homes

The table below lists permissions and API operations for `db-homes`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_HOME_INSPECT	ListDBHome GetDBHome ListDbHomePatches ListDbHomePatchHistoryEntries GetDbHomePatch GetDbHomePatchHistoryEntry	*none*
read	*no extra*	*no extra*	*none*
use	DB_HOME_UPDATE	UpdateDBHome	ChangeCloudVmClusterCompartment (also needs use cloud-vmclusters, use databases, and inspect backups)
manage	USE + DB_HOME_CREATE DB_HOME_DELETE	no extra	CreateCloudVmCluster, DeleteCloudVmCluster (both also need manage cloud-vmclusters, manage databases, use vnics, and use subnets). If automatic backups are enabled on the default database, also needs manage backups CreateDbHome, (also needs use cloud-vmclusters and manage databases). If creating the Database Home by restoring from a backup, also needs read backups DeleteDbHome, (also needs use cloud-vmclusters and manage databases). If automatic backups are enabled on the default database, also needs manage backups. If the performFinalBackup option is selected, also needs manage backups and read databases.

Permissions and API operation details for Database Software Image

The table below lists permissions and API operations for database-software-image.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_SOFTWARE_IMG_INSPECT	ListDatabaseSoftwareImages GetDatabaseSoftwareImage	*none*
read	no extra	none	*none*
use	*READ* + DB_SOFTWARE_IMG_UPDATE	UpdateDatabaseSoftwareImage ChangeDatabaseSoftwareImageCompartment	*none*
manage	*USE* + DB_SOFTWARE_IMG_CREATE DB_SOFTWARE_IMG_DELETE	CreateDatabaseSoftwareImage DeleteDatabaseSoftwareImage	*none*

Permissions and API operation details for VM Clusters

Review the list of permissions and API operations for the `exadb-vm-clusters` resource-type.

Table 6-7 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
EXADB_VM_CLUSTER_INSPECT	ListExadbVmClusters GetExadbVmCluster ListExadbVmClusterUpdates GetExadbVmClusterUpdate ListExadbVmClusterUpdateHistoryEntries GetExadbVmClusterUpdateHistoryEntry	None

Table 6-8 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	*No extra*	None

Table 6-9 USE

Permissions	APIs Fully Covered	APIs Partially Covered
inspect + EXADB_VM_CLUSTER_UPDATE	RemoveVirtualMachineFromExadbVmClusterDetails	ChangeExadbVmClusterCompartment (also needs use db-homes, use databases, and inspect db-backups)

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-10 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered

`use + EXADB_VM_CLUSTER_CREATE, EXADB_VM_CLUSTER_DELETE`	*No extra*	`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`)

Permissions and API operation details for Exascale DB Storage Vaults

Review the list of permissions and API operations for the `exascale-db-storage-vaults` resource-type.

Table 6-11 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
`EXASCALE_DB_STORAGE_VAULT_INSPECT`	`ListExascaleDbStorageVaults` `GetExascaleDbStorageVault`	None

Table 6-12 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	*No extra*	None

Table 6-13 USE

Permissions	APIs Fully Covered	APIs Partially Covered
`inspect + EXASCALE_DB_STORAGE_VAULT_UPDATE`	`ChangeExascaleDbStorageVaultCompartment` `UpdateExascaleDbStorageVault`	None

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-14 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
use + EXASCALE_DB_STORAGE_VAULT_CREATE EXASCALE_DB_STORAGE_VAULT_DELETE	CreateExascaleDbStorageVault DeleteExascaleDbStorageVault	None

Permissions and API operation details for Data Guard Group

The table below lists permissions and API operations for Data Guard with multiple standby databases.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	DATABASE_INSPECT	ListDataGuardAssociations, GetDatabase	CreateDatabase
READ	*no extra*	*no extra*	*none*
USE	*READ* + **Standby:** CLOUD_VM_CLUSTER_INSPECT + DB_HOME_INSPECT + DATABASE_CREATE **Primary:** DATABASE_INSPECT + DATABASE_CONTENT_READ + DATABASE_UPDATE	CreateDatabase (Standby database) DeleteDatabase	DataguardAction (Switchover, Failover, Reinstate,, UpdateDatabase) Needs only DATABASE_INSPECT + DATABASE_UPDATE
MANAGE	*USE* + DATABASE_DELETE	DeleteDatabase	*none*

Permissions and API operation details for Key Stores

The table below lists permissions and API operations for keystores.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	KEY_STORE_INPSECT	GetKeyStore	ChangeKeyStoreCompartment
READ	*no extra*	*no extra*	*none*
USE	*READ* + KEY_STORE_UPDATE	UpdateKeyStore	ChangeKeyStoreCompartment
MANAGE	*USE* + KEY_STORE_CREATE KEY_STORE_DELETE	CreateKeyStore DeleteKeyStore	*none*

Copyright © 2023, 2026, Oracle and/or its affiliates.

Permissions and API operation details for DB Connection Bundles

This table lists permissions and API operations for `db-connection-bundle`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`DB_CONNECTION_BUNDLE_INSPECT`	`ListDbConnectionBundles`	`DownloadDbConnectionBundle`
READ	*INSPECT* + `DB_CONNECTION_BUNDLE_READ`	`GetDbConnectionBundle`	`ListDbConnectionBundle`, `DownloadDbConnectionBundle`
USE	*READ* + `DB_CONNECTION_BUNDLE_DOWNLOAD`	`DownloadDbConnectionBundle`	*none*
MANAGE	*USE* + *no extra*	*none*	*none*

Permissions Required for Each API Operation

The following tables list the API operations for Oracle Exadata Database Service on Exascale Infrastructure instances in a logical order, grouped by resource type.

Database API Operations

For information about permissions, see:

Permissions.

The following tables list of API operations and permissions by API operation.

Table 6-15 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`
`GetDbConnectionBundle`	`DB_CONNECTION_BUNDLE_READ` and `DB_CONNECTION_BUNDLE_INSPECT`
`ListDbConnectionBundle`	`DB_CONNECTION_BUNDLE_INSPECT`
`DownloadDbConnectionBundle`	`DB_CONNECTION_BUNDLE_READ`, `DB_CONNECTION_BUNDLE_INSPECT`, and `DB_CONNECTION_BUNDLE_DOWNLOAD`

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-16 Cloud VM Cluster

API Operation	Permissions Required to Use the Operation
ListCloudVmClusters	CLOUD_VM_CLUSTER_INSPECT
GetCloudVmCluster	CLOUD_VM_CLUSTER_INSPECT
CreateCloudVmCluster	CLOUD_VM_CLUSTER_CREATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and VNIC_CREATE and VNIC_ATTACH and SUBNET_ATTACH and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_CREATE DNS_VIEW_INSPECT)
ChangeCloudVmClusterCompartment	CLOUD_VM_CLUSTER_UPDATE
UpdateCloudVmCluster	CLOUD_VM_CLUSTER_UPDATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
GetCloudVmClusterIormConfig	CLOUD_VM_CLUSTER_INSPECT
UpdateCloudVmClusterIormConfig	CLOUD_VM_CLUSTER_UPDATE
DeleteCloudVmCluster	CLOUD_VM_CLUSTER_DELETE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and DB_HOME_DELETE and VNIC_DELETE and SUBNET_DETACH and VNIC_DETACH and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_DELETE)
AddVmToCloudVmCluster	CLOUD_VM_CLUSTER_UPDATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_CREATE, DNS_VIEW_INSPECT)
RemoveVmFromCloudVmCluster	CLOUD_VM_CLUSTER_UPDATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_DELETE)

Table 6-17 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 6-18 Virtual Machines / Nodes

API Operation	Permissions Required to Use the Operation
ListDbNodes	DB_NODE_INSPECT
GetDbNode	DB_NODE_INSPECT
DbNodeAction	DB_NODE_POWER_ACTIONS

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-19 Database Homes

API Operation	Permissions Required to Use the Operation
ListDbHomes	DB_HOME_INSPECT
GetDbHome	DB_HOME_INSPECT
ListDbHomePatches	DB_HOME_INSPECT
ListDbHomePatchHistoryEntries	DB_HOME_INSPECT
GetDbHomePatch	DB_HOME_INSPECT
GetDbHomePatchHistoryEntry	DB_HOME_INSPECT
CreateDbHome	DB_SYSTEM_INSPECT and DB_SYSTEM_UPDATE and DB_HOME_CREATE and DATABASE_CREATE To enable automatic backups for the database, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
UpdateDbHome	DB_HOME_UPDATE
DeleteDbHome	DB_SYSTEM_UPDATE and DB_HOME_DELETE and DATABASE_DELETE If automatic backups are enabled, also need DELETE_BACKUP If performing a final backup on termination, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ

Table 6-20 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 To enable oracle managed software updates, also need VM_CLUSTER_SUBSCRIBE_DB_SW_UPDATES
UpdateDatabase	DATABASE_UPDATE To enable automatic backups, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ To enable oracle managed software updates, also need VM_CLUSTER_SUBSCRIBE_DB_SW_UPDATES
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

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-21 Pluggable Databases (PDBs)

API Operation	Permissions Required to Use the Operation
ListPluggableDatabase	PLUGGABLE_DATABASE_INSPECT
GetPluggableDatabase	PLUGGABLE_DATABASE_INSPECT
CreatePluggableDatabase	PLUGGABLE_DATABASE_CREATE and DATABASE_INSPECT and DATABASE_UPDATE
UpdatePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
StartPluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
StopPluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
DeletePluggableDatabase	PLUGGABLE_DATABASE_DELETE and DATABASE_INSPECT and DATABASE_UPDATE
LocalClonePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE and PLUGGABLE_DATABASE_CONTENT_READ and PLUGGABLE_DATABASE_CONTENT_WRITE and PLUGGABLE_DATABASE_CREATE and DATABASE_INSPECT and DATABASE_UPDATE
RemoteClonePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE and PLUGGABLE_DATABASE_CONTENT_READ and PLUGGABLE_DATABASE_CONTENT_WRITE and PLUGGABLE_DATABASE_CREATE and DATABASE_INSPECT and DATABASE_UPDATE
enableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE

Table 6-22 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 6-23 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-24 Backups and Database Restore

API Operation	Permissions Required to Use the Operation
GetBackup	DB_BACKUP_INSPECT
ListBackups	DB_BACKUP_INSPECT
CreateBackup	DB_BACKUP_CREATE and DATABASE_CONTENT_READ
DeleteBackup	DB_BACKUP_DELETE and DB_BACKUP_INSPECT
RestoreDatabase	DB_BACKUP_INSPECT and DB_BACKUP_CONTENT_READ and DATABASE_CONTENT_WRITE

Table 6-25 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 6-26 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`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	PLUGGABLE_DATABASE_INSPECT	ListPluggableDatabases GetPluggableDatabase	UpdatePluggableDatabase StartPluggableDatabase StopPluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase RefreshPluggableDatabase ConvertRefreshablePluggableDatabase
	DATABASE_INSPECT	*no extra*	CreatePluggableDatabase DeletePluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase
read	*INSPECT* + PLUGGABLE_DATABASE_CONTENT_READ	*no extra*	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	*READ* + PLUGGABLE_DATABASE_CONTENT_WRITE	*no extra*	LocalClonePluggableDatabase RemoteClonePluggableDatabase

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
	PLUGGABLE_DATABASE_UPDATE	*no extra*	UpdatePluggableDatabase StartPluggableDatabase StopPluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase RefreshPluggableDatabase ConvertRefreshablePluggableDatabase
	DATABASE_UPDATE	*no extra*	CreatePluggableDatabase DeletePluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase
manage	*USE* + PLUGGABLE_DATABASE_CREATE	*no extra*	CreatePluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase
	PLUGGABLE_DATABASE_DELETE	*no extra*	DeletePluggableDatabase

Permissions and API operation details for Oracle DB Azure Connectors

The table below lists permissions and API operations for `oracle-db-azure-connectors`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_CONNECTOR_INSPECT	ListOracleDbAzureConnectors	*none*
READ	INSPECT+ ORACLE_DB_AZURE_CONNECTOR_READ	GetOracleDbAzureConnector	*none*
USE	READ+ ORACLE_DB_AZURE_CONNECTOR_UPDATE	UpdateOracleDbAzureConnector	*none*

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
MANAGE	USE + ORACLE_DB_AZURE_CONNECTOR_CREATE ORACLE_DB_AZURE_CONNECTOR_MOVE ORACLE_DB_AZURE_CONNECTOR_DELETE	*no extra*	CreateOracleDbAzureConnector ChangeOracleDbAzureConnectorCompartment DeleteOracleDbAzureConnector

Permissions and API operation details for Oracle DB Azure Blob Mounts

The table below lists permissions and API operations for `oracle-db-azure-blob-mounts`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_BLOB_MOUNT_INSPECT	ListOracleDbAzureBlobMounts	*none*
READ	INSPECT+ ORACLE_DB_AZURE_BLOB_MOUNT_READ	GetOracleDbAzureBlobMount	*none*
USE	READ+ ORACLE_DB_AZURE_BLOB_MOUNT_UPDATE	UpdateOracleDbAzureBlobMount	*none*
MANAGE	USE + ORACLE_DB_AZURE_BLOB_MOUNT_CREATE ORACLE_DB_AZURE_BLOB_MOUNT_MOVE ORACLE_DB_AZURE_BLOB_MOUNT_DELETE	*no extra*	CreateOracleDbAzureBlobMount ChangeOracleDbAzureBlobMountCompartment DeleteOracleDbAzureBlobMount

Permissions and API operation details for Oracle DB Azure Blob Containers

The table below lists permissions and API operations for `oracle-db-azure-blob-containers`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_BLOB_CONTAINER_INSPECT	ListOracleDbAzureBlobContainers	*none*
READ	INSPECT+ ORACLE_DB_AZURE_BLOB_CONTAINER_READ	GetOracleDbAzureBlobContainer	*none*
USE	READ+ ORACLE_DB_AZURE_BLOB_CONTAINER_UPDATE	UpdateOracleDbAzureBlobContainer	*none*

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
MANAGE	USE + ORACLE_DB_AZURE_BLOB_CONTAINER_CREATE ORACLE_DB_AZURE_BLOB_CONTAINER_MOVE ORACLE_DB_AZURE_BLOB_CONTAINER_DELETE	*no extra*	CreateOracleDbAzureBlobContainer ChangeOracleDbAzureBlobContainerCompartment DeleteOracleDbAzureBlobContainer

Permissions and API operation details for Oracle DB MCI Work Requests

The table below lists permissions and API operations for `oracle-db-mci-work-requests`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_INSPECT	*no extra*	ListWorkRequests ListWorkRequestErrors ListWorkRequestLogs
READ	INSPECT+ ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_READ	GetWorkRequest	*none*
USE	*no extra*	*no extra*	*none*
MANAGE	USE + ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_CANCEL	CancelWorkRequest	*none*

Permissions and API operation details for Multicloud Resource Discoveries

The table below lists permissions and API operations for `multi-cloud-resource-discoveries`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	MULTICLOUD_DISCOVERY_INSPECT	ListMultiCloudResourceDiscoveries	*none*
READ	INSPECT+ MULTICLOUD_DISCOVERY_READ	GetMultiCloudResourceDiscovery	*none*
USE	READ+ MULTICLOUD_DISCOVERY_UPDATE	UpdateMultiCloudResourceDiscovery	*none*

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

MANAGE	USE +
	`MULTICLOUD_DISCOVERY_CREATE`
	`MULTICLOUD_DISCOVERY_MOVE`
	`MULTICLOUD_DISCOVERY_DELETE`
		no extra	`CreateMultiCloudResourceDiscovery`
			`ChangeMultiCloudResourceDiscoveryCompartment`
			`DeleteMultiCloudResourceDiscovery`

Permissions and API operation details for Oracle DB Azure Vaults

The table below lists permissions and API operations for `oracle-db-azure-vaults`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`ORACLE_DB_AZURE_VAULT_INSPECT`	`ListOracleDbAzureVaults`	*none*
READ	INSPECT+ `ORACLE_DB_AZURE_VAULT_READ`	`GetOracleDbAzureVault`	*none*

USE	READ+
	`ORACLE_DB_AZURE_VAULT_UPDATE`
	`ORACLE_DB_AZURE_VAULT_REFRESH`
		no extra	`UpdateOracleDbAzureVault`
			`RefreshOracleDbAzureVault`

MANAGE	USE +
	`ORACLE_DB_AZURE_VAULT_CREATE`
	`ORACLE_DB_AZURE_VAULT_MOVE`
	`ORACLE_DB_AZURE_VAULT_DELETE`
		no extra	`CreateOracleDbAzureVault`
			`ChangeOracleDbAzureVaultCompartment`
			`DeleteOracleDbAzureVault`

Permissions and API operation details for Oracle DB Azure Keys

The table below lists permissions and API operations for `oracle-db-azure-keys`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`ORACLE_DB_AZURE_KEY_INSPECT`	`ListOracleDbAzureKeys`	*none*
READ	INSPECT+ `ORACLE_DB_AZURE_KEY_READ`	`GetOracleDbAzureKey`	*none*
USE	*no extra*	*no extra*	*none*
MANAGE	*no extra*	*no extra*	*none*

Copyright © 2023, 2026, Oracle and/or its affiliates.

Permissions and API operation details for Oracle DB Azure Vault Associations

The table below lists permissions and API operations for `oracle-db-azure-vault-associations`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_ASSOCIATION_INSPECT	ListOracleDbAzureVaultAssociations	*none*
READ	INSPECT+ ORACLE_DB_AZURE_ASSOCIATION_READ	GetOracleDbAzureVaultAssociation	*none*
USE	READ+ ORACLE_DB_AZURE_ASSOCIATION_UPDATE	UpdateOracleDbAzureVaultAssociation	*none*
MANAGE	ORACLE_DB_AZURE_ASSOCIATION_CREATE ORACLE_DB_AZURE_ASSOCIATION_MOVE ORACLE_DB_AZURE_ASSOCIATION_DELETE	*no extra*	CreateOracleDbAzureVaultAssociation ChangeOracleDbAzureVaultAssociationCompartment DeleteOracleDbAzureVaultAssociation CascadingDeleteOracleDbAzureVaultAssociation

Permissions and API operation details for Oracle DB GCP Connectors

The table below lists permissions and API operations for `oracle-db-gcp-connectors`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_GCP_IDENTITY_CONNECTOR_INSPECT	ListOracleDbGcpIdentityConnectors	*none*
READ	INSPECT+ ORACLE_DB_GCP_IDENTITY_CONNECTOR_READ	GetOracleDbGcpIdentityConnector	*none*
USE	READ+ ORACLE_DB_GCP_IDENTITY_CONNECTOR_UPDATE	UpdateOracleDbGcpIdentityConnector	*none*
MANAGE	USE+ ORACLE_DB_GCP_KEY_RING_CREATE ORACLE_DB_GCP_KEY_RING_MOVE ORACLE_DB_GCP_KEY_RING_DELETE	*no extra*	CreateOracleDbGcpIdentityConnector ChangeOracleDbGcpIdentityConnectorCompartment DeleteOracleDbGcpIdentityConnector

Copyright © 2023, 2026, Oracle and/or its affiliates.

Permissions and API operation details for Oracle DB GCP Keyrings

The table below lists permissions and API operations for `oracle-db-gcp-keyrings`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`ORACLE_DB_GCP_KEY_RING_INSPECT`	`ListOracleDbGcpKeyRings`	*none*
READ	INSPECT+ `ORACLE_DB_GCP_KEY_RING_READ`	`GetOracleDbGcpKeyRing`	*none*

USE	READ+
	`ORACLE_DB_GCP_KEY_RING_UPDATE`
	`ORACLE_DB_GCP_KEY_RING_REFRESH`
		`UpdateOracleDbGcpKeyRing`
		`RefreshOracleDbGcpKeyRing`
			none

MANAGE	USE+
	`ORACLE_DB_GCP_KEY_RING_CREATE`
	`ORACLE_DB_GCP_KEY_RING_MOVE`
	`ORACLE_DB_GCP_KEY_RING_DELETE`
		no extra	`CreateOracleDbGcpKeyRing`
			`ChangeOracleDbGcpKeyRingCompartment`
			`DeleteOracleDbGcpKeyRing`

Permissions and API operation details for Oracle DB GCP Keys

The table below lists permissions and API operations for `oracle-db-gcp-keys`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`ORACLE_DB_GCP_KEY_INSPECT`	`ListOracleDbGcpKeys`	*none*
READ	INSPECT+ `ORACLE_DB_GCP_KEY_READ`	`GetOracleDbGcpKey`	*none*
USE	READ+ *no extra*	*no extra*	*none*
MANAGE	USE+ *no extra*	*no extra*	*none*

Permissions and API operation details for Oracle DB AWS Identity Connectors

The table below lists permissions and API operations for `OracleDbAwsIdentityConnectors`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`ORACLE_DB_AWS_IDENTITY_CONNECTOR_INSPECT`	`ListOracleDbAwsIdentityConnectors`	*none*

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
READ	INSPECT+ ORACLE_DB_AWS_IDENTITY_CONNECTOR_READ	GetOracleDbAwsIdentityConnector	*none*
USE	READ+ ORACLE_DB_AWS_IDENTITY_CONNECTOR_UPDATE	*no extra*	UpdateOracleDbAwsIdentityConnector RefreshOracleDbAwsIdentityConnector
MANAGE	USE+ ORACLE_DB_AWS_IDENTITY_CONNECTOR_CREATE ORACLE_DB_AWS_IDENTITY_CONNECTOR_MOVE ORACLE_DB_AWS_IDENTITY_CONNECTOR_DELETE	*no extra*	CreateOracleDbAwsIdentityConnector ChangeOracleDbAwsIdentityConnectorCompartment DeleteOracleDbAwsIdentityConnector

Permissions and API operation details for Oracle DB AWS Key

The table below lists permissions and API operations for `OracleDbAwsKey`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AWS_KEY_INSPECT	ListOracleDbAwsKeys	*none*
READ	INSPECT+ ORACLE_DB_AWS_KEY_READ	GetOracleDbAwsKey	*none*
USE	READ+ ORACLE_DB_AWS_KEY_UPDATE	*no extra*	UpdateOracleDbAwsKey RefreshOracleDbAwsKey
MANAGE	USE+ ORACLE_DB_AWS_KEY_CREATE ORACLE_DB_AWS_KEY_MOVE ORACLE_DB_AWS_KEY_DELETE	*no extra*	CreateOracleDbAwsKey ChangeOracleDbAwsKeyCompartment DeleteOracleDbAwsKey

Permissions and API operation details for Oracle DB AWS Key Pool

This table lists permissions and API operations for `OracleDbAwsKeyPool`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AWS_KEY_POOL_INSPECT	ListOracleDbAwsKeyPools	*none*

Copyright © 2023, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
READ	*INSPECT* + ORACLE_DB_AWS_KEY_POOL_READ	GetOracleDbAwsKeyPool	*none*
USE	*READ* + ORACLE_DB_AWS_KEY_POOL_UPDATE	*no extra*	UpdateOracleDbAwsKeyPool
MANAGE	*USE* + ORACLE_DB_AWS_KEY_POOL_CREATE ORACLE_DB_AWS_KEY_POOL_MOVE ORACLE_DB_AWS_KEY_POOL_DELETE	*no extra*	CreateOracleDbAwsKeyPool ChangeOracleDbAwsKeyPoolCompartment DeleteOracleDbAwsKeyPool

Oracle Cloud Infrastructure Operations Insights

Oracle Cloud Infrastructure Operations Insights allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases deployed in Exadata Cloud Infrastructure.

Using Operations Insights on Oracle Cloud Databases allows you to:

- Analyze resource usage of databases across cloud databases
- Forecast future demand for database resources such as CPU, memory, and storage based on historical trends
- Improve resource utilization by identifying under and over utilized resources
- Identify Exadata systems projected to reach high utilization
- Identify total lead time to expand capacity using machine learning based forecast based on long term historic data to project future resource growth
- Compare SQL performance across databases and identify common patterns

Related Topics

- Enabling Database Cloud Service Databases

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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Features of Enterprise Manager Cloud Control
  Familiarize yourself with the features of Enterprise Manager Cloud Control to manage and monitor Oracle Exadata Database Service on Exascale Infrastructure 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.

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 Oracle Exadata Database Service on Exascale Infrastructure resources.

Enterprise Manager Target for Exadata Cloud

The target for Oracle Cloud Infrastructure Exadata resources, which supports the Oracle Exadata Database Service on Exascale Infrastructure services, 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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 and Oracle Exadata Database Service on Exascale Infrastructure 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 Oracle Exadata Database Service on Exascale Infrastructure 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.

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, Database Management must be enabled for the database. When enabling a database, the database administrator can choose from two database management options: Basic Management and Full Management.

> **Note**
>
> Using Identity and Access Management (IAM), you can create a policy that grants users access to Performance Hub while limiting actions they can take on Autonomous Databases, Oracle Database Cloud Service, and external databases. For information about IAM policies and Oracle Exadata Database Service on Exascale Infrastructure databases, see *Required IAM Policy*. For information about policies and how to use them, see How Policies Work.

Related Topics

- Required IAM Policy for 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

- Known Issues for Oracle Exadata Database Service on Exascale Infrastructure
  General known issues.
- Troubleshooting Oracle Managed Database Software Updates Issues
  Learn to identify and resolve issues in databases that are enrolled for Oracle Managed Database Software Updates.
- Troubleshooting Oracle Data Guard
  Learn to identify and resolve Oracle Data Guard issues.
- Obtaining Further Assistance

Known Issues for Oracle Exadata Database Service on Exascale Infrastructure

General known issues.

Troubleshooting Oracle Managed Database Software Updates Issues

Learn to identify and resolve issues in databases that are enrolled for Oracle Managed Database Software Updates.

Databases enrolled for Oracle Managed Database Software Updates have periodic readiness checks before they are updated.

This table lists all the issues and their resolutions:

Table 6-27 Readiness Checks

Name	Description	Affected Resource	Recommended Solution
Sudo Access	You have no sudo access	Virtual Machine	Log in with a user who has sudo privileges, and then use the `usermod` command to grant sudo access to the specified user.
Incorrect UID	User '*username*' has incorrect uid	Virtual Machine	Use `usermod` command to set uid for the specified user.
Incorrect umask	User '*username*' has incorrect umask	Virtual Machine	Add or update the umask value in the user's `~/.bash_profile`, save the file, then run `source ~/.bash_profile` to apply the changes.
Incorrect ownership	Required directory `/tmp` has incorrect ownership	Virtual Machine	Verify if directory `/tmp`with owner as `<dir_owner>`, group as `<group_owner>` and <minimum/exact> permission `<permission>` is present on the impacted `<resource_type>`.
Incorrect ownership	Required directory `/tmp` has incorrect group ownership	Virtual Machine	Verify if directory `/tmp`with owner as `<dir_owner>`, group as `<group_owner>` and <minimum/exact> permission `<permission>` is present on the impacted `<resource_type>`.
Incorrect ownership	Required file `<file_full_path>` has incorrect ownership	Virtual Machine	Verify if file `<file_full_path>` with owner as `<file_owner>`, group as `<group_owner>` and <minimum/exact> permission `<permission>` is present on the impacted `<resource_type>`.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-27 (Cont.) Readiness Checks

Name	Description	Affected Resource	Recommended Solution
Incorrect ownership	Required file <file_full_path> has incorrect group ownership	Virtual Machine	Verify if file <file_full_path> with owner as <file_owner>, group as <group_owner> and <minimum/exact> permission <permission> is present on the impacted <resource_type>.
Incorrect permission	Required directory <dir_full_path> has incorrect permission	Virtual Machine	Verify if directory <dir_full_path> with owner as <dir_owner>, group as <group_owner> and <minimum/exact> permission <permission> is present on the impacted <resource_type>.
Insufficient filesystem free space	Available space in '<dir_name>' directory is '1.95 GiB'. Minimum free space required is '2.00 GiB'	Virtual Machine	Verify if directory <dir_name> has minimum free space of <minimum_dir_space>. Free up more space to meet minimum free space requirement.
Missing file	Required file <file_full_path> is missing	Virtual Machine	Verify if file <file_full_path> with owner as <file_owner>, group as <group_owner> and <minimum/exact> permission <permission> is present on the impacted <resource_type>.
Passwordless SSH connectivity failed	User '*username*' has no passwordless SSH connectivity	Virtual Machine	Ensure that user <username> has passwordless SSH connectivity to all the nodes.
Version check	Executable 'java' version is below the expected version	Virtual Machine	Executable 'java' should have minimum '2.8.*' version.
Missing directory	Required directory /tmp is missing	Virtual Machine	Verify if directory /tmp with owner as <dir_owner>, group as <group_owner> and <minimum/exact> permission <permission> is present on the impacted <resource_type>.
User does not exist	User '*username*' does not exist	Virtual Machine	Verify if user <username> exists on all the nodes. If not, create user on that node.
Insufficient permission	Required directory <dir_full_path> has incorrect permission	Virtual Machine	Verify if directory <directory_full_path> with owner as <directory_owner>, group as <group_owner> and <minimum/exact> permission <permission> is present on the impacted <resource_type>.
Internal Issue	An internal issue that will be investigated by Oracle operations	Database	No customer action required.
Online Database instance count check	Verify whether the database is running on multiple nodes to ensure compatibility with rolling updates	Database	Make sure that database is running on more than one node to support rolling patching.

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-27 (Cont.) Readiness Checks

Name	Description	Affected Resource	Recommended Solution
Oracle Network Services configuration check	Check if the files in the Oracle Database home directory have the expected permissions, owner, and group	Database	Review and correct the Oracle Network Services configuration files (listener.ora, tnsnames.ora, and sqlnet.ora). Verify that each file has the proper syntax, correct host, port, and service names, and adheres to Oracle's configuration standards. After making necessary corrections, restart the listener and test connectivity to ensure the network services function properly. For more details, see https://docs.oracle.com/en/error-help/db/index.html.
Software home check	Check if the files in the Oracle Database Home directory have the expected permissions, owner, and group	Database	Make sure all files in the Oracle Database Home directory are assigned the correct permissions, owner, and group. For more details, see https://docs.oracle.com/en/error-help/db/index.html.
Database backup job check	Verify that no RMAN jobs are running in the database	Database	Consider patching the database when RMAN jobs are not running. For more details, refer to MOS Doc ID 2975965.1
Central Inventory check	Verify that the Central Inventory is defined correctly and the Oracle home is registered in it	Database	Make sure that the Central Inventory is defined correctly and that the Oracle home is registered in it. For more details, refer to MOS Doc ID 2975965.1
Check sys public grants	Verify that database PUBLIC user group is granted the required privilege on dictionary objects	Database	Grant EXECUTE privilege to PUBLIC user group for listed dictionary objects. For more details, refer to MOS Note 247093.1
Oracle Database Keystore check	Verify that Oracle Database Keystore is open	Database	Open the Oracle Database Keystore after the database has been opened in the Oracle home. For more details, refer to MOS Doc ID 2975965.1
Oracle Database Vault check	Verify that Oracle Database Vault is enabled and that the DV_PATCH_ADMIN role has been granted to the SYS user	Database	Log in as the Database Vault (DV) administrator and grant the DV_PATCH_ADMIN role to SYS with 'container=all' clause in CDB$ROOT. For non-CDB or a single PDB, log in as the Database Vault (DV) administrator and grant the DV_PATCH_ADMIN role to SYS.
Queryable Inventory DBA directories check	Verify that Oracle directory objects are accessible and can be queried	Database	Verify that the directories used by Queryable Patch Inventory are defined correctly and are relative to the ORACLE HOME. Check both physical paths and directory objects; and create or correct as needed. For more details, refer to MOS Note 1602089.1
Queryable Inventory external table check	Verify that opatch_xml_inv table can be queried	Database	Verify external tables work correctly in your environment. For more details, refer to MOS Note 1602089.1

Copyright © 2023, 2026, Oracle and/or its affiliates.

Table 6-27 (Cont.) Readiness Checks

Name	Description	Affected Resource	Recommended Solution
Queryable Inventory locks check	Check whether the inventory is queryable by ensuring the ORA$QP_CONTROL_LOCK lock is not allocated	Database	Release lock ORA$QP_CONTROL_LOCK before running datapatch. For more details, refer to MOS Doc ID 2975965.1
Queryable Inventory package check	Verify that queryable Inventory package is able to retrieve OPatch Inventory information	Database	Refer to MOS Note 1602089.1 for details on verifying the Queryable Inventory package before patching.
Symlinks on oracle home path check	Check if symbolic links are present in the directory object paths used with BFILE data types, the UTL_FILE package, or external tables	Database	Recreate directory objects to remove any symbolic links from directory paths. To identify paths that contain symbolic links before patching use OS commands. For more details, refer to MOS Doc ID 2975965.1
Temporary tablespace status check	Verify that temporary tablespaces have sufficient space for patching	Database	Ensure there is adequate tablespace for patching. Minimum recommended space is 2GB. For more details, refer to MOS Doc ID 2975965.1
Temporary file exists check	Verify that the default temporary tablespace has at least one temporary file associated with it	Database	Add at least one temporary file to the indicated temporary tablespace.
Temporary file online check	Verify that the default temporary tablespace or tablespace group includes at least one online temporary file	Database	Bring at least one temporary file online in the default temporary tablespace or tablespace group.

Troubleshooting Oracle Data Guard

Learn to identify and resolve Oracle Data Guard issues.

When troubleshooting Oracle Data Guard, you must first determine whether the problem occurs during the Data Guard setup and initialization or during Data Guard operation, when lifecycle commands are entered. The steps to identify and resolve the issues are different, depending on the scenario in which they are used.

There are three lifecycle operations: switchover, failover, and 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

Copyright © 2023, 2026, Oracle and/or its affiliates.

- 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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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`

Copyright © 2023, 2026, Oracle and/or its affiliates.

- **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.

Copyright © 2023, 2026, Oracle and/or its affiliates.

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@exa1-****** ~]$ srvctl status
                scan
    ```
  - ```
    [grid@exa1-****** ~]$ 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 = exa1-*****-
   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:

Copyright © 2023, 2026, Oracle and/or its affiliates.

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>
   DGMGRL> VALIDATE CONFIGURATION VERBOSE
   DGMGRL> VALIDATE DATABASE VERBOSE <PRIMARY>
   DGMGRL> VALIDATE DATABASE VERBOSE <STANDBY>
   ```

2. Consult the Oracle Database documentation to check for the respective error message. For example:
   - **ORA-16766:** Redo apply is stopped.
   - **ORA-16853**: Apply lag has exceeded specified threshold.
   - **ORA-16664**: Unable to receive the result from a member (under the standby database).
   - **ORA-12541**: TNS: no listener (under the primary database)

Obtaining Further Assistance

Copyright © 2023, 2026, Oracle and/or its affiliates.

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

Copyright © 2023, 2026, Oracle and/or its affiliates.