Oracle Exadata

Oracle Exadata Database Service on Cloud@Customer Administrator's Guide

G27855-62
June 2026

Oracle Exadata Oracle Exadata Database Service on Cloud@Customer Administrator's Guide,

G27855-62

Copyright © 2008, 2026, Oracle and/or its affiliates.

Primary Authors: Nirmal Kumar, James Spiller

Contributing Authors: Aneesh Khandelwal, Kasey Parker, Sridhar Murala

Contributors: Abhilash Gudasi, Akhila Prabhu, Alan Williams, Alexander Prince Mathew, Ankit Mahale, Ankur Raina, Ankur Sinha, Ashutosh Chetal, Barb Lundhild, Behkam Aminzadeh, Bhargavi Amancharla, Bob Thome, Boming Zhang, Brad Nowrouzi, Bryce Cracco, Charan Chaudary Lekkalapudi, Cynthia Varela, David Jimenez Alvarez, Douglas Williams, Gabriel Melendez, Gorev Khanna, Guruprasad Hegde, Jai Krishnani, Jean-Francois Verrier, Jose Ricardo Ortiz Olivares, Karina Ledesma, Kevin Deihl, Kishore Sridhar, Kris Bhanushali, Lakshmi Sneha Kandukuri, Lisa Grant, Luqman, Manini Chattopadhyay, Manish Mahawar, Manish Shah, Manu Sinha, Mark Bauer, Michael Fitch, Monika Gupta, Namratha Mandya Subramanya, Nayana Vishwa, Neha Avasthy, Neil Hebert, Nishith Pandey, Nithin Kovoor, Omar Briseno Safa, Oscar Gallegos, Pablo Sainz Albanez, Peter Fusek, Peter Liu, Peter Sciarra, Pradeep Bhat, Pravin Jha, Raj Ojha, Ranganath Srirangapatna Ramachandra, Rhonda Day, Robert Greene, Rodrigo Gonzalez Alba, Rohan Anand, Rudregowda Mallegowda, Sanjay Narvekar, Santosh Uttam Bobade, Saravanan Sabapathy, Sarita Nori, Sham Rao Pavan, Sheila Ray, Shravan Thatikonda, Smitha Gurunathakrishnan, Somnath Lahiri, Sonali Purdhani, Sujit Dhawalkar, Syam Vasa, Tammy Bednar, Terence Buencamino, Tushar Pandit, Vedavathi Eagala, Vedika Joshi, Vira Goorah, Vivek Nama, Vrishali Hajare, Xiao Pan, Youngju Cho

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

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

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

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

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

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

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

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

Contents

1 Oracle Exadata Database Service on Cloud@Customer Overview

Oracle Exadata Database Service on Cloud@Customer Service Description 1
- About Oracle Exadata Database Service on Cloud@Customer 2
- Licensing Considerations for Oracle Exadata Database Service on Cloud@Customer 3
- Per-Second Billing for OCPU Usage 3
- Supported Database Edition and Versions 3
- System and Shape Configuration Options 3
  - System Configuration Options for Oracle Exadata Database Service on Cloud@Customer 4
  - Oracle Exadata X11M System Model Specifications 4
  - Oracle Exadata X10M System Model Specifications 5
  - Oracle Exadata X9M-2 System Model Specifications 6
  - Oracle Exadata X8M-2 System Model Specifications 8
  - Oracle Exadata X8-2 System Model Specifications 9
  - Oracle Exadata X7-2 System Model Specifications 9
- Support Requirements 10

Exadata Cloud Management Interfaces 11
- Introduction to Exadata Cloud Management Interfaces 11
- OCI Control Plane Interfaces 11
- Local VM Command-Line Interfaces 13

Exadata Cloud@Customer Technical Architecture 14

2 What’s New in Oracle Exadata Database Service on Cloud@Customer

Monthly Infrastructure Maintenance Update Preferences 4
UX Updates for Autonomous AI Database on Cloud@Customer 5
Custom ASM Disk Group Sizing 5
Exadata Storage Server X11M Extreme Flash (EF) 6
Customize SCAN Listener Port (TCPS) 6
Support for a 100Gbps NIC for the Exadata Cloud@Customer (ExaDB-C@C) Client Network 7
Exascale VM storage for Exadata Database Service on Cloud@Customer 7
Pluggable Database Snapshots for Exadata Database Service on Cloud@Customer 9
Display ASM Rebalance Progress During Storage Scaling 10
Standalone Backup 10

Copyright © 2008, 2026, Oracle and/or its affiliates.

Real-Time Data Protection on Oracle Zero Data Loss Recovery Appliance (ZDLRA) 11
Database Update Enhancement: Skip Datapatch and Run Independently Later 11
Enhanced Controls to Configure Automatic Backups 12
Enhanced Data Guard Health Insights 13
Support for a 100Gbps NIC for the Exadata Cloud@Customer (ExaDB-C@C) Backup Network 13
Convert a Physical Standby Database to a Snapshot Standby Database and vice versa 13
Exadata System Software 25.1.0.0.0 14
Select Grid Infrastructure Minor Version (N-3) While Provisioning a VM Cluster 15
Expanded Tagging Support Across Database Workflows 15
Support for Concurrent Data Guard Operations in Multiple Standby Environments 15
Schedule VM Cloud Automation Updates 16
Classify and Control Access to Critical OCI REST APIs 16
Oracle Data Guard Setup with Precheck Validation 17
Backup and Restore from a Standby Database in a Data Guard Environment 17
Concurrently Create or Delete Pluggable Databases (PDBs) 18
Exadata Database Service for Developers 18
Multiple Standby Databases 19
Key Management Using an External Keystore 20
X11M System Support 21
Exadata Exascale with Exadata Database Service on Cloud@Customer 22
Microsoft Entra ID (MS-EI) Integration with Oracle Exadata Database Service on Cloud@Customer 23
Support for Concurrent Data Guard, Container Database (CDB), and Pluggable Database (PDB) Operations 24
Delegate Access Control for ExaDB-C@C 25
Cost and Usage Attribution for Pluggable Databases (PDBs) 25
Different RUs for Primary and Standby DB Homes in Data Guard Associations, Switchover, and Failover Operations 26
Granular Permissions for VM Cluster Update Operations 27
Enhancements to Quarterly Exadata Infrastructure Maintenance Planning and Execution 27
Enable Unified Auditing While Creating a Database Home 28
Enhancements to Oracle Key Vault (OKV) Integration with ExaDB-C@C to Manage Transparent Data Encryption (TDE) Keys 28
Ability to Increase the Size of Guest VM Local File Systems 29
Create and Use Custom Software Images 29
Oracle AI Database 26ai on Oracle Exadata Database Service on Cloud@Customer 30
Change the Home Region of ExaDB-C@C Infrastructure 30
Enhancements to Serial Console Functionality 31
Provision a VM Cluster with Either an OL7 or OL8-Based Image 32
VM Cluster on a Single VM 32
Enhancement to Pluggable Database (PDB) Management 32
Manage Administrator (SYS User) and TDE Wallet Passwords 33

Copyright © 2008, 2026, Oracle and/or its affiliates.

Update Guest VM Operating System to Oracle Linux 8 33
Exadata Fleet Update 33
Oracle Key Vault (OKV) Integration with ExaDB-C@C to Manage Transparent Data Encryption (TDE) Keys 34
X10M System Support 34
Manage Serial Console Access to ExaDB-C@C Systems 35
Interim Software Updates 36
Link Aggregation Control Protocol (LACP) Support for Client and Backup Networks 36
VM Cluster and Database Health and Performance Metrics in the OCI Console 37
Enable Data Guard with Database Software Image (Custom Image) 38
Use Identity and Access Management (IAM) Authentication with Oracle Exadata Database Service on Cloud@Customer 38
Create Data Guard Associations Across Regions 39
Multi-Rack Elastic Compute and Storage Expansion 39
Automatic Diagnostic Collection 41
Elastic Compute Expansion 42
Oracle Standard Tagging for Resources on Oracle Exadata Database Service on Cloud@Customer 43
Exadata Infrastructure Maintenance History 44
Monthly ExaDB-C@C Infrastructure Security Maintenance 44
Enhanced Control to Rollback or Retry Failed Guest VM Operating System Update 45
Concurrently Create or Terminate Oracle Databases in a VM Cluster 45
Support for Rack Serial Number as a System Tag 45
Support for DB Home Minor Version Selection (N-3) 46
VM Guest Exadata OS Image Major Version Update 46
Database Service Events 46
Control Plane Server (CPS) Offline Diagnostic Report 47
Enhanced Infrastructure Maintenance Controls 48
Manage Pluggable Databases on Exadata Cloud@Customer 49
Allow Customers to Choose Data Guard Type 49
Specify the Same SID for Primary and Standby Databases in a Data Guard Association 50
Specify db_unique_name and SID for Primary and Standby Databases in Data Guard Association 50
VM Cluster Node Subsetting 50
X9M-2 System Support 51
Customize SCAN Listener Port 52
Creating DG Association/Standby Database Using Existing Database Home 52
Upgrading Oracle Grid Infrastructure on an Exadata Cloud@Customer VM Cluster 52
Updating Guest VM Operating System 53
Upgrading Oracle Databases 53
Download Network Validation Report 53
Elastic Storage Expansion 53
Oracle Database Software Images 55

Copyright © 2008, 2026, Oracle and/or its affiliates.

Exadata Cloud@Customer Infrastructure Patching Automation 55
Customer Maintenance Contacts 55
X8M-2 System Support 56
Enable and Manage Data Guard Associations 56
Oracle Exadata Cloud@Customer Deployment Assistant 57
Oracle Grid Infrastructure and Oracle Database Patching 57
Per-Second Billing for OCPU Usage 58
Shared Database Home Resource Tags 58
Create and Manage Multiple Virtual Machines per Exadata System (MultiVM) 58
Scale OCPUs Without Cloud Connectivity 59
Configure Oracle Database Character Set and National Character Set 59
Specify a Time Zone While Provisioning Oracle Exadata Database Service on Cloud@Customer Infrastructure 60
Shared Database Homes for Oracle Exadata Database Service on Cloud@Customer 60
X7-2 System Support 61

3 Preparing for Oracle Exadata Database Service on Cloud@Customer

Oracle Cloud Infrastructure (OCI) Requirements for Oracle Exadata Database Service on Cloud@Customer 1
Required IAM Policy for Oracle Exadata Database Service on Cloud@Customer 2
Site Requirements for Oracle Exadata Database Service on Cloud@Customer 3
Space Requirements for Oracle Exadata Database Service on Cloud@Customer 4
Weight of Oracle Exadata Database Service on Cloud@Customer Racks 4
Receiving, Unpacking, and Access for Oracle Exadata Database Service on Cloud@Customer Racks 4
Flooring for Oracle Exadata Database Service on Cloud@Customer Racks 5
Electrical Power for Oracle Exadata Database Service on Cloud@Customer Racks 5
Temperature and Humidity Ranges for Oracle Exadata Database Service on Cloud@Customer 7
Ventilation for Oracle Exadata Database Service on Cloud@Customer Racks 8
Network Requirements for Oracle Exadata Database Service on Cloud@Customer 8
Network Requirements for Oracle Exadata Database Service on Cloud@Customer 9
Data Center Network Services for Oracle Exadata Database Service on Cloud@Customer 17
IP Addresses and Subnets for Oracle Exadata Database Service on Cloud@Customer 17
Uplinks for Oracle Exadata Database Service on Cloud@Customer 19
Network Cabling for Oracle Exadata Database Service on Cloud@Customer 20
Establish a Secure Connection Between Your CPS and OCI Using OCI's FastConnect Service 20
Storage Configuration Requirements for Oracle Exadata Database Service on Cloud@Customer 23
Configuring Exadata ASM Storage 23

Copyright © 2008, 2026, Oracle and/or its affiliates.

About ASM Storage Configuration for Oracle Exadata Database Service on Cloud@Customer 24
Allocation of Storage Space Options on Oracle Exadata Storage Servers 24
Allocation Proportions for DATA, RECO and SPARSE Disk Groups 25
Virtual Machine File System Structure for Oracle Exadata Database Service on Cloud@Customer 26
Configuring Exadata Exascale Storage 26
About Exascale Storage Configuration Oracle Exadata Database Service on Cloud@Customer 27
Using the Console to Configure Exascale Storage on Oracle Exadata Database Service on Cloud@Customer Infrastructure 28
Using the Console to Create an Exascale Storage Vault 29
Using the Console to Scale an Exascale Storage Vault 30
Using the Console to Move an Exascale Storage Vault to Another Compartment 30
Using the Console to Delete an Exascale Storage Vault 31
Using the Console to Configure Exascale VM Storage on Oracle Exadata Database Service on Cloud@Customer Infrastructure 31
Configure All Future VM Backups to Use Exascale Storage on a VM Cluster 32
Create an Advanced Cluster File System (ACFS) Using Exascale Block Volume 32
Deconfigure Exascale Storage on the Exascale Infrastructure 33
Checklists for Oracle Exadata Database Service on Cloud@Customer Deployments 33
System Components Checklist for Oracle Exadata Database Service on Cloud@Customer 34
Data Center Room Checklist for Oracle Exadata Database Service on Cloud@Customer 34
Data Center Environment Checklist for Oracle Exadata Database Service on Cloud@Customer 35
Access Route Checklist for Oracle Exadata Database Service on Cloud@Customer 35
Facility Power Checklist for Oracle Exadata Database Service on Cloud@Customer 36
Safety Checklist for Oracle Exadata Database Service on Cloud@Customer 36
Logistics Checklist for Oracle Exadata Database Service on Cloud@Customer 36
Network Configuration Checklist for Oracle Exadata Database Service on Cloud@Customer 38
Reracking Checklist for Oracle Exadata Database Service on Cloud@Customer 38

4 Getting Started with Oracle Exadata Database Service on Cloud@Customer Deployment

Using the Console to Create API Access Control 2
About Provisioning Oracle Exadata Database Service on Cloud@Customer Systems 2
Tagging Oracle Exadata Database Service on Cloud@Customer Resources 3
Overview of Elastic Storage Expansion 8
Using the Console to Provision Oracle Exadata Database Service on Cloud@Customer 12
Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure 12

Copyright © 2008, 2026, Oracle and/or its affiliates.

Enabling or Disabling the Control Plane Server Diagnostic Offline Report 22
Viewing the Control Plane Server Diagnostic Offline Report 22
Using the Console to View Exadata Infrastructure Network Configuration Details 23
Using the Console to Edit Oracle Exadata Database Service on Cloud@Customer Infrastructure Networking Configuration 23
Using the Console to Download a File Containing Configuration Data 25
Using the Console to Activate Oracle Exadata Database Service on Cloud@Customer Infrastructure 26
Using the Console to Check the Status of Oracle Exadata Database Service on Cloud@Customer Infrastructure 26
Creating First VM Cluster Network on Oracle Exadata Database Service on Cloud@Customer 27
About Managing VM Cluster Networks on Oracle Exadata Database Service on Cloud@Customer 27
Using the Console to Create a VM Cluster Network 28
Using the Console to Edit a VM Cluster Network 31
Using the Console to Download a File Containing the VM Cluster Network Configuration Details 33
Using the Console to Validate a VM Cluster Network 33
Provisioning the First VM Cluster on an Oracle Exadata Database Service on Cloud@Customer System 34
About Managing VM Clusters on Oracle Exadata Database Service on Cloud@Customer 34
Prerequisites for VM Clusters on Oracle Exadata Database Service on Cloud@Customer 35
Using the Console to Create an ASM VM Cluster 35
Using the Console to Create an Exascale VM Cluster 44
Creating Database Backup Destinations for Oracle Exadata Database Service on Cloud@Customer 52
About Managing Backup Destinations for Oracle Exadata Database Service on Cloud@Customer 52
Prerequisites for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer 53
Using the Console for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer 54
Using the Console to Create a Backup Destination 55
Using the Console to Edit a Backup Destination 57
Using the Console to Move a Backup Destination to Another Compartment 58
Using the Console to Delete a Backup Destination 58
Creating First Database Home on an Oracle Exadata Database Service on Cloud@Customer System 59
About Creating Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer System 59
Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer 60

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the API to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer 62
Creating First Database on an Oracle Exadata Database Service on Cloud@Customer System 62
Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer 63
Oracle Database Releases Supported by Oracle Exadata Database Service on Cloud@Customer 63
About Provisioning and Configuring Oracle Databases on Oracle Exadata Database Service on Cloud@Customer 64
Using the Console to Create a Database 65
Connecting to an Oracle Exadata Database Service on Cloud@Customer System 72
Connecting to a Virtual Machine with SSH 72
Prerequisites for Connecting to an Oracle Exadata Database Service on Cloud@Customer System 73
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY 73
Accessing a Database After You Connect to the Virtual Machine 74
Connecting from a Unix-Style System 76
Connecting to a Database with Oracle Net Services 76
Using Oracle Net Services to Connect to a Database 77
Prerequisites for Connecting to a Database with Oracle Net Services 77
Connecting to a Database Using SCAN 77
Connecting to a Database Using a Node Listener 79

5 How-to Guides

Manage Database Security with Oracle Data Safe 2
About Oracle Data Safe 2
Get Started 3
Using Oracle Data Safe 4
Connect to the Oracle Exadata Database Service on Cloud@Customer Service 5
Connecting to a Virtual Machine with SSH 6
Prerequisites for Connecting to an Oracle Exadata Database Service on Cloud@Customer System 6
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY 7
Connecting from a Unix-Style System 8
Accessing a Database After You Connect to the Virtual Machine 8
Connecting to a Database with Oracle Net Services 9
Using Oracle Net Services to Connect to a Database 10
Prerequisites for Connecting to a Database with Oracle Net Services 10
Connecting to a Database Using SCAN 11
Connecting to a Database Using a Node Listener 12
Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure 13
About Provisioning Oracle Exadata Database Service on Cloud@Customer Systems 14

Copyright © 2008, 2026, Oracle and/or its affiliates.

Defining Maintenance Contacts for Exadata Cloud@Customer 14
Add Your Primary Maintenance Contact 15
Add Secondary Contacts for Cloud Notifications (CNs) 16
Edit Maintenance Contacts 16
Remove a Secondary Contact 17
About Changing the Home Region of Oracle Exadata Database Service on Cloud@Customer Infrastructure 17
Overview of Elastic Storage Expansion 19
Using the Console to Provision Oracle Exadata Database Service on Cloud@Customer 23
Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure 24
Using the Console to Edit Network Bonding 34
Using the Console to View Exadata Infrastructure Network Configuration Details 35
Using the Console to Edit Oracle Exadata Database Service on Cloud@Customer Infrastructure Networking Configuration 35
Using the Console to Download a File Containing Configuration Data 37
Using the Console to Activate Oracle Exadata Database Service on Cloud@Customer Infrastructure 38
Using the Console to Check the Status of Oracle Exadata Database Service on Cloud@Customer Infrastructure 38
Using the Console to Scale Infrastructure Compute and Storage 39
Using the Console to Download Scale Infrastructure Compute and Storage Configuration File 42
Using the Console to Activate New Database and Storage Servers 42
Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption 43
Using the Console to View Details of Oracle Exadata Database Service on Cloud@Customer Infrastructure with Scaled Compute and Storage Capacity 44
Using the Console to Move Oracle Exadata Database Service on Cloud@Customer Infrastructure 44
Using the Console to Terminate Oracle Exadata Database Service on Cloud@Customer Infrastructure 45
Using the Console to Manage Tags for Your Oracle Exadata Database Service on Cloud@Customer Resources 45
Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure 46
Configure Oracle-Managed Infrastructure Maintenance 47
About Oracle Managed Oracle Exadata Database Service on Cloud@Customer Infrastructure Maintenance Updates 48
Overview of the Quarterly Infrastructure Maintenance Process 49
Overview of Monthly Security Maintenance 50
Understanding Monthly and Quarterly Maintenance in the Same Month 51
Infrastructure Maintenance Contacts 53
Maintenance Scheduling Policy 54
View the List of Maintenance Scheduling Policy 55

Copyright © 2008, 2026, Oracle and/or its affiliates.

Create a Maintenance Scheduling Policy 55
View Details of a Maintenance Scheduling Policy 57
Edit a Maintenance Scheduling Policy 57
Edit a Maintenance Scheduling Policy that's in Needs Attention State 57
View the Maintenance Windows of a Maintenance Scheduling Policy 57
Edit a Maintenance Window of a Maintenance Scheduling Policy 58
Add Additional Maintenance Windows to a Maintenance Scheduling Policy 58
View Resources Associated with a Policy 58
Delete a Maintenance Window of a Maintenance Scheduling Policy 59
Move a Maintenance Scheduling Policy to a Different Compartment 59
Add Tags to a Maintenance Scheduling Policy 60
Delete a Maintenance Scheduling Policy 60
Using the Console to Configure Oracle-Managed Infrastructure Updates 61
View or Edit Quarterly Maintenance Preferences 61
View or Edit Monthly Maintenance Preferences 65
Manage Quarterly/Monthly Maintenance Plan using Scheduling Policy 66
View or Edit a Scheduled Quarterly Maintenance 71
View or Edit a Scheduled Security Maintenance 75
View the Maintenance History 76
Manage Quarterly/Monthly Maintenance Run created from Scheduling Plan 77
View Maintenance Windows Associated with a Maintenance Run 77
Edit Maintenance Window Associated with a Maintenance Run 79
View Maintenance Actions Associated with a Maintenance Run 79
Edit Maintenance Actions of a Maintenance Window Associated with a Maintenance Run 80
View and Edit Maintenance While Maintenance is In Progress 82
(Applies to quarterly maintenance) View and Edit Maintenance While Maintenance is Waiting for Custom Action 83
Cancel a Maintenance Run While In Progress 84
Start a Maintenance On Demand 85
View Maintenance Activity in a Compartment 85
View the Details of a Maintenance Activity 85
View Maintenance History in a Compartment 86
View Maintenance History Details in a Compartment 86
Review and Respond to Unplanned Maintenance Activity 86
Monitor Infrastructure Maintenance Using Lifecycle State Information 87
Receive Notifications about Your Infrastructure Maintenance Updates 88
Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure Maintenance Controls 88
Manage VM Cluster Networks 90
About Managing VM Cluster Networks on Oracle Exadata Database Service on Cloud@Customer 91
Using the Console to Create a VM Cluster Network 91

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to View VM Cluster Network Details 94
Using the Console to View SCAN Listener Port Configured 95
Using the Console to Edit a VM Cluster Network 95
Using the Console to Download a File Containing the VM Cluster Network Configuration Details 97
Using the Console to Validate a VM Cluster Network 97
Using the Console to Download Network Validation Report 98
Using the Console to Terminate a VM Cluster Network 99
Using the Console to Add Another DB Server to the VM Cluster Network 100
Using the Console to Remove DB Servers from a VM Cluster Network 100
Manage VM Clusters 101
About Managing VM Clusters on Oracle Exadata Database Service on Cloud@Customer 101
Overview of VM Cluster Node Subsetting 102
Overview of Automatic Diagnostic Collection 103
Incident Logs and Trace Files 104
Health Metrics 109
Introduction to Scale Up or Scale Down Operations 114
Scaling Up or Scaling Down the VM Cluster Resources 114
Resizing Memory and Large Pages 116
Calculating the ASM Storage 117
Estimating How Much Local Storage You Can Provision On Your VMs 119
Scaling Local Storage 120
Using the Console to Manage VM Clusters on Oracle Exadata Database Service on Cloud@Customer 121
Using the Console to Create an ASM VM Cluster 122
Using the Console to Create an Exascale VM Cluster 131
VM Cloud Automation Software Update Management 139
Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection 142
Using the Console to Add VMs to a Provisioned Cluster 143
Using the Console to View a List of DB Servers on an Exadata Infrastructure 144
Using the Console to Remove a VM from a VM Cluster 145
Using the Console to Update the License Type on a VM Cluster 145
Using the Console to Add SSH Keys After Creating a VM Cluster 146
Using the Console to Scale the Resources on a VM Cluster 146
Using the Console to Stop, Start, or Reboot a VM Cluster Virtual Machine 149
Using the Console to Check the Status of a VM Cluster Virtual Machine 149
Using the Console to Move a VM Cluster to Another Compartment 150
Using the Console to Terminate a VM Cluster 150
Using the API to Manage Oracle Exadata Database Service on Cloud@Customer VM Clusters 151
Troubleshooting Virtual Machines Using Console Connections 152
Required IAM Policies 153

Copyright © 2008, 2026, Oracle and/or its affiliates.

Prerequisites 153
Create the Virtual Machine Serial Console Connection 157
Make an SSH Connection to the Serial Console 157
Using Cloud Shell to Connect to the Serial Console 162
Displaying the Console History for a Virtual Machine 163
Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems 166
Exiting the Virtual Machine Serial Console Connection 168
Manage Software Images 169
Creation and Storage of Software Images 170
Using a Software Image with an Oracle Exadata Database Service on Cloud@Customer System 170
Using the Console to View the List of Software Images 171
Using the Console to Create a Database Software Image 171
Using the Console to Create a Grid Infrastructure Software Image 172
Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Database Home or Oracle Grid Infrastructure Home 173
Using the Console to Delete a Software Image 174
Using the Console to View the Update Information of a Software Image 174
Using the Console to Move a Software Image to a Different Compartment 174
Using the API for Managing Software Images 175
Create Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer System 175
About Creating Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer System 176
Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer 176
Using the API to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer 178
Manage Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems 179
About Managing Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems 179
Manage Database Home Using the Console 179
Using the Console to View Information About an Oracle Database Home 180
Using the Console to Delete an Oracle Database Home 180
Using the API to Manage Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer 181
Differences Between Managing Resources with dbaascli and the Database API 181
Manage Databases on Oracle Exadata Database Service on Cloud@Customer 182
Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer 183
Oracle Database Releases Supported by Oracle Exadata Database Service on Cloud@Customer 183

Copyright © 2008, 2026, Oracle and/or its affiliates.

About Provisioning and Configuring Oracle Databases on Oracle Exadata Database Service on Cloud@Customer 184
Using the Console to Manage Databases on Oracle Exadata Database Service on Cloud@Customer 185
Using the Console to Create a Database 186
Using the Console to Manage SYS User and TDE Wallet Passwords 192
Using the Console to Migrate Encryption Keys from Oracle Wallet to Other Encryption Methods (Oracle Key Vault, External Keystore) 193
Using the Console to Rotate Encryption Keys 194
Using the Console to Move a Database to Another Database Home 194
To run datapatch on a container database (CDB) 195
To run datapatch on a pluggable database (PDB) 196
Using the Console to Terminate a Database 197
Using the API to Manage Oracle Database Components 197
Changing the Database Passwords 198
To Change the SYS Password for an Exadata Database Service on Cloud@Customer Database 199
To Change Database Passwords in a Data Guard Environment 199
To Change the TDE Wallet Password for an Exadata Database Service on Cloud@Customer Database 199
Manage Pluggable Databases on Oracle Exadata Database Service on Cloud@Customer 199
Pluggable Database Operations 200
Concurrently Create or Delete Pluggable Databases (PDBs) 227
Cost and Usage Attribution for Pluggable Databases (PDBs) 228
Connect to an Oracle Database using Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users 230
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database 230
Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database 233
Enabling the Database and Clients for IAM Integration 235
Authenticating and Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Cloud@Customer 236
About Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Cloud@Customer 236
Configuring the Oracle Database for Microsoft Entra ID (MS-EI) Integration 237
Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer 240
Backup Destinations 240
About Managing Backup Destinations for Oracle Exadata Database Service on Cloud@Customer 241
Prerequisites for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer 242
Using the Console for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer 243

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Backup Destinations 247
Oracle Database Backup Methods in Exadata Cloud 247
Oracle Managed Backup 248
User Configured Backup 261
Configuring and Customizing Backups with dbaascli 269
Customizing Backup Settings by Using a Generated Configuration File 270
Creating an On-Demand Backup by Using dbaascli 277
Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management 278
Customizing Real Time Redo Transport (RTRT) Behavior for Recovery Appliance Backups 278
Alternative Backup Methods 279
Recovering a Database Using Oracle Recovery Manager (RMAN) 280
Standalone Backup 281
Using the Console to Create a Standalone Backup 281
Patch and Update an Oracle Exadata Database Service on Cloud@Customer System 282
Perform User Managed Maintenance Updates 282
Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System 283
Patching and Updating VM Clusters and Database Homes 283
Updating Guest VM Operating System 289
Upgrading Oracle Grid Infrastructure on an Oracle Exadata Database Service on Cloud@Customer VM Cluster 293
Upgrading Oracle Databases 296
Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System Manually 301
Updating Software Manually 302
Updating the Guest VM Operating System Manually 302
Resolving Dependency Issues Associated with Additional Non-Exadata Software Packages for Guest VM Upgrade 311
Interim Software Updates 312
Create an Interim Software Update 312
Download an Interim Software Update 313
Delete an Interim Software Update 313
Move an Interim Software Update Resource to Another Compartment 313
Using the API to Manage Interim Software Updates 314
Use Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer 314
About Using Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer 315
Prerequisites for Using Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer 316
VM Clusters 316
Password 316
Adding a Node to a VM Cluster 316

Copyright © 2008, 2026, Oracle and/or its affiliates.

Removing a Node from a VM Cluster 316
Working with Data Guard 317
Switchover 317
Failover 317
Reinstate 317
Support for Concurrent Data Guard Operations in Multiple Standby Environments 318
Enhanced Data Guard Health Status Reporting 318
Using the Console to Manage an Oracle Data Guard Group 320
Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System 321
Using the Console to View Data Guard Group of Databases in an Exadata VM Cluster 325
Using the Console To View and Edit Data Guard Group Details 326
Using the Console To Perform a Database Switchover 326
Using the Console To Perform a Database Failover 327
Using the Console To Reinstate a Database 327
Using the Console To Terminate a Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer System 328
Using the API to Manage Data Guard Associations on an Oracle Exadata Database Service on Cloud@Customer System 328
Using the API to Manage Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer System 329
Manage Keys Using an External Keystore 346
Configuring an External Keystore 346
Storing the Keys in an External Keystore 347
Restrictions on Adding a Virtual Machine to a VM Cluster Configured with an External Keystore 347
Managing Encryption Keys on External Devices 348
Customer-Managed Keys in Oracle Exadata Database Service on Cloud@Customer 349
About Oracle Key Vault 349
Overview of Key Store 350
Required IAM Policy for Managing OKV on Oracle Exadata Database Service on Cloud@Customer 350
Tagging Resources 351
Moving Resources to a Different Compartment 351
Setting Up Your Oracle Exadata Database Service on Cloud@Customer to Work With Oracle Key Vault 351
Step 1: Create a Vault in OCI Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password 352
Step 2: Create a Policy Statement for Group to Access Secret in OCI Vault Service 352
Step 3: Create Key Store 352
Managing Your Key Store 353
View Key Store Details 354
Edit Key Store Details 354

Copyright © 2008, 2026, Oracle and/or its affiliates.

Move a Key Store to Another Compartment 354
Delete a Key Store 355
View Key Store Associated Container Database Details 355
Using the API to Manage Key Store 355
Administer Transparent Data Encryption (TDE) Keys 356
How to Manually Clone a Pluggable Database (PDB) from a Remote Container Database (CDB) When Data is Encrypted with Master Encryption Key (MEK) in Oracle Key Vault (OKV) 357
Source CDB and Target CDB are Encrypted with MEK in the Same OKV Server 358
Source CDB and Target CDB are Encrypted with MEK in a Different OKV Server 368
How to Upgrade Oracle Key Vault (OKV) Home in ExaDB-C@C 378
Migrate to Oracle Exadata Database Service on Cloud@Customer 383
Moving to Oracle Cloud Using Zero Downtime Migration 383
Overview of Exadata Cloud@Customer Gen1 to Out-of-Place Cloud Upgrade to Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure 383
Scope for Exadata Cloud@Customer Gen1 to Gen2 Out-of-Place Cloud Upgrade 384
Hardware and Software Required for Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure 385
Using Oracle Zero Downtime Migration (ZDM) to Migrate Oracle Databases 386
During Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure 391
Post Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure 391
Best Practices for Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure 392

6 Oracle Autonomous AI Database on Exadata Cloud@Customer

What’s New in Oracle Autonomous AI Database on Exadata Cloud@Customer 1
Introduction to Oracle Autonomous AI Database on Exadata Cloud@Customer 1

7 Reference Guides for Oracle Exadata Database Service on Cloud@Customer

Using the dbaascli Utility with Oracle Exadata Database Service on Cloud@Customer 1
About Using the dbaascli Utility on Oracle Exadata Database Service on Cloud@Customer 2
Creating Oracle Database Using dbaascli 3
Listing Available Software Images and Versions for Database 3
Creating Oracle Database Home 4
Creating Oracle Database In the Specified Oracle Database Home 5
Changing the Database Passwords 8
To Change the SYS Password for an Exadata Database Service on Cloud@Customer Database 8

Copyright © 2008, 2026, Oracle and/or its affiliates.

To Change Database Passwords in a Data Guard Environment 8
To Change the TDE Wallet Password for an Exadata Database Service on Cloud@Customer Database 9
Managing Oracle Exadata Database Service on Cloud@Customer Software Images Using the Dbaascli Utility 9
Listing Available Software Images and Versions for Database and Grid Infrastructure 9
To download a software image 11
Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli 11
Patching Databases using dbaascli 12
Patching Oracle Grid Infrastructure 14
Listing Available Software Images and Versions for Database and Grid Infrastructure 15
Performing a Precheck Before Patching Databases and Grid Infrastructure 16
Resuming or Rolling Back a Patching Operation 19
Updating Cloud Tooling Using dbaascli 20
Creating a Duplicate Database 21
Using dbaascli to Duplicate a Cloud Database 21
Duplicate an On-Premises Database 22
Release Notes 24
Release 26.1.1.0.0 (260302) 25
Release 25.4.1.0.0 (251125.0831, 251125.0837) 25
Release 25.3.1.0.0 (250826) 25
Release 25.2.1.0.0 (250612) 25
Release 25.1.2.0.0 (250505) 25
Release 25.1.1.0.0 (250220) 25
Release 24.4.1.0.0 (241212) 25
Release 24.3.2.0.0 (240912) 26
Release 24.3.1.0.0 (240730) 26
Release 24.2.1.0.0 (240620) 26
Release 24.1.2.0.0 (240327, 240424, 240502) 26
Release 24.1.1.0.0 (240118, 240219) 26
Release 23.4.1.0.0 (231219) 26
Release 23.3.2.0.0 (231115) 27
Release 23.3.1.0.0 (230817, 231020) 27
Release 23.2.1.0.0 (230708, 230724) 27
Release 23.1.2.0.0 (230411, 230616) 28
Release 23.1.1.0.1 (230302) 28
Release 22.4.1.0.1 (221214) 28
Release 22.3.1.1.0 (221003) 29
Release 22.3.1.0.1 (220831) 30
Release 22.2.1.1.0 (220713) 30
Release 22.2.1.0.1 (220504) 31
Release 22.1.1.1.0 (220301) 32

Copyright © 2008, 2026, Oracle and/or its affiliates.

Release 22.1.1.0.1 (220223) 32
Release 21.4.1.1.0 33
Release 21.3.1.2.0 34
Release 21.3.1.1.0 34
Release 21.3.1.0.1 34
Release 21.2.1.x.x 35
dbaascli Command Reference 36
Administration and Configuration 38
CPU Scaling 41
Cloud Software Library (CSWLIB) Management 45
Database Management 51
Data Guard Management 88
Database Home Management 104
Diagnostics and Health Checks 112
Grid Infrastructure Management 118
Patching and Upgrading 130
Pluggable Database (PDB) Management 154
System Management 188
Transparent Data Encryption (TDE) Management 190
Deprecated dbaascli commands 226
Monitoring and Managing Exadata Storage Servers with ExaCLI 230
About the ExaCLI Command 230
Exadata Storage Server Username and Password 230
ExaCLI Command 231
Connecting to a Storage Server with ExaCLI 237
Monitor Metrics for VM Cluster Resources 237
Prerequisites for Using Metrics 238
View Metrics for a VM Cluster 239
View Metrics for a Database 240
View Metrics for VM Clusters in a Compartment 241
View Metrics for Databases in a Compartment 242
Manage Oracle Trace File Analyzer 242
Metrics for Oracle Exadata Database Service on Cloud@Customer in the Monitoring Service 243
Oracle Exadata Database Service on Cloud@Customer Events 249
About Event Types on Oracle Exadata Database Service on Cloud@Customer 251
Exadata Infrastructure Event Types 251
Exascale DB Storage Vaults Event Types 267
VM Cluster Network Event Types 267
VM Cluster Event Types 269
Backup Destination Event Types 270
Database Node Event Types (Oracle Exadata Database Service on Cloud@Customer) 271
Database Home Event Types (Oracle Exadata Database Service on Cloud@Customer) 272

Copyright © 2008, 2026, Oracle and/or its affiliates.

Database Event Types (Oracle Exadata Database Service on Cloud@Customer) 273
Database and Grid Infrastructure Patching Event Types 275
Data Guard Event Types 278
Key Store Event Types 281
Exadata Infrastructure Maintenance Event Types 282
Storage Expansion Event Types 288
Database Software Images Event Types 290
Database Upgrade Event Types 298
Pluggable Database Event Types 306
VM Node Subsetting Event Types 316
Database Service Events 320
Overview of Database Service Events 320
Receive Notifications about Database Service Events 323
Database Service Event Types 324
Temporarily Restrict Automatic Diagnostic Collections for Specific Events 335
Remediation 339
Compute Expansion Event Types 360
Serial Console Connection Event Types 365
Viewing Audit Log Events 369
Serial Console History Event Types 370
Interim Software Updates Event Types 375
Policy Details for Oracle Exadata Database Service on Cloud@Customer 381
About Resource-Types 381
Resource-Types for Oracle Exadata Database Service on Cloud@Customer 382
Supported Variables 384
Details for Verb + Resource-Type Combinations 384
Database-Family Resource Types 385
Soft Deprecation of Certain Database Service APIs and Fields 385
Permissions and API operation details for Exadata Infrastructures 386
Permissions and API operation details for exascale-db-storage-vaults 387
Permissions and API operation details for VM Cluster Networks 388
Permissions and API operation details for VM Clusters 388
Permissions and API operation details for Backup Destinations 390
Permissions and API operation details for DB Nodes 391
Permissions and API operation details for DB Homes 391
Permissions and API operation details for Databases 392
Permissions and API operation details for Backups 393
Permissions and API operation details for Database Software Image 393
Permissions and API operation details for Key Stores 394
Permissions and API operation details for Pluggable Databases (PDBs) 394
Permissions and API operation details for DB Servers 396
Permissions and API operation details for DB Node Console Connection 396

Copyright © 2008, 2026, Oracle and/or its affiliates.

Permissions and API operation details for DB Node Console History 397
Permissions and API operation details for Interim Software Updates 397
Permissions and API operation details for Scheduling Policies 398
Permissions and API operation details for Scheduling Windows 398
Permissions and API operation details for Scheduling Plan 399
Permissions and API operation details for Scheduled Action 400
Permissions and API operation details for Execution Windows 402
Permissions and API operation details for Execution Action 403
Permissions and API operation details for Data Guard Group 405
Permissions Required for Each API Operation 405
Managing Exadata Resources with Oracle Enterprise Manager Cloud Control 411
Overview of Oracle Enterprise Manager Cloud Control 412
Features of Enterprise Manager Cloud Control 412
Observability and Management for Oracle Exadata Database Service on Cloud@Customer 413
Security Guide for Oracle Exadata Database Service on Cloud@Customer Systems 413
Security Configurations and Default Enabled Features 413
Responsibilities 414
Guiding Principles Followed for Security Configuration Defaults 414
Security Features 415
Guest VM Default Fixed Users 416
Guest VM Default Fixed Groups 418
Guest VM Default Security Settings 421
Guest VM Default Processes 422
Guest VM Network Security 423
OS/VM backups 425
Additional Procedures for Updating Security Posture 425
Customer Responsibilities 425
Enabling Additional Security Capabilities 426
EU Data Act Compliance 430
Troubleshooting Oracle Exadata Database Service on Cloud@Customer Systems 430
Patching Failures on Oracle Exadata Database Service on Cloud@Customer Systems 431
Determining the Problem 431
Troubleshooting and Diagnosis 431
Obtaining Further Assistance 432
Collecting Cloud Tooling Logs 433
Collecting Oracle Diagnostics 433
VM Operating System Update Hangs During Database Connection Drain 433
Adding a VM to a VM Cluster Fails 435
Nodelist is not Updated for Data Guard-Enabled Databases 435
CPU Offline Scaling Fails 436
Standby Database Fails to Restart After Switchover in Oracle Database 11g Oracle Data Guard Setup 436

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using Custom SCAN Listener Port With Data Guard On Disaster Recovery Network Causes Data Guard Association Verification Failures 437
PDB Creation Fails After Moving Database to a New DB Home (26ai) 437
Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel 437
Data Guard Association Creation Fails When Using DR Network 438

Copyright © 2008, 2026, Oracle and/or its affiliates.

1

Oracle Exadata Database Service on Cloud@Customer Overview

This topic is an overview of the Oracle Exadata Database Service on Cloud@Customer service.

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

Oracle Exadata Database Service on Cloud@Customer Service Description

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

> **Note**
>
> For information about Oracle Autonomous AI Database on Exadata Cloud@Customer, see Oracle Autonomous AI Database on Exadata Cloud@Customer.

- About Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is one of the Database services offered on Oracle Cloud Infrastructure. Oracle offers both autonomous and co-managed Oracle Database cloud solutions. For more information, see Overview of the Database Service.
- Licensing Considerations for Oracle Exadata Database Service on Cloud@Customer
  Subscription to Exadata Database Service on Cloud@Customer 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 Cloud@Customer.
- Per-Second Billing for OCPU Usage
  Oracle Exadata Database Service on Cloud@Customer Gen2 uses for OCPUs.
- Supported Database Edition and Versions
  Learn about the versions of Oracle Database that Oracle Exadata Database Service on Cloud@Customer supports.
- System and Shape Configuration Options
  Review the list of Exadata System Shapes
- Support Requirements

Copyright © 2008, 2026, Oracle and/or its affiliates.

About Oracle Exadata Database Service on Cloud@Customer

Oracle Exadata Database Service on Cloud@Customer is one of the Database services offered on Oracle Cloud Infrastructure. Oracle offers both autonomous and co-managed Oracle Database cloud solutions. For more information, see Overview of the Database Service.

> **Note**
>
> For more information on technical architecture, see *Oracle Exadata Cloud@Customer (ExaC@C) Technical Architecture*.

With Oracle Exadata Database Service on Cloud@Customer, you can maintain absolute control over your data while leveraging the combined capabilities of Oracle Exadata and Oracle Cloud Infrastructure managed by Oracle.

Oracle Exadata Database Service on Cloud@Customer enables you to apply the combined power of Oracle Exadata and Oracle Cloud Infrastructure inside your own data center. You have full access to the features and capabilities of Oracle Database along with the intelligent performance and scalability of Oracle Exadata, but with Oracle owning and managing the Exadata infrastructure. You can use the Oracle Cloud Infrastructure console and APIs to manage Oracle Exadata Database Service on Cloud@Customer just as with any other cloud resource, while maintaining absolute sovereignty over your data.

Each Oracle Exadata Database Service on Cloud@Customer system configuration contains Oracle Exadata Database Servers and Oracle Exadata Storage Servers that are interconnected using a high-speed, low-latency RDMA fabric network, and intelligent Oracle Exadata software.

Oracle Exadata Database Service on Cloud@Customer uses virtual machine (VM) technology to separate the customer-managed and Oracle-managed components on each database server. You have `root` privilege for the Oracle Exadata database server VMs, so you can manage the Oracle Database, Oracle Grid Infrastructure, and Oracle Exadata system software. However, you do not have administrative access to the physical database server hardware, which Oracle administers.

Oracle Exadata Database Service on Cloud@Customer uses Oracle Exadata Storage Servers for database storage. The storage is allocated to disk groups managed by Oracle Automatic Storage Management (Oracle ASM). You have full administrative access to the Oracle ASM disk groups, but Oracle administers the Oracle Exadata Storage Server hardware and software.

In addition to the database server hardware and Oracle Exadata Storage Servers, Oracle also manages other Oracle Exadata Database Service on Cloud@Customer infrastructure components, including the network switches, power distribution units (PDUs), and integrated lights-out management (ILOM) interfaces.

On each Oracle Exadata Database Service on Cloud@Customer system, you can create one or more databases. Apart from the inherent storage and processing capacity of your Oracle Exadata system, there is no set maximum for the number of databases that you can create.

Related Topics

- Overview of the Database Service
- Oracle Exadata Cloud@Customer (ExaC@C) Technical Architecture

Copyright © 2008, 2026, Oracle and/or its affiliates.

Licensing Considerations for Oracle Exadata Database Service on Cloud@Customer

Subscription to Exadata Database Service on Cloud@Customer 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 Cloud@Customer.

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

Per-Second Billing for OCPU Usage

Oracle Exadata Database Service on Cloud@Customer Gen2 uses for OCPUs.

Per-second billing means that OCPU usage is billed by the second, with a minimum usage period of 1 minute.

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

Supported Database Edition and Versions

Learn about the versions of Oracle Database that Oracle Exadata Database Service on Cloud@Customer supports.

Oracle Exadata Database Service on Cloud@Customer supports the following Oracle Database software releases:

- Oracle AI Database 26ai
- Oracle Database 19c (19.x)
- Oracle Database 12c Release 2 (12.2.0.1) (requires a valid Upgrade Support contract).
- Oracle Database 12c Release 1 (12.1.0.2) (requires a valid Upgrade Support contract).
- Oracle Database 11g Release 2 (11.2.0.4) (requires a valid Upgrade Support contract).

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)

System and Shape Configuration Options

Review the list of Exadata System Shapes

Copyright © 2008, 2026, Oracle and/or its affiliates.

- System Configuration Options for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.
- Oracle Exadata X11M System Model Specifications
  Review the technical specifications of X11M system model.
- Oracle Exadata X10M System Model Specifications
  Review the technical specifications of X10M system model.
- Oracle Exadata X9M-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.
- Oracle Exadata X8M-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.
- Oracle Exadata X8-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.
- Oracle Exadata X7-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.

System Configuration Options for Oracle Exadata Database Service on Cloud@Customer

Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.

Oracle Exadata X11M offers the following **database server** shapes:

- Base: 120 ECPUs and 660 GB of memory
- X11M (standard): 760 ECPUs and 1390 GB of memory
- X11M-L (large): 760 ECPUs and 2090 GB of memory
- X11M-XL (extra large): 760 ECPUs and 2800 GB of memory

Oracle Exadata X11M offers the following **storage server** shapes:

- Base: 35.6 TB usable storage capacity
- X11M-HC: 80 TB usable storage capacity
- X11M-EF: 37.2 TB usable storage capacity

Start with 2 database and 3 storage servers, and independently add up to a total of 16 servers in the initial rack and up to a total of 32 database and 64 storage servers across multiple racks. Any database server (Base, X11M, X11M-L, or X11M-XL) can be combined with Base, High Capacity (HC), or Extreme Flash (EF) storage servers. All database and storage servers in a system must be the same type.

Review the technical specifications of available Exadata System Shapes.

Oracle Exadata X11M System Model Specifications

Review the technical specifications of X11M system model.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 1-1 Oracle Exadata X11M System Model Specifications

Property	Base System Example	Quarter Rack Equivalent Example (min config)
Number of DB servers per System	2x Base	2x X11M, X11M-L, or X11M-XL
Number of Storage Servers per System	3x Base	3x High Capacity or 3x Extreme Flash
Total Usable ECPUs per System	240	1520
Total Memory Available for VMs (GB)	1320	• X11M: 2780 • X11M-L: 4180 • X11M-XL: 5600
Max Usable Local Storage Per DB Server (GB)	1084	2243
Max Usable File System Size Per VM (GB)	900	900
VM Image Size Minimum and Default (GB)	244	244
Max Number of VM Clusters per System	12	12
Maximum number of VMs per DB server	4	8
Total Flash Cache Capacity (TB)	38.4	81.6
Total Usable Disk Storage Capacity (TB)	106	240 High Capacity or 111 Extreme Flash

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

A maximum of 12 VM Clusters can be created on a single X11M system with 2 database servers and a maximum of 24 VM Clusters can be created on a system that contains greater than 2 DB servers. The maximum number of VM Clusters per system is the same regardless of whether the system contains base or standard database servers.

For more information, see Estimating How Much Local Storage You Can Provision to Your VMs and Scaling Local Storage.

Oracle Exadata X10M System Model Specifications

Review the technical specifications of X10M system model.

Oracle Exadata Cloud@Customer X10M is offered in the following Exadata X10M System Shapes:

- X10M (standard): Starts with two standard memory (1390 GB) database servers and three Oracle Exadata storage servers.
- X10M-L (large): Starts with two large memory (2090 GB) database servers and three Oracle Exadata storage servers.
- X10M-XL (extra large): Starts with two extra large memory (2800 GB) database servers and three Oracle Exadata storage servers.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

- A single DB server contains 190 usable cores with a choice of X10M (1390 GB), X10M-L (2090 GB), or X10M-XL (2800 GB) memory.
- A single storage server contains 80 TB of usable disk storage capacity.

Base System contains 2 base DB servers and 3 base storage servers. It can be expanded in a single rack up to a total of 16 servers.

- A single base DB server contains 30 usable cores and 660 GB memory.
- A single base storage server contains 35.6 TB of usable disk storage capacity.

Table 1-2 Oracle Exadata X10M System Model Specifications

Property	Base System	Quarter Rack X10M (min config)
Number of DB servers per System	2	2
Number of Storage Servers per System	3	3
Total Usable Cores in DB Servers per System	60	380
Total Memory Available for VMs (GB)	1320	• X10M: 2780 • X10M-L: 4180 • X10M-XL: 5600
Max Usable Local Storage Per DB Server (GB)	1084	2243
Max Usable File System Size Per VM (GB)	900	900
VM Image Size Minimum and Default (GB)	244	244
Max Number of VM Clusters per System	12	12
Maximum number of VMs per DB server	4	8
Total Flash Capacity (TB)	38.4	81.6
Total Usable Disk Storage Capacity (TB)	106	240

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

A maximum of 12 VM Clusters can be created on a single X10M system with 2 database servers and a maximum of 24 VM Clusters can be created on a system that contains greater than 2 DB servers. The maximum number of VM Clusters for systems containing Base database servers is 12 regardless of the number of Base database servers in the system.

For more information, see Estimating How Much Local Storage You Can Provision On Your VMs and Scaling Local Storage.

Oracle Exadata X9M-2 System Model Specifications

Review the technical specifications of available Exadata System Shapes.

Oracle Exadata X9M-2 is available in Quarter, Half, and Full Rack shapes. Independently add compute and/or storage servers to these shapes up to a total of 32 DB servers and 64 storage servers.

- A single DB server contains 62 usable cores and 1390 GB memory.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- A single storage server contains 63.6 TB of usable disk storage capacity.

Base System contains 2 base DB servers and 3 base storage servers. It can be expanded in a single rack up to 8 base DB servers and 12 base storage servers.

- A single base DB server contains 24 usable cores and 328 GB memory.
- A single base storage server contains 24.8 TB of usable disk storage capacity.

Table 1-3 Oracle Exadata X9M-2 System Model Specifications

Property	Base System	Quarter Rack	Half Rack	Full Rack
Number of DB servers per System	2	2	4	8
Number of Storage Servers per System	3	3	6	12
Total Usable Cores in DB Servers per System	48	124	248	496
Total Memory Available for VMs (GB)	656	2780	5560	11120
Max Usable Local Storage Per DB Server (GB)	1077	2243	2243	2243
Max Usable File System Size Per VM (GB)	892	900	900	900
VM Image Size Minimum and Default (GB)	244	244	244	244
Max Number of VM Clusters per System	12	12	24	24
Max number of VMs per DB server	4	8	8	8
Total Flash Capacity (TB)	38.4	76.8	153.6	307.2
Total Usable Disk Storage Capacity (TB)	73	191	383	766

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

A maximum of 12 VM Clusters can be created on a single X9M system with 2 database servers and a maximum of 24 VM Clusters can be created on a system that contains greater than 2 DB servers. The maximum number of VM Clusters for systems containing Base database servers is 12 regardless of the number of Base database servers in the system.

For more information, see Estimating How Much Local Storage You Can Provision On Your VMs and Scaling Local Storage.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Oracle Exadata X8M-2 System Model Specifications

Review the technical specifications of available Exadata System Shapes.

Oracle Exadata X8M-2 is available in Quarter, Half, Full Rack shapes. They can be expanded in a single rack up to a total of 8 DB servers and 12 storage servers.

- A single DB server contains 50 usable cores and 1390 GB of memory.
- A single storage server contains 49.6 TB of usable disk storage capacity.

Base System contains 2 base DB servers and 3 base storage servers. It can be expanded in a single rack up to a total of 8 base DB servers and 12 base storage servers.

- A single base DB server contains 24 usable cores and 328 GB of memory.
- A single base storage server contains 24.6 TB of usable disk storage capacity.

Table 1-4 Oracle Exadata X8M-2 System Model Specifications

Property	Base System	Quarter Rack	Half Rack	Full Rack
Number of DB servers per System	2	2	4	8
Number of Storage Servers per System	3	3	6	12
Total Usable Cores in DB Servers per System	48	100	200	400
Total Memory Available for VMs (GB)	656	2780	5560	11120
Max Usable Local Storage Per DB Server (GB)	1237	2500	2500	2500
Max Usable File System Size Per VM (GB)	900	900	900	900
VM Image Size Minimum and Default (GB)	244	244	244	244
Max Number of VM Clusters per System	8	8	16	16
Max number of VMs per DB server	4	8	8	8
Total Flash Capacity (TB)	38.4	76.8	153.6	307.2
Total Usable Disk Storage Capacity (TB)	73	148	297	595

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

A maximum of 8 VM Clusters can be created on a single X8M system with 2 database servers and a maximum of 16 VM Clusters can be created on a system that contains greater than 2 DB servers. The maximum number of VM Clusters for systems containing Base database servers is 8 regardless of the number of Base database servers in the system.

Copyright © 2008, 2026, Oracle and/or its affiliates.

For more information, see Estimating How Much Local Storage You Can Provision On Your VMs and Scaling Local Storage.

Oracle Exadata X8-2 System Model Specifications

Review the technical specifications of available Exadata System Shapes.

Table 1-5 Oracle Exadata X8-2 System Model Specifications

Property	Base System	Quarter Rack	Half Rack	Full Rack
Number of DB servers per System	2	2	4	8
Number of Storage Servers per System	3	3	6	12
Total Usable Cores in DB Servers per System	48	100	200	400
Total Memory Available for VMs (GB)	720	1440	2880	5760
Max Usable Local Storage Per DB Server (GB)	1237	1037	1037	1037
Max Usable File System Size Per VM (GB)	1100	900	900	900
VM Image Size Minimum and Default (GB)	244	244	244	244
Total Number of VM Clusters per System	5	5	10	10
Maximum Number of VMs per DB server	5	5	5	5
Total Flash Capacity (TB)	38.4	76.8	153.6	307.2
Total Usable Disk Storage Capacity (TB)	73	148	297	595

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

For more information, see Estimating How Much Local Storage You Can Provision On Your VMs and Scaling Local Storage.

Oracle Exadata X7-2 System Model Specifications

Review the technical specifications of available Exadata System Shapes.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 1-6 Oracle Exadata X7-2 System Model Specifications

Property	Base System	Quarter Rack	Half Rack	Full Rack
Number of DB servers per System	2	2	4	8
Number of Storage Servers per System	3	3	6	12
Total Usable Cores in DB Servers per System	44	92	184	368
Total Memory Available for VMs (GB)	480	1440	2880	5760
Max Usable Local Storage Per DB Server (GB)	1237	1237	1237	1237
Max Usable File System Size Per VM (GB)	1100	1100	1100	1100
VM Image Size Minimum and Default (GB)	244	244	244	244
Total Number of VM Clusters per System	6	6	12	12
Maximum Number of VMs per DB server	6	6	6	6
Total Flash Capacity (TB)	19.2	76.8	153.6	307.2
Total Usable Disk Storage Capacity (TB)	42	105	211	423

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

For more information, see Estimating How Much Local Storage You Can Provision On Your VMs and Scaling Local Storage.

Support Requirements

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

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

the Oracle Technical Support Policies, including, but not limited to, the Matching Service Level requirement.

Exadata Cloud Management Interfaces

Oracle Exadata Database Service on Cloud@Customer 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
- Local VM Command-Line Interfaces

Introduction to Exadata Cloud Management Interfaces

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

The various interfaces include:

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

The management interfaces are grouped into two primary categories:

- OCI Control Plane Interfaces
- Local Exadata Cloud VM CLIs

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

Related Topics

- Using the Console
- https://support.oracle.com/epmos/faces/DocContentDisplay?id=2768569.1

OCI Control Plane Interfaces

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 apply to Exadata Cloud resources and its infrastructure. Each of these interfaces can be used to accomplish the same functionality, all calling the OCI APIs, and are provided to enable flexibility and choice depending on preference and use case.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Command Line Interface (CLI):** The OCI CLI is a small footprint tool that you can use on its own or with the Console to complete Exadata Cloud resource 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 the developing 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.
- **Appendix and Reference:** This general reference shows how to configure the SDKs and other developer tools to integrate with Oracle Cloud Infrastructure services.
- **REST APIs:** This complete reference provides details on the Oracle Cloud Infrastructure REST APIs, including descriptions, syntax, endpoints, errors, and signatures. Exadata Cloud@Customer specific OCI REST APIs can be found throughout the documentation in the *Using the API* sections:
  - *Using the API to Manage Exadata Cloud@Customer Infrastructure*
  - *Using the API to Manage Exadata Cloud@Customer Backup Destinations*
  - *Using the API to Manage Exadata Cloud@Customer VM Clusters*
  - *Using the API to Create Oracle Database Home on Exadata Cloud@Customer*
  - *Using the API to Manage Oracle Database Home on Exadata Cloud@Customer*
  - *Using the API to Manage Oracle Database Components*
  - *Using the API to Manage Data Guard Associations on an Exadata Cloud@Customer System*
  - *Using the API to Manage Database Backup and Recovery*
  - *Using the API to Patch an Exadata Cloud@Customer System*

Related Topics

- Command Line Interface (CLI)
- Software Development Kits
- DevOps Tools and Plug-ins
- Terraform Provider
- Resource Manager
- Ansible Collection
- Cloud Shell
- Appendix and Reference
- REST APIs
- Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure
  Oracle Exadata Database Service on Cloud@Customer uses the same API as Oracle Cloud Infrastructure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Backup Destinations
  Review the list of API calls to manage your Oracle Exadata Database Service on Cloud@Customer backup destinations.
- Using the API to Manage Oracle Exadata Database Service on Cloud@Customer VM Clusters
  Review the list of API calls to manage your Oracle Exadata Database Service on Cloud@Customer VM cluster networks and VM clusters.
- Using the API to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home, review the list of API calls.
- Using the API to Manage Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  Review the list of API calls to manage Oracle Database home.
- Using the API to Manage Oracle Database Components
  Use various API features to help manage your databases on Oracle Exadata Database Service on Cloud@Customer.
- Using the API to Manage Data Guard Associations on an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to use the API to manage Data Guard associations on an Oracle Exadata Database Service on Cloud@Customer system.
- Using the API to Manage Database Backup and Recovery
  Learn how to use the API to manage database backup and recovery with Oracle Exadata Database Service on Cloud@Customer.
- Using the API for Patching and Updating VM Cluster and Database Homes
  Use various API features to help manage patching an Oracle Exadata Database Service on Cloud@Customer system.

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 when a corresponding OCI API is not available or the Exadata Cloud@Customer is in a disconnected mode.

The utilities include:

- **dbaascli:** Use the `dbaascli` utility to perform various database lifecycle and administration operations on the Exadata Cloud Service such as
  - changing the password of a database user
  - starting a database
  - managing pluggable databases (PDBs)
  - scaling the CPU core count in disconnected mode
- **bkup_api:** Use the `bkup_api` utility to perform various backup and recovery operations on the Exadata Cloud Service such as creating an on-demand backup of a complete database or an individual pluggable database (PDB), or to *customize backup settings* used by the automatic backup configuration

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

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

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

Related Topics

- About Using the dbaascli Utility on Oracle Exadata Database Service on Cloud@Customer
  You can use the `dbaascli` utility to perform various database lifecycle and administration operations on Oracle Exadata Database Service on Cloud@Customer such as creating an Oracle Database, patching an Oracle Database, managing pluggable databases (PDBs), scaling the CPU core count in disconnected mode, and more.
- Creating an On-Demand Backup by Using dbaascli
  You can use the `dbaascli` to create an on-demand backup of a complete database or an individual pluggable database (PDB):
- Customizing Backup Settings by Using a Generated Configuration File
  You can customize backup settings for a database deployment by generating a file containing the current customizable settings, editing the file, and then using the file to update the backup settings.
- Monitoring and Managing Exadata Storage Servers with ExaCLI
  Learn to use the ExaCLI command-line utility to perform monitoring and management functions on Exadata storage servers in the Exadata Cloud Service.
- Updating Cloud Tooling Using dbaascli
  To update the cloud tooling release for Oracle Exadata Database Service on Cloud@Customer, complete this procedure.

Exadata Cloud@Customer Technical Architecture

Explore the Exadata Cloud@Customer architecture, featuring virtual machine clusters, control plane servers, and seamless connectivity between client applications, OCI services, and Oracle-managed admin networks. The interactive diagrams further illustrate the role of automation tunnels and the clear separation of responsibilities in managing both databases and infrastructure. For detailed insights, refer to Exadata Cloud@Customer Technical Architecture.

Copyright © 2008, 2026, Oracle and/or its affiliates.

2

What’s New in Oracle Exadata Database Service on Cloud@Customer

Oracle is constantly adding new capabilities to Oracle Exadata Database Service on Cloud@Customer. This section provides a brief overview of new features as they are released.

- Monthly Infrastructure Maintenance Update Preferences
- UX Updates for Autonomous AI Database on Cloud@Customer
- Custom ASM Disk Group Sizing
- Exadata Storage Server X11M Extreme Flash (EF)
- Customize SCAN Listener Port (TCPS)
- Support for a 100Gbps NIC for the Exadata Cloud@Customer (ExaDB-C@C) Client Network
- Exascale VM storage for Exadata Database Service on Cloud@Customer
- Pluggable Database Snapshots for Exadata Database Service on Cloud@Customer
- Display ASM Rebalance Progress During Storage Scaling
- Standalone Backup
- Real-Time Data Protection on Oracle Zero Data Loss Recovery Appliance (ZDLRA)
- Database Update Enhancement: Skip Datapatch and Run Independently Later
- Enhanced Controls to Configure Automatic Backups
- Enhanced Data Guard Health Insights
- Support for a 100Gbps NIC for the Exadata Cloud@Customer (ExaDB-C@C) Backup Network
- Convert a Physical Standby Database to a Snapshot Standby Database and vice versa
- Exadata System Software 25.1.0.0.0
- Select Grid Infrastructure Minor Version (N-3) While Provisioning a VM Cluster
- Expanded Tagging Support Across Database Workflows
- Support for Concurrent Data Guard Operations in Multiple Standby Environments
- Schedule VM Cloud Automation Updates
- Classify and Control Access to Critical OCI REST APIs
- Oracle Data Guard Setup with Precheck Validation
- Backup and Restore from a Standby Database in a Data Guard Environment
- Concurrently Create or Delete Pluggable Databases (PDBs)
- Exadata Database Service for Developers
- Multiple Standby Databases
- Key Management Using an External Keystore

Copyright © 2008, 2026, Oracle and/or its affiliates.

- X11M System Support
- Exadata Exascale with Exadata Database Service on Cloud@Customer
- Microsoft Entra ID (MS-EI) Integration with Oracle Exadata Database Service on Cloud@Customer
- Support for Concurrent Data Guard, Container Database (CDB), and Pluggable Database (PDB) Operations
- Delegate Access Control for ExaDB-C@C
- Cost and Usage Attribution for Pluggable Databases (PDBs)
- Different RUs for Primary and Standby DB Homes in Data Guard Associations, Switchover, and Failover Operations
- Granular Permissions for VM Cluster Update Operations
- Enhancements to Quarterly Exadata Infrastructure Maintenance Planning and Execution
- Enable Unified Auditing While Creating a Database Home
- Enhancements to Oracle Key Vault (OKV) Integration with ExaDB-C@C to Manage Transparent Data Encryption (TDE) Keys
- Ability to Increase the Size of Guest VM Local File Systems
- Create and Use Custom Software Images
- Oracle AI Database 26ai on Oracle Exadata Database Service on Cloud@Customer
- Change the Home Region of ExaDB-C@C Infrastructure
- Enhancements to Serial Console Functionality
- Provision a VM Cluster with Either an OL7 or OL8-Based Image
- VM Cluster on a Single VM
- Enhancement to Pluggable Database (PDB) Management
- Manage Administrator (SYS User) and TDE Wallet Passwords
- Update Guest VM Operating System to Oracle Linux 8
- Exadata Fleet Update
- Oracle Key Vault (OKV) Integration with ExaDB-C@C to Manage Transparent Data Encryption (TDE) Keys
- X10M System Support
- Manage Serial Console Access to ExaDB-C@C Systems
- Interim Software Updates
- Link Aggregation Control Protocol (LACP) Support for Client and Backup Networks
- VM Cluster and Database Health and Performance Metrics in the OCI Console
- Enable Data Guard with Database Software Image (Custom Image)
- Use Identity and Access Management (IAM) Authentication with Oracle Exadata Database Service on Cloud@Customer
- Create Data Guard Associations Across Regions
- Multi-Rack Elastic Compute and Storage Expansion
- Automatic Diagnostic Collection
- Elastic Compute Expansion

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Oracle Standard Tagging for Resources on Oracle Exadata Database Service on Cloud@Customer
- Exadata Infrastructure Maintenance History
- Monthly ExaDB-C@C Infrastructure Security Maintenance
- Enhanced Control to Rollback or Retry Failed Guest VM Operating System Update
- Concurrently Create or Terminate Oracle Databases in a VM Cluster
- Support for Rack Serial Number as a System Tag
- Support for DB Home Minor Version Selection (N-3)
- VM Guest Exadata OS Image Major Version Update
- Database Service Events
- Control Plane Server (CPS) Offline Diagnostic Report
- Enhanced Infrastructure Maintenance Controls
- Manage Pluggable Databases on Exadata Cloud@Customer
- Allow Customers to Choose Data Guard Type
- Specify the Same SID for Primary and Standby Databases in a Data Guard Association
- Specify db_unique_name and SID for Primary and Standby Databases in Data Guard Association
- VM Cluster Node Subsetting
- X9M-2 System Support
- Customize SCAN Listener Port
- Creating DG Association/Standby Database Using Existing Database Home
- Upgrading Oracle Grid Infrastructure on an Exadata Cloud@Customer VM Cluster
- Updating Guest VM Operating System
- Upgrading Oracle Databases
- Download Network Validation Report
- Elastic Storage Expansion
- Oracle Database Software Images
- Exadata Cloud@Customer Infrastructure Patching Automation
- Customer Maintenance Contacts
- X8M-2 System Support
- Enable and Manage Data Guard Associations
- Oracle Exadata Cloud@Customer Deployment Assistant
- Oracle Grid Infrastructure and Oracle Database Patching
- Per-Second Billing for OCPU Usage
- Shared Database Home Resource Tags
- Create and Manage Multiple Virtual Machines per Exadata System (MultiVM)
- Scale OCPUs Without Cloud Connectivity
- Configure Oracle Database Character Set and National Character Set

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Specify a Time Zone While Provisioning Oracle Exadata Database Service on Cloud@Customer Infrastructure
- Shared Database Homes for Oracle Exadata Database Service on Cloud@Customer
- X7-2 System Support

Monthly Infrastructure Maintenance Update Preferences

- **Services**: Database
- **Release Date**: May 27, 2026

> **Note**
>
> This feature has been rolled out only to the following regions:
>
> - `uk-cardiff-1`
> - `eu-zurich-1`
> - `us-dallas-1`
> - `ap-batam-1`
> - `ap-delhi-1`
>
> It will be rolled out to the remaining regions in phases.

Customers can now specify monthly infrastructure maintenance update preferences, giving greater control over when maintenance is performed. Based on customer-defined time slots that align with business needs, Oracle automation performs maintenance on designated infrastructure components and ensures that all components receive required software updates to meet compliance guidelines.

With monthly maintenance update preferences for Exadata infrastructure components, customers can:

- Perform security maintenance in smaller, more manageable monthly time windows.
- Group infrastructure components and streamline maintenance across database and storage servers to align with business priorities (for example, development first, followed by test and production).
- Abort an in-progress maintenance run to accommodate unplanned operational requirements.

Related Topics

- Maintenance Scheduling Policy
  Learn how to use the OCI Console to configure and manage maintenance scheduling policies. If a maintenance scheduling policy is used, all the scheduling preferences for the infrastructure maintenance are derived from the policy.
- Using the Console to Configure Oracle-Managed Infrastructure Updates
  Full Exadata infrastructure software updates are scheduled on a quarterly basis. In addition, important security updates are scheduled monthly. While you cannot opt-out of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal and allows scheduling flexibility to help you plan for them.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Manage Quarterly/Monthly Maintenance Run created from Scheduling Plan
- Exadata Infrastructure Maintenance Event Types
  Review the list of event types that Exadata Infrastructure Maintenance emits.

UX Updates for Autonomous AI Database on Cloud@Customer

- **Services**: Database
- **Release Date**: May 20, 2026

The landing page experience for Autonomous AI Database on Cloud@Customer has been enhanced to improve visibility and simplify resource management across Autonomous AI Database and Exadata Database Services.

With this update, Autonomous AI Database on Cloud@Customer resources now have a dedicated landing page, separate from the Exadata Database Service on Cloud@Customer landing page.

From the Autonomous AI Database landing page, the following resources redirect to their corresponding pages under Exadata Database Service on Cloud@Customer:

- Exadata VM Clusters
- Exadata Infrastructure
- Backup Destinations
- Key Stores

These UX enhancements provide clearer navigation and improve the overall management experience for Autonomous AI Database environments.

Related Topics

- Introduction to Oracle Autonomous AI Database on Exadata Cloud@Customer

Custom ASM Disk Group Sizing

- **Services**: Database
- **Release Date**: May 15, 2026

The ASM Disk Group Resizing feature is now generally available for Oracle Exadata Database Service on Cloud@Customer (ExaDB-C@C).

This feature enables you to define custom space allocations for your ASM disk groups.

Previously, Exadata Cloud supported only fixed default percentage allocations for the DATA, RECO, and SPARSE disk groups. This limitation may have prevented you from fully utilizing available storage as your workload requirements evolved.

With this capability, you can specify custom space allocations for DATA, RECO, and SPARSE disk groups during VM Cluster provisioning and modify those allocations after provisioning. This flexibility is subject to certain system restrictions described in the Planning Essentials section below.

This capability complements the existing Oracle-recommended default allocations, allowing you to either use the simplified defaults or tailor disk group allocations to meet your operational requirements.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Key Benefits

- Optimal Storage Utilization
  You can maximize the use of your available storage by starting with Oracle default allocations or defining custom allocations during VM Cluster creation. You can also modify disk group sizes over time as your workload requirements change. For example, you can reallocate unused space in the RECO disk group to the DATA disk group to improve storage utilization and potentially reduce costs.
- Increased Control and Flexibility
  You gain greater control over how storage is allocated across disk groups. If you are accustomed to the flexibility available in on-premises Exadata environments, this feature enables you to apply similar allocation strategies in your cloud deployments. This allows you to better align storage configuration with your workload and operational requirements.

Related Topics

- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Scaling Up or Scaling Down the VM Cluster Resources
  You can scale up or scale down the memory, local disk size (`/u02`), ASM Storage, and CPUs (ECPUs for X11M).
- Calculating the ASM Storage

Exadata Storage Server X11M Extreme Flash (EF)

- **Services**: Database
- **Release Date**: April 01, 2026

This release adds support for Exadata Storage Server X11M Extreme Flash (EF), delivering deterministic low-latency performance for large random I/O and accelerating write-intensive operations through faster persistence to capacity-optimized flash. This configuration is ideal for modern, latency-critical workloads that demand consistent all-flash behavior.

Related Topics

- Oracle Exadata X11M System Model Specifications
  Review the technical specifications of X11M system model.
- System Configuration Options for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.

Customize SCAN Listener Port (TCPS)

- **Services**: Database
- **Release Date**: March 03, 2026

You can now specify a SCAN listener port (TCPS) within the permissible range while creating a VM cluster network resource. For more information, see:

- *Using the Console to Create a VM Cluster Network*
- *Using the Console to View SCAN Listener Port Configured*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Using the Console to Create a VM Cluster Network
  To create your VM cluster network with the Console, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to View SCAN Listener Port Configured
  You can only edit a VM cluster network that is not associated with a VM cluster.

Support for a 100Gbps NIC for the Exadata Cloud@Customer (ExaDB-C@C) Client Network

- **Services**: Database
- **Release Date**: February 25, 2026

For new Exadata Cloud@Customer subscriptions, customers can choose to upgrade the client NIC from 25Gbps to 100Gbps. Oracle will ship a dual 100Gbps QSFP28 network card and the necessary transceivers to connect to a 100Gbps network. The 100Gbps client network is available with both Exadata Database Service and Autonomous Database running on ExaC@C.

Customers will benefit from higher network speed for large data transfers between applications and database.

The 100Gbps NIC is available for both client and backup networks. 100Gbps support for the backup network was previously released.

Availability

The 100Gbps client NIC is available immediately in all commercial regions.

Exascale VM storage for Exadata Database Service on Cloud@Customer

- **Services**: Database
- **Release Date**: February 17, 2026

> **Note**
>
> This feature has been rolled out to all regions except: `UAE Central (Abu Dhabi)` **`AUH`**, `UAE East (Dubai)` **`DXB`**, `Saudi Arabia West (Jeddah)` **`JED`**, `Saudi Arabia Central (Riyadh)` **`RUH`**, `South Korea North (Chuncheon)` **`YNY`**, `India North (Delhi)` **`ONM`**, `Malaysia West 2 (Kulai)` **`JBP`**, `Australia Southeast (Melbourne)` **`MEL`**, `India West (Mumbai)` **`BOM`**, `Japan Central (Osaka)` **`KIX`**, `Singapore West (Singapore)` **`XSP`**, `Canada Southeast (Toronto)` **`YYZ`**, `Ireland East (Dublin)` **`ZQO`**, `Spain Central (Madrid)` **`MAD`**, `Italy Northwest (Milan)` **`LIN`**, `France Central (Paris)` **`CDG`**, `Italy North (Turin)` **`NRQ`**, `Switzerland North (Zurich)` **`ZRH`**, `Mexico Central (Queretaro)` **`QRO`**, `Brazil East (Sao Paulo)` **`GRU`**, and `Chile West (Valparaiso)` **`VAP`**.
>
> It will be rolled out to the remaining regions in phases.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Oracle is excited to announce the General Availability (GA) of Exascale VM Storage for Exadata Database Service on Cloud@Customer (ExaDB-C@C). You can now use Exascale storage on Exadata Infrastructure not just for Oracle Database 26ai, but also for Virtual Machines - to store VM images and Oracle-managed VM backups with greater flexibility. This capability extends to Virtual Machines on ASM VM clusters too, letting you use Exascale storage for Virtual Machines while still provisioning Oracle Database 19c and 26ai with ASM. Additionally, existing VM clusters with Oracle-managed backups stored locally will also receive a one-time opportunity to migrate future VM backups to Exascale storage.

Key Benefits

- Expanded VM storage: Exascale offers greater storage capacity for provisioning and expanding VM and general-purpose file systems, eliminating local disk storage constraints.
- More VM backups: Exascale supports two full Oracle-managed VM backups, compared to only one with local storage. Future releases will introduce new features like on-demand VM Backups and VM snapshots for even more flexibility.
- ACFS Support: Create and manage Advanced Cluster File System (ACFS) on Exascale VM cluster.

Related Topics

- Using the Console to Configure Exascale Storage on Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To configure Exascale storage on your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the infrastructure configuration.
- Using the Console to Configure Exascale VM Storage on Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To configure Exascale VM storage on your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the infrastructure configuration.
- Configure All Future VM Backups to Use Exascale Storage on a VM Cluster
  If a VM Cluster (ASM or Exascale) is using local database server storage for Oracle-managed VM backups, This option allows you to migrate all future VM backups to Exascale storage.
- Create an Advanced Cluster File System (ACFS) Using Exascale Block Volume
  Before you begin, ensure that the VM cluster is configured to use Exascale for database storage and that sufficient storage capacity is available in the Exascale storage vault associated with the cluster.
- Deconfigure Exascale Storage on the Exascale Infrastructure
  Before you begin, ensure that the VM cluster is configured to use Exascale for database storage and that sufficient storage capacity is available in the Exascale storage vault associated with the cluster.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create an Exascale VM Cluster
  To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Pluggable Database Snapshots for Exadata Database Service on Cloud@Customer

- **Services**: Database
- **Release Date**: February 17, 2026

> **Note**
>
> This feature has been rolled out to all regions except: UAE Central (Abu Dhabi) **AUH**, UAE East (Dubai) **DXB**, Saudi Arabia West (Jeddah) **JED**, Saudi Arabia Central (Riyadh) **RUH**, South Korea North (Chuncheon) **YNY**, India North (Delhi) **ONM**, Malaysia West 2 (Kulai) **JBP**, Australia Southeast (Melbourne) **MEL**, India West (Mumbai) **BOM**, Japan Central (Osaka) **KIX**, Singapore West (Singapore) **XSP**, Canada Southeast (Toronto) **YYZ**, Ireland East (Dublin) **ZQO**, Spain Central (Madrid) **MAD**, Italy Northwest (Milan) **LIN**, France Central (Paris) **CDG**, Italy North (Turin) **NRQ**, Switzerland North (Zurich) **ZRH**, Mexico Central (Queretaro) **QRO**, Brazil East (Sao Paulo) **GRU**, and Chile West (Valparaiso) **VAP**.
>
> It will be rolled out to the remaining regions in phases.

Oracle is excited to announce the General Availability (GA) of Pluggable Database (PDB) Snapshots for Exadata Database Service on Cloud@Customer (ExaDB-C@C).

> **Note**
>
> The feature requires Oracle Database version 23.9.0.25.07 (and later).

With Oracle Database 26ai using Exascale storage on ExaDB-C@C, you can now create point-in-time PDB snapshots and reuse them as needed. These snapshots enable you to quickly create thin, space-efficient clones that deliver full Exadata performance, making them ideal for development and testing environments.

PDB Snapshots simplify data refresh operations and improve developer productivity by making database cloning easy, scalable, and efficient.

Key Benefits

- Agility for Dev/Test: Rapidly create and refresh development and test environments with minimal overhead.
- Efficiency: Snapshots are fast to create and require very little initial storage.
- Scale: Powered by Exascale redirect-on-write technology, you can create virtually unlimited snapshots.
- Parallel Development: Use multiple snapshots and clones to support independent development or testing branches.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Create an Exascale PDB Snapshot
- Terminate an Exascale PDB Snapshot
- Create a Thick or Thin Local or Remote Clone from a PDB Snapshot

Display ASM Rebalance Progress During Storage Scaling

- **Services**: Database
- **Release Date**: February 17, 2026

This release enhances visibility into ASM rebalance progress during storage scaling operations for Exadata VM Clusters, and improves operational flexibility while scaling is in progress.

Key enhancements include:

- Scale VM Cluster operations can now be initiated to scale CPU resources even while a storage attachment is in progress and the VM Cluster is in the UPDATING state.
- VM Clusters where the storage attachment has completed and the ASM rebalance is running can now transition back to AVAILABLE, even if other clusters are still in progress. This enables customers to submit new operations without waiting for all clusters to complete.
- Exadata VM Clusters automatically transition from UPDATING to AVAILABLE as soon as the storage attachment workflow completes.
- Enhanced visibility into storage attachment and ASM rebalance progress, including:
  - Identification of VM Clusters that have completed
  - Real-time rebalance status, including progress percentage and messages, displayed in the Log Messages section of the Work Request Details page

Standalone Backup

- **Services**: Database
- **Release Date**: February 17, 2026

ExaDB-C@C now supports standalone (ad-hoc) backups, in addition to automatic backups. Customers can create a backup on demand, available through OCI Console, OCI API, and Terraform.

Key benefits and use cases include:

- Taking a backup before a critical event (e.g., refresh, patching, or upgrade)
- Meeting business or compliance requirements

Related Topics

- Standalone Backup
- Using the Console to Create a Standalone Backup
  To take a standalone backup, complete this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Real-Time Data Protection on Oracle Zero Data Loss Recovery Appliance (ZDLRA)

- **Services**: Database
- **Release Date**: January 27, 2026

ExaDB-C@C now supports real-time data protection for database backups on the Recovery Appliance (RA) through the OCI Console and API. Previously, this capability was only available via dbaascli.

Real-time data protection continuously captures database changes, enabling near-zero data loss recovery even during outages or ransomware attacks. This significantly reduces downtime and ensures minimal disruption to business operations.

> **Note**
>
> Real-time data protection is supported only for Oracle Database server versions 19.18 and above.

Key Benefits

- Zero data loss through continuous capture of redo changes
- Reduced downtime with faster recovery operations
- Enhanced security against data corruption or ransomware
- Cost savings by optimizing recovery operations and reducing business impact

This enhancement improves the overall user experience by making real-time data protection accessible directly from the UI and API.

For more information about RTRT, see Real-Time Redo Transport.

Related Topics

- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.

Database Update Enhancement: Skip Datapatch and Run Independently Later

- **Services**: Database
- **Release Date**: January 21, 2026

During a Database Release Update, the Database Home is first upgraded to the newer version, and the database is then restarted on the updated Database Home using a rolling restart. After all database instances are running from the updated Database Home, the datapatch phase is executed to apply the SQL updates and synchronize the database metadata with the newly installed changes.

Copyright © 2008, 2026, Oracle and/or its affiliates.

With this enhancement, customers now have the option to skip the datapatch step for all container databases or closed PDBs running from the Database Home being updated. Note that if datapatch is skipped, the database version will be lower than the Database Home version. To address this version mismatch, a new Run Datapatch action is provided for container and pluggable databases where datapatch was previously skipped. Because datapatch runs as an online operation, customers can initiate it at any time, even while the database is open and running.

The Skip datapatch option is also available during a Database Move operation, allowing you to skip the datapatch phase on the database or on all closed PDBs being moved from the source to the target Database Home.

The Skip datapatch and Run datapatch options are not available for standby databases. When datapatch is executed on the primary database, the standby database is automatically updated through the application of redo received from the primary.

Related Topics

- Using the Console to Perform Update Operation on a Database Home
  Learn to apply patches on a Database Home.
- Using the Console to Move a Database to Another Database Home
  Learn to move a database to another Database Home.
- To run datapatch on a container database (CDB)
  To run datapatch on a container database, use this procedure.
- To run datapatch on a pluggable database (PDB)
  To run datapatch on a pluggable database, use this procedure.

Enhanced Controls to Configure Automatic Backups

- **Services**: Database
- **Release Date**: January 14, 2026

Automatic Backups can be enabled either during the database creation process or added later after the database is set up.

With this enhancement, you have the flexibility to:

- Choose the specific day of the week for full backups to begin
- Schedule full and incremental backups to start independently, using either the default 6-hour window or a specified 2-hour window for more precise control
- Decide whether the first backup runs immediately or at a scheduled time
- Modify backup schedules for future backups as needed
- Select your preferred backup retention policy when terminating a database

Related Topics

- About Managing Database Backup for Oracle Exadata Database Service on Cloud@Customer
  Learn how configure backup when creating the database on Oracle Exadata Database Service on Cloud@Customer.
- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Enhanced Data Guard Health Insights

- **Services**: Database
- **Release Date**: November 19, 2025

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

Related Topics

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

Support for a 100Gbps NIC for the Exadata Cloud@Customer (ExaDB-C@C) Backup Network

- **Services**: Database
- **Release Date**: October 17, 2025

For new Exadata Cloud@Customer subscriptions, customers can choose to upgrade the backup NIC speed from 25Gbps to 100Gbps. Oracle will ship a dual 100Gbps capable network card and the necessary transceivers to connect to a 100Gbps network. The 100Gbps backup network is available with both Exadata Database Service and Autonomous Database running on ExaDB-C@C.

Exadata Cloud@Customer backups are sent over the backup network to minimize the impact on client network traffic. All customers must take a full backup the first time they back up their database. Customers backing up to an NFS server, or object storage, must also take a full backup once per week. These full backups can be very large. Restoration of a backup over the network may also be time consuming, as the entire backup may need to be transferred over the network to the ExaDB-C@C. The faster backup network speeds the transfer of those large data files, shortening the overall backup and restore time. Benefits include:

- Shorter duration for full backups eliminates likelihood backup will extend into production hours and impact production workloads
- Faster restoration of a backup reduces downtime when a database restore is required
- Improved administrative productivity with less time spent waiting for data to transfer

Availability

The 100Gbps backup NIC is available immediately in all commercial regions.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Services**: Database
- **Release Date**: October 07, 2025

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

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

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

Related Topics

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

Exadata System Software 25.1.0.0.0

- **Services**: Database
- **Release Date**: December, 2024

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

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

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

To learn more, refer to the Exadata System Software Release 25.1 documentation.

For update instructions, see My Oracle Support Note 3021895.1 — Exadata System Software 25.1.0.0.0 Update.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

Select Grid Infrastructure Minor Version (N-3) While Provisioning a VM Cluster

- **Services**: Database
- **Release Date**: July 23, 2025

This enhancement eliminates the previous limitation of provisioning a VM cluster with the default Grid Infrastructure version provided by the system. You can now select the latest or any of the last three Grid Infrastructure releases (four options in total) during VM cluster creation. For example, when provisioning a 19c VM cluster, the available Grid Infrastructure versions include: 19.24.0.0.0 (latest), 19.23.0.0.0, 19.22.0.0.0, and 19.21.0.0.0.

Related Topics

- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create an Exascale VM Cluster
  To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Expanded Tagging Support Across Database Workflows

- **Services**: Database
- **Release Date**: July 22, 2025

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

Related Topics

- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.

Support for Concurrent Data Guard Operations in Multiple Standby Environments

- **Services**: Database
- **Release Date**: July 15, 2025

In addition to the Support for Concurrent Data Guard, Container Database (CDB), and Pluggable Database (PDB) Operations, these enhancements enable you to perform concurrent operations on CDBs and PDBs alongside Data Guard migration operations in environments with multiple standby databases.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Support for Concurrent Data Guard Operations in Multiple Standby Environments

Schedule VM Cloud Automation Updates

- **Services**: Database
- **Release Date**: July 11, 2025

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

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

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

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

Related Topics

- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create an Exascale VM Cluster
  To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- VM Cloud Automation Software Update Management

Classify and Control Access to Critical OCI REST APIs

- **Service**: Database
- **Release Date**: June 17, 2025

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

Related Topics

- Using the Console to Create API Access Control
  Use this procedure to create API Access Control for your ExaDB-C@C resources—Exadata Infrastructure, VM Cluster, Database, and Pluggable Database.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Oracle Data Guard Setup with Precheck Validation

- **Services**: Database
- **Release Date**: June 17, 2025

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

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

Related Topics

- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.

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

- **Services**: Database
- **Release Date**: May 20, 2025

> **Note**
>
> The ability to perform backups on more than one standby database using Zero Data Loss Recovery Appliance (ZDLRA) as the backup destination is on the roadmap for Oracle Exadata Database Service on Cloud@Customer.

This enhancement enables you to:

- Offload backups to the standby database using Oracle Cloud Infrastructure Object Storage, Zero Data Loss Recovery Appliance (ZDLRA), or Network File System (NFS) in a Data Guard environment, freeing up resources in the production database environment.
- Schedule automatic backups on the standby database in a Data Guard environment and configure the retention period.
- Restore and recover a standby database using a backup of the standby database.
- Choose to take backups only on the primary database, only on the standby database, or both, based on your needs.

Note:

- Restore/recover operations:
  - You cannot use backups on the standby database to perform restore/recover operations on the primary database.
- Switchover scenarios:
  - If automatic backups was configured on the primary with backup destination of Oracle Cloud Infrastructure Object Storage, Zero Data Loss Recovery Appliance (ZDLRA), or

Copyright © 2008, 2026, Oracle and/or its affiliates.

Network File System (NFS), upon switchover, the backups will continue on the new standby database.

  - If automatic backups was configured on the standby with backup destination of Oracle Cloud Infrastructure Object Storage, Zero Data Loss Recovery Appliance (ZDLRA), or Network File System (NFS), upon switchover, the backups will continue on the new primary database.
- Failover scenarios:
  - If automatic backups was configured on the primary with backup destination of Oracle Cloud Infrastructure Object Storage, Zero Data Loss Recovery Appliance (ZDLRA), or Network File System (NFS), upon failover the backups will be disabled on the new Disabled Standby database.
  - If automatic backups was configured on the standby with backup destination of Oracle Cloud Infrastructure Object Storage, Zero Data Loss Recovery Appliance (ZDLRA), or Network File System (NFS), upon failover the backups will continue on the new primary database.

Related Topics

- Automatic Backup and Recovery

Concurrently Create or Delete Pluggable Databases (PDBs)

- **Services**: Database
- **Release Date**: March 19, 2025

You can now create or delete up to 10 PDBs concurrently, even when the container database (CDB) is in an updating state. However, you cannot create or delete PDBs while a CDB is in an updating state if other operations—excluding PDB creation or deletion—are modifying its metadata or structure.

Related Topics

- Concurrently Create or Delete Pluggable Databases (PDBs)
  You can now create or delete up to 10 PDBs concurrently, even when the container database (CDB) is in an updating state. However, you cannot create or delete PDBs while a CDB is in an updating state if other operations—excluding PDB creation or deletion—are modifying its metadata or structure.
- Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel

Exadata Database Service for Developers

- **Services**: Database
- **Release Date**: March 19, 2025

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

You can allocate as many OCPU's/ECPU's, memory, and storage to the Developer VM Cluster as you require. Within the VM cluster you can have multiple CDB's, and each CDB can have Multiple PDB's.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

- **Single VM Cluster:** Exadata Database-Developer VM clusters are limited to a single VM.
- **Threads per PDB:** 2 threads (1 core) per PDB.
- **Memory per PDB:** 8 GB memory per PDB.
- **Database size per PDB:** 20 GB storage per PDB.
- **Sessions per PDB:** 30 sessions per PDB.
- **Data Guard:** Creating a Data Guard standby database is prohibited when using the Developer VM Clusters.
- **Cross-region backups:** Cross-region backups are not supported.
- **Versions:** Only Oracle Database versions 19.26 and 23.6 and later are supported

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

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

Related Topics

- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create an Exascale VM Cluster
  To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Multiple Standby Databases

- **Services**: Database
- **Release Date**: January 22, 2025

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

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

- **Local Standby:** A standby database in the same region as the production database is ideal for failover scenarios, offering zero data loss for local failures (such as database, cluster, or availability domain failures). Application failover impact is reduced in this case,

Copyright © 2008, 2026, Oracle and/or its affiliates.

as applications continue operating without the performance overhead of communicating with a remote region.

- **Remote (Cross-Region) Standby:** A remote standby database, located in a different region, is typically used for disaster recovery or to offload read-only query processing. A remote standby database setup ensures data protection against regional failures.

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

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

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

Related Topics

- Use Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer
  Learn to configure and manage Data Guard Groups in your VM cluster.
- Using the API to Manage Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to use the API to manage Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer system.
- Permissions and API operation details for Data Guard Group
- Data Guard Event Types
  Review the list of event types that Data Gurad groups and Data Guard Associations emit.

Key Management Using an External Keystore

- **Services**: Database
- **Release Date**: January 22, 2025

With this enhancement, you can now store and manage database encryption keys in an external keystore provider of your choice. This feature offers flexibility in choosing external keystore providers to enhance key security and control.

Related Topics

- Manage Keys Using an External Keystore
  Review the use cases and implementation details of an external keystore.
- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.
- Using the Console to Migrate Encryption Keys from Oracle Wallet to Other Encryption Methods (Oracle Key Vault, External Keystore)
  Learn to migrate encryption keys between different encryption methods.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Rotate Encryption Keys
  Learn how to rotate encryption keys. Note that key rotation is only available for databases that do not use Oracle Wallet as the encryption method.

X11M System Support

- **Services**: Database
- **Release Date**: January 07, 2025

> **Note**
>
> This feature is rolled out only to the `MTY`, `BOG`, `HYD`, `DXB`, `YUL`, `KIX`, `SYD`, `ZRH`, `AMS`, `MEL`, `SCL`, `YNY`, `VCP`, `CWL`, `SJC`, `MRS`, `ARN`, `MTZ`, `AUH`, `LIN`, `SIN`, `JNB`, `ORD`, `CDG`, `QRO`, `MAD`, `AGA`, `VAP`, `XSP`, `RUH`, `LHR`, `NRT`, `BOM`, `JED`, `ICN`, `PHX`, `IAD`, `FRA`, `YYZ`, and `GRU` regions. It will be rolled out to other regions in phases.

Oracle Exadata Cloud@Customer has been extended to support Exadata Infrastructure X11M.

Related Topics

- System Configuration Options for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.
- Oracle Exadata X11M System Model Specifications
  Review the technical specifications of X11M system model.
- Scaling Up or Scaling Down the VM Cluster Resources
  You can scale up or scale down the memory, local disk size (`/u02`), ASM Storage, and CPUs (ECPUs for X11M).
- Estimating How Much Local Storage You Can Provision On Your VMs
- Overview of Elastic Storage Expansion
  With elastic storage expansion, you can dynamically increase your storage capacity to meet your growing workload requirements.
- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create an Exascale VM Cluster
  To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Scale the Resources on a VM Cluster
  Starting in Oracle Exadata Database Service on Cloud@Customer, you can scale up or down multiple resources at the same time. You can also scale up or down resources one at a time.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Exadata Exascale with Exadata Database Service on Cloud@Customer

- **Services**: Database
- **Release Date**: December 17, 2024

Introducing Exadata Exascale with Exadata Database Service on Cloud@Customer (ExaDB-C@C), the world's only intelligent data architecture for the cloud is a revolutionary leap that combines the best of Exadata and the best of the cloud. Exascale reimagines how compute and storage resources are managed on Exadata platforms by decoupling and simplifying storage management, paving the way for innovative capabilities. It ensures industry-leading database performance, availability, and security standards that organizations expect from Exadata.

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

Minimum requirements for configuring Exascale on Exadata Cloud@Customer:

- This feature is only supported on Exadata Infrastructure Model X8M (and later) and requires Exadata system software release 24.1 and later.
- This feature requires Oracle Grid Infrastructure version 26ai (23.6) and supports Oracle AI Database versions 26ai (23.6) and later.

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

Exadata Exascale with Exadata Database Service on Cloud@Customer offers the following benefits:

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

For more information about Oracle Exadata Exascale, see the Oracle® Exadata Exascale User's Guide.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Storage Configuration Requirements for Oracle Exadata Database Service on Cloud@Customer
  With the introduction of Exascale technology in Exadata Database Service on Cloud@Customer, you can configure the Exadata infrastructure to use ASM, Exascale, or a combination of both. When configured to coexist with ASM, you must allocate Exascale storage capacity from the total Exadata storage available on the Infrastructure. The allocated storage capacity will be solely available for Exascale use.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create an Exascale VM Cluster
  To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.
- Clone a Pluggable Database (PDB)
  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.
- Resource-Types for Oracle Exadata Database Service on Cloud@Customer
  Review the list of resource-types specific to Oracle Exadata Database Service on Cloud@Customer.
- Permissions and API operation details for Exadata Infrastructures
- Permissions and API operation details for exascale-db-storage-vaults
- Exadata Infrastructure Event Types
  Review the list of event types that Exadata Infrastructure instances emit.
- Exascale DB Storage Vaults Event Types
  Review the list of event types that exascale-db-storage-vaults emit.

Microsoft Entra ID (MS-EI) Integration with Oracle Exadata Database Service on Cloud@Customer

- **Services**: Database
- **Release Date**: November 08, 2024

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

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

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

- Authenticating and Authorizing Microsoft Azure Active Directory Users for Oracle Databases in the Oracle Database 19c Security Guide.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Authenticating and Authorizing Microsoft Azure Users for Oracle Databases in the Oracle AI Database 26ai Security Guide.

Related Topics

- Authenticating and Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Cloud@Customer
  An Oracle Database can be configured for Microsoft Azure users of Microsoft Entra ID to connect using single sign-on authentication.

Support for Concurrent Data Guard, Container Database (CDB), and Pluggable Database (PDB) Operations

- **Services**: Database
- **Release Date**: November 05, 2024

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

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

Related Topics

- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Terminate a Database
  You can terminate a database and thereby remove the terminated database from the Cloud Control Plane.
- Using the Console to Create a Pluggable Database
  To create a pluggable database with the console, use this procedure.
- Using the Console to Delete a Pluggable Database
  The PDB must be available and stopped to use this procedure.
- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.
- Using the Console To Perform a Database Switchover
  You can initiate a switchover operation on a standby database that is a member of the Data Guard Group.
- Using the Console To Perform a Database Failover
  You can initiate a failover operation on a standby database that is a member of the Data Guard Group.
- Using the Console 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.

Delegate Access Control for ExaDB-C@C

- **Services**: Database
- **Release Date**: November 01, 2024

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

Related Topics

- Overview of Delegate Access Control
- Exadata Cloud@Customer Security Controls
- Oracle Engineered System Deployment and Infrastructure Support

Cost and Usage Attribution for Pluggable Databases (PDBs)

- **Services**: Database
- **Release Date**: October 23, 2024

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

Related Topics

- Cost and Usage Attribution for Pluggable Databases (PDBs)

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Generate Attributed Cost Analysis Report for Pluggable Databases
  Follow the steps below to view the attributed costs based on CPU utilization for all pluggable databases within a VM Cluster.
- Network Requirements for Oracle Exadata Database Service on Cloud@Customer
  To provide secure and reliable network connectivity for different application and management functions, Exadata Database Service on Cloud@Customer uses different networks.
- Overview of VM Cluster Node Subsetting
  VM Cluster Node Subsetting enables you to allocate a subset of database servers to new and existing VM clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.

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

- **Services**: Database
- **Release Date**: October 05, 2024

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

Create Data Guard Associations:

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

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

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

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

Related Topics

- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.
- Using the Console To Perform a Database Switchover
  You can initiate a switchover operation on a standby database that is a member of the Data Guard Group.
- Using the Console To Perform a Database Failover
  You can initiate a failover operation on a standby database that is a member of the Data Guard Group.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Granular Permissions for VM Cluster Update Operations

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

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

You can now assign specific permissions for VM Cluster operations, such as allowing a selected set of users to only scale memory or CPU, scale local/Exadata storage, or add SSH keys to a VM cluster.

For more information, see Permissions and API operation details for VM Clusters.

Enhancements to Quarterly Exadata Infrastructure Maintenance Planning and Execution

- **Services**: Database
- **Release Date**: August 27, 2024

Infrastructure maintenance is performed as a single scheduled activity based on customer preferences and includes all infrastructure components. Depending on the number of components, infrastructure maintenance runs can take anywhere between 12 - 30 hours (or longer for infrastructures with elastic expansion).

Infrastructure components include:

- DB Servers
- Storage Servers
- Network Switches

With this enhancement, you will have the flexibility to plan and apply quarterly infrastructure updates to fit smaller maintenance windows. Based on customer-preferred timeslots best suited to your business needs, Oracle automation will perform maintenance on specific infrastructure components across these maintenance windows to ensure all components have software updates applied to meet compliance guidelines.

Related Topics

- Configure Oracle-Managed Infrastructure Maintenance
  Oracle performs the updates to all of the Oracle-managed infrastructure components on Oracle Exadata Database Service on Cloud@Customer.
- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Exadata Infrastructure Maintenance Event Types
  Review the list of event types that Exadata Infrastructure Maintenance emits.
- Resource-Types for Oracle Exadata Database Service on Cloud@Customer
  Review the list of resource-types specific to Oracle Exadata Database Service on Cloud@Customer.
- Permissions and API operation details for Scheduling Policies

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Permissions and API operation details for Scheduling Windows
- Permissions and API operation details for Scheduling Plan
- Permissions and API operation details for Scheduled Action
- Permissions and API operation details for Execution Windows
- Permissions and API operation details for Execution Action
- Permissions Required for Each API Operation
  Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.

Enable Unified Auditing While Creating a Database Home

- **Services**: Database
- **Release Date**: July 23, 2024

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

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

> **Note**
>
> You cannot disable Unified Auditing after provisioning the Database Home

Related Topics

- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

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

- **Services**: Database
- **Release Date**: July 17, 2024

Integrate your on-premises Oracle Key Vault (OKV) with Oracle Exadata Database Service on Cloud@Customer and use customer-managed keys stored in Oracle Key Vault to secure your critical data.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Managing Encryption Keys on External Devices
  Learn how to store and manage database encryption keys.

Ability to Increase the Size of Guest VM Local File Systems

- **Services**: Database
- **Release Date**: May 09, 2024

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

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

Related Topics

- Estimating How Much Local Storage You Can Provision On Your VMs
- Scaling Local Storage
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Scale the Resources on a VM Cluster
  Starting in Oracle Exadata Database Service on Cloud@Customer, you can scale up or down multiple resources at the same time. You can also scale up or down resources one at a time.
- Oracle Exadata X10M System Model Specifications
  Review the technical specifications of X10M system model.
- Oracle Exadata X9M-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.
- Oracle Exadata X8M-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.
- Oracle Exadata X8-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.
- Oracle Exadata X7-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.

Create and Use Custom Software Images

- **Services**: Database
- **Release Date**: May 06, 2024

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

Related Topics

- Manage Software Images
- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- Using the Console to Perform Update Operation on a Database Home
  Learn to apply patches on a Database Home.
- Using the Console to Perform Grid Infrastructure Updates
  Learn to apply Grid Infrastructure updates on a VM cluster.
- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.

Oracle AI Database 26ai on Oracle Exadata Database Service on Cloud@Customer

- **Services**: Database
- **Release Date**: May 02, 2024

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

Related Topics

- Using the Console to Upgrade Oracle Grid Infrastructure of a VM Cluster
- Restoring a Database with the Console
  To restore a database to a point in time, to a system change number (SCN), or to the latest backup, use the Oracle Exadata Database Service on Cloud@Customer Console.
- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

Change the Home Region of ExaDB-C@C Infrastructure

- **Services**: Database

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Release Date**: February 21, 2024

With this enhancement, you can change the home OCI region that your ExaDB-C@C infrastructure connects to. This is a Field Engineer assisted operation and there is no service downtime while the home region change is in progress.

> **Note**
>
> Changing the home region of your ExaDB-C@C Infrastructure will not affect billing.

Related Topics

- About Changing the Home Region of Oracle Exadata Database Service on Cloud@Customer Infrastructure

Enhancements to Serial Console Functionality

- **Services**: Database
- **Release Date**: December 13, 2023

These new features include:

- Serial Console access via OCI Cloud Shell
- Console History

With this new feature, you can now easily connect to the serial console of your virtual machines to perform corrective actions as well as review and audit previous activities performed via the serial console by other users.

> **Note**
>
> - You cannot concurrently connect to more than one DB node using Cloud Shell. As an example, if you have an open connection to *DBnode1* and want to connect to *DBnode2*, you must first exit the active Cloud Shell from *DBnode1* and then establish a connection to *DBnode2*.
> - Cloud Shell access to the serial console require proper IAM permissions for Cloud Shell, see OCI Cloud Shell documentation for details. Also, to access the serial console and to use console history, firewall rules must be configured so that the Control Plane Server (CPS) can access the necessary OCI endpoints. Please review Table 3-2 details for Object Storage and VM console connectivity requirements.

Related Topics

- Using Cloud Shell to Connect to the Serial Console
- Displaying the Console History for a Virtual Machine
- Permissions Required for Each API Operation
  Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- Serial Console History Event Types
  Review the list of new event types that serial console history emits.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Resource-Types for Oracle Exadata Database Service on Cloud@Customer
  Review the list of resource-types specific to Oracle Exadata Database Service on Cloud@Customer.
- Permissions and API operation details for DB Node Console History

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

- **Services**: Database
- **Release Date**: December 07, 2023

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

Related Topics

- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

VM Cluster on a Single VM

- **Services**: Database
- **Release Date**: October 24, 2023

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

Related Topics

- Manage VM Clusters
  Learn how to manage your VM clusters on Oracle Exadata Database Service on Cloud@Customer.

Enhancement to Pluggable Database (PDB) Management

- **Services**: Database
- **Release Date**: October 11, 2023

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

Related Topics

- Manage Pluggable Databases on Oracle Exadata Database Service on Cloud@Customer
  Learn to manage pluggable databases on Oracle Exadata Database Service on Cloud@Customer.
- Permissions Required for Each API Operation
  Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- Permissions and API operation details for Pluggable Databases (PDBs)
- Pluggable Database Event Types
  Review the list of event types that Pluggable Databases emit.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Manage Administrator (SYS User) and TDE Wallet Passwords

- **Services**: Database
- **Release Date**: October 02, 2023

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

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

Related Topics

- Using the Console to Manage SYS User and TDE Wallet Passwords
  Learn to manage administrator (SYS user) and TDE wallet passwords.

Update Guest VM Operating System to Oracle Linux 8

- **Services**: Database
- **Release Date**: August 03, 2023

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

Related Topics

- Updating Guest VM Operating System
  Learn to update the operating system image on Oracle Exadata Database Service on Cloud@Customer VM cluster nodes in an automated manner from the OCI console and APIs.
- Supported Software Versions and Update Restrictions
  Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):
- Using the Console to Update Guest VM Operating System
  To update the guest VM operating system with the Console, be prepared to provide values for the fields required.
- Using the Console to Add VMs to a Provisioned Cluster
  To add virtual machines to a provisioned cluster, use this procedure.

Exadata Fleet Update

- **Services**: Database
- **Release Date:** August 02, 2023

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

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

For more information, see:

- Exadata Fleet Update Overview
- Exadata Fleet Update service API

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

- **Services**: Database
- **Release Date**: June 28, 2023

Integrate your on-premises Oracle Key Vault (OKV) with Oracle Exadata Database Service on Cloud@Customer to secure your critical data on-premises.

Oracle Key Vault integration enables you to take complete control of your encryption keys and store them securely on an external, centralized key management device.

Related Topics

- Managing Encryption Keys on External Devices
  Learn how to store and manage database encryption keys.

X10M System Support

- **Services**: Database
- **Release Date**: June 21, 2023

Oracle Exadata Cloud@Customer has been extended to support X10M system.

Related Topics

- System Configuration Options for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.
- Oracle Exadata X10M System Model Specifications
  Review the technical specifications of X10M system model.
- Estimating How Much Local Storage You Can Provision On Your VMs
- Overview of Elastic Storage Expansion
  With elastic storage expansion, you can dynamically increase your storage capacity to meet your growing workload requirements.
- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.

Manage Serial Console Access to ExaDB-C@C Systems

- **Services**: Database
- **Release Date**: June 14, 2023

You can create and delete serial console connections to your ExaDB-C@C systems to diagnose and resolve VM guest operating system issues using an SSH connection in case standard SSH access to the VMs is not possible.

Requirements: The use of the serial console feature requires Exadata Infrastructure version 22.1.10 or higher for 22.X users and version 23.1.1 or higher for 23.X users. The serial console feature will be available on any new VM Clusters created immediately but will only be available on previously existing VM Clusters after the next quarterly maintenance cycle. 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 for meeting 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.

> **Note**
>
> The following two endpoints have been added for the Control Plane Server (CPS). Use these URL formats, replacing `oci_region` with your region.
>
> - `console1.exacc.`*`oci_region`*`.oci.oraclecloud.com`
> - `console2.exacc.`*`oci_region`*`.oci.oraclecloud.com`

These must be allowed at your firewall in order for serial console connections to function. See *Network Requirements for Oracle Exadata Database Service on Cloud@Customer* Table 3-2 for further details.

Related Topics

- Network Requirements for Oracle Exadata Database Service on Cloud@Customer
  To provide secure and reliable network connectivity for different application and management functions, Exadata Database Service on Cloud@Customer uses different networks.
- Troubleshooting Virtual Machines Using Console Connections
  You can troubleshoot malfunctioning virtual machines using console connections. For example, a previously working Guest VM stops responding.
- Create the Virtual Machine Serial Console Connection
- Resource-Types for Oracle Exadata Database Service on Cloud@Customer
  Review the list of resource-types specific to Oracle Exadata Database Service on Cloud@Customer.
- Permissions and API operation details for DB Nodes

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Permissions Required for Each API Operation
  Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- Permissions and API operation details for DB Node Console Connection
- Serial Console Connection Event Types
  Review the list of event types that serial console connection emits.
- Viewing Audit Log Events
  Oracle Cloud Infrastructure Audit service provides records of API operations performed against supported services as a list of log events.

Interim Software Updates

- **Services**: Database
- **Release Date**: June 07, 2023

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

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

Related Topics

- Interim Software Updates
  For authorized environments, learn how to download interim software updates.
- Permissions Required for Each API Operation
  Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- Permissions and API operation details for Interim Software Updates
- Interim Software Updates Event Types

Link Aggregation Control Protocol (LACP) Support for Client and Backup Networks

- **Services**: Database
- **Release Date**: May 23, 2023

Oracle Exadata Database Service on Cloud@Customer physical networks (client and backup) are configured to use active-backup by default. This mode is recommended for most customers. However, support for LACP (802.3ad active/active bonding mode) has now been added also for those customers who require it.

You can configure client and backup networks with LACP while provisioning the Exadata Infrastructure. The network bonding mode is set at the infrastructure level and will apply globally to all VM Cluster Networks, future or existing. Since the setting gets applied to each network interface individually, you have the flexibility to configure client and backup network interfaces independently from one another. For example, you can configure the backup network to LACP and the client network to Active-Backup, if desired.

Copyright © 2008, 2026, Oracle and/or its affiliates.

On an existing Exadata Infrastructure that is now in use, you may change the network bonding mode from Active-Backup to LACP (or vice versa). However, note that this is a non-rolling update process and bonding modes are changed simultaneously on all Database Servers. You are required to manage your own switch settings on your network accordingly. You may experience a network outage and should plan for application downtime until your switch settings and ExaDB-C@C settings are made to match.

The use of LACP requires both the server and the switch(es) have compatible settings to allow LACP to work properly. To use LACP for ExaDB-C@C, you must configure your network switches such that they are compatible with the following parameter from the Linux `ifcfg-bondethx` configuration file:

```
BONDING_OPTS="mode=802.3ad miimon=100 downdelay=200 updelay=200 lacp_rate=1
xmit_hash_policy=layer3+4"
```

Note that the `BONDING_OPTS` line in the host operating system cannot be changed, so customer switch settings must be compatible with the parameters stated above without alteration.

Related Topics

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Edit Network Bonding
  To edit the client and backup network bonding mode, follow these steps.

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

- **Services**: Database
- **Release Date**: May 08, 2023

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

> **Note**
>
> When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.
>
> Potentially one hour of metrics will be lost between network restore and the first metric posted.

Related Topics

- Monitor Metrics for VM Cluster Resources
  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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Metrics for Oracle Exadata Database Service on Cloud@Customer in the Monitoring Service
  Review the metrics emitted by the Oracle Exadata Database Service on Cloud@Customer service in the `oci_database_cluster` and `oci_database` namespaces for Oracle Databases.

Enable Data Guard with Database Software Image (Custom Image)

- **Services**: Database
- **Release Date**: April 11, 2023

If you have provisioned DB Homes with a Database Software Image (DSI), then enabling Data Guard operations will default to the same DSI as the primary DB Home.

With this enhancement.

- You can change the automatically selected primary DSI to a different DSI for the standby database.
- If the DSI used to create the primary database is not available, then enabling Data Guard will use the latest Oracle-published image. Optionally, you can configure the standby database to have a DSI of your choice.

In both cases, the system will warn you of potential issues with having different images for the primary and standby databases.

Related Topics

- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.

Use Identity and Access Management (IAM) Authentication with Oracle Exadata Database Service on Cloud@Customer

- **Services**: Database
- **Release Date**: March 28, 2023

You can configure Oracle Database on Oracle Exadata Database Service on Cloud@Customer system to use Oracle Cloud Infrastructure Identity and Access Management (OCI IAM) authentication and authorization to allow IAM users to access the database with IAM credentials.

> **Note**
>
> Oracle Exadata Database Service on Cloud@Customer integration with OCI IAM is supported in commercial tenancies with identity domains as well as the legacy OCI IAM, which does not include identity domains. OCI IAM with identity domains was introduced with new OCI tenancies created after November 8, 2021. Only default domain OCI IAM users are supported with the new identity domains.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Connect to an Oracle Database using Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users
  You can configure Oracle Exadata Database Service on Cloud@Customer 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.

Create Data Guard Associations Across Regions

- **Services**: Database
- **Release Date**: February 14, 2023

Create Data Guard associations across regions within your tenancy. This will help you implement an effective disaster recovery plan to protect your data against natural calamities.

> **Note**
>
> Active Data Guard or Data Guard associations can only be created for Transparent Data Encryption (TDE) enabled databases, not for Oracle Key Vault (OKV) enabled databases.

Related Topics

- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.

Multi-Rack Elastic Compute and Storage Expansion

- **Services**: Database
- **Release Date**: January 18, 2023

For certain deployments, you may want to expand beyond the maximum number of compute and storage servers supported within a single Exadata Infrastructure rack. With this enhancement to compute and storage expansion, you can now have additional compute and storage spanning multiple racks available for VM Clusters provisioned on the Exadata Infrastructure.

Base systems are not eligible for multirack expansion. Multi-rack applies only to the following shapes:

- X8M-2 standard shapes
- X9M-2 standard shapes

Note: Exadata Infrastructure with more than 8 compute or 12 storage servers requires a multi-rack deployment.

The "multi-rack" selection will drive the maximum number of compute and storage servers you can specify in the server configuration.

- **Single Rack:**
  - All Systems: 8 compute and 12 storage

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Mulit-rack:**
  - **X7, X8, and all Base Systems:** Not Applicable
  - **X8M and X9M Sytems:** 32 compute and 64 storage servers

Pre-deployment process:

- If multi-rack is required or desired, the FE will provide the customer with a JSON configuration file.
- The customer uploads the config file when creating or expanding the Exadata Infrastructure.

Deployment process:

- When the customer creates the infrastructure and selects multi-rack, they must upload the JSON so the control plane can correctly generate the config bundle for download. The JSON file will be used to properly configure the additional components. Uploading the JSON will only be mandatory for deployments identified as multi-rack.
- If a customer fails to identify an infrastructure as multi-rack and subsequently they determine it is multi-rack, they will have to delete that infrastructure and recreate it.
- If a customer expands an existing infrastructure and remains within a single rack (only adding servers), then there is no change to the current deployment process.
- When a customer expands an existing infrastructure and adds a new rack, or expands an infrastructure with an existing expansion rack, they will need to choose the multi-rack deployment type and upload a new Multi-Rack Configuration File (JSON).

Related Topics

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Download a File Containing Configuration Data
  To download an Oracle Exadata Database Service on Cloud@Customer configuration file, complete this procedure.
- Using the Console to Activate Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To activate Oracle Exadata Database Service on Cloud@Customer infrastructure, ensure that you meet the prerequisites, and complete this procedure.
- Using the Console to Create a VM Cluster Network
  To create your VM cluster network with the Console, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Validate a VM Cluster Network
  You can only validate a VM cluster network if its current state is **Requires Validation**, and if the underlying Exadata infrastructure is activated.
- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.
- Using the Console to Download Scale Infrastructure Compute and Storage Configuration File
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Activate New Database and Storage Servers
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.
- Using the Console to Add Another DB Server to the VM Cluster Network
  To a DB server to the VM Cluster Network, complete this procedure.
- Using the Console to Add VMs to a Provisioned Cluster
  To add virtual machines to a provisioned cluster, use this procedure.
- Using the Console to Remove a VM from a VM Cluster
  To remove a virtual machine from a provisioned cluster, use this procedure.
- Using the Console to Remove DB Servers from a VM Cluster Network
  To remove a database server, use this procedure.
- Using the Console to Edit a VM Cluster Network
  To edit a VM Cluster Network, use this procedure. You can only edit a VM cluster network that is not associated with a VM cluster.
- Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption
  To make storage capacity from the new servers for VM clusters consumption, complete this procedure.
- Using the Console to View Details of Oracle Exadata Database Service on Cloud@Customer Infrastructure with Scaled Compute and Storage Capacity
  To view the storage capacity from the new storage server, or scaled compute capacity, use this procedure.
- Compute Expansion Event Types
  Review the list of event types that compute expansion emits.

Automatic Diagnostic Collection

- **Services**: Database
- **Release Date**: October 25, 2022

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

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

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

Related Topics

- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Database Service Events
  The Database Service emits events, which are structured messages that indicate changes in resources.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Elastic Compute Expansion

- **Services**: Database
- **Release Date**: October 08, 2022

Elastic Compute Expansion enables the addition of an arbitrary number of DB Servers to Exadata Cloud@Customer infrastructure. In conjunction with the previously released Elastic Storage Expansion feature, new infrastructure installations can now be better customized by provisioning a discrete number of DB Servers and Storage Servers. Also, compute capacity for the existing infrastructure deployments can now be expanded by adding individual DB Servers in a manner similar to Elastic Storage Expansion.

In addition, significant changes have been made to the VM Cluster Network object to facilitate the ability to subset network resources to certain DB Servers.

Follow the links below for instructions on how to provision infrastructure and activate. Once the DB Servers are activated, they are recognized and available for use immediately, but the additional resources need to be added to your VM Clusters. This is not done automatically. First, add VM Cluster Network resources and then you can add additional VMs to your VM Clusters, following the instructions in the links below.

Related Topics

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Download a File Containing Configuration Data
  To download an Oracle Exadata Database Service on Cloud@Customer configuration file, complete this procedure.
- Using the Console to Activate Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To activate Oracle Exadata Database Service on Cloud@Customer infrastructure, ensure that you meet the prerequisites, and complete this procedure.
- Using the Console to Create a VM Cluster Network
  To create your VM cluster network with the Console, be prepared to provide values for the fields required for configuring the infrastructure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Validate a VM Cluster Network
  You can only validate a VM cluster network if its current state is **Requires Validation**, and if the underlying Exadata infrastructure is activated.
- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.
- Using the Console to Download Scale Infrastructure Compute and Storage Configuration File
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.
- Using the Console to Activate New Database and Storage Servers
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.
- Using the Console to Add Another DB Server to the VM Cluster Network
  To a DB server to the VM Cluster Network, complete this procedure.
- Using the Console to Add VMs to a Provisioned Cluster
  To add virtual machines to a provisioned cluster, use this procedure.
- Using the Console to Remove a VM from a VM Cluster
  To remove a virtual machine from a provisioned cluster, use this procedure.
- Using the Console to Remove DB Servers from a VM Cluster Network
  To remove a database server, use this procedure.
- Using the Console to Edit a VM Cluster Network
  To edit a VM Cluster Network, use this procedure. You can only edit a VM cluster network that is not associated with a VM cluster.
- Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption
  To make storage capacity from the new servers for VM clusters consumption, complete this procedure.
- Using the Console to View Details of Oracle Exadata Database Service on Cloud@Customer Infrastructure with Scaled Compute and Storage Capacity
  To view the storage capacity from the new storage server, or scaled compute capacity, use this procedure.
- Compute Expansion Event Types
  Review the list of event types that compute expansion emits.

Oracle Standard Tagging for Resources on Oracle Exadata Database Service on Cloud@Customer

- **Services**: Database
- **Release Date**: August 31, 2022

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Tagging Oracle Exadata Database Service on Cloud@Customer 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.

Exadata Infrastructure Maintenance History

- **Services**: Database
- **Release Date**: August 02, 2022

From the Maintenance History page, you can now click on a maintenance to view maintenance history details, containing the same maintenance detail information available for a scheduled or in-progress maintenance. Maintenance history is available for both successful and failed maintenance.

Related Topics

- View the Maintenance History
  Learn how to view the maintenance history for an Oracle Exadata Database Service on Cloud@Customer Infrastructure.

Monthly ExaDB-C@C Infrastructure Security Maintenance

- **Services**: Database
- **Release Date**: July 27, 2022

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

Related Topics

- Overview of Monthly Security Maintenance
  Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.
- View or Edit a Scheduled Security Maintenance
  Learn how to view and edit the next scheduled security maintenance.
- Exadata Infrastructure Maintenance Event Types
  Review the list of event types that Exadata Infrastructure Maintenance emits.
- Overview of the Quarterly Infrastructure Maintenance Process
  By default, infrastructure maintenance updates the Exadata database server hosts in a rolling fashion, followed by updating the storage servers.
- Using the Console to Configure Oracle-Managed Infrastructure Updates
  Full Exadata infrastructure software updates are scheduled on a quarterly basis. In addition, important security updates are scheduled monthly. While you cannot opt-out of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal and allows scheduling flexibility to help you plan for them.
- View or Edit Quarterly Maintenance Preferences
  To edit your Oracle Exadata Database Service on Cloud@Customer infrastructure quarterly maintenance preferences, be prepared to provide values for the infrastructure

Copyright © 2008, 2026, Oracle and/or its affiliates.

configuration. The changes you make will only apply to future maintenance runs, not those already scheduled.

- View or Edit a Scheduled Quarterly Maintenance
  Learn how to view and edit the time of the next scheduled maintenance.
- View the Maintenance History
  Learn how to view the maintenance history for an Oracle Exadata Database Service on Cloud@Customer Infrastructure.
- View and Edit Quarterly/Monthly Maintenance While Maintenance is In Progress
  While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for a custom action, you can resume the maintenance prior to the timeout or extend the timeout.

Enhanced Control to Rollback or Retry Failed Guest VM Operating System Update

- **Services**: Database
- **Release Date**: July 27, 2022

You need not mandatorily roll back if applying the Guest VM operating system update fails. A new option, in addition to the current rollback option, has been added to retry and apply the failed update. If you want to apply a different operating system image update on failure, then you will have to first roll back and then apply.

Related Topics

- Using the Console to Rollback or Retry Failed Guest VM Operating System Update
  To update the guest VM operating system with the Console, be prepared to provide values for the fields required.

Concurrently Create or Terminate Oracle Databases in a VM Cluster

- **Services**: Database
- **Release Date**: July 20, 2022

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

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

Support for Rack Serial Number as a System Tag

- **Services**: Database

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Release Date**: July 19, 2022

This enhancement is to display the Exadata cloud@Customer rack serial number in the OCI console under the **Infrastructure details** page in the **General Information** section. The serial number may be required when creating an SR or during a service call.

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

- **Services**: Database
- **Release Date**: May 23, 2022

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

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

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

Related Topics

- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.
- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

VM Guest Exadata OS Image Major Version Update

- **Services**: Database
- **Release Date**: April 20, 2022

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

Related Topics

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

Database Service Events

- **Services**: Database
- **Release Date**: April 12, 2022

Copyright © 2008, 2026, Oracle and/or its affiliates.

Database Service Events feature implementation enables you to get notified about health issues with your Oracle Databases or other components on the Guest VM.

Related Topics

- Overview of Database Service Events
- 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.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Control Plane Server (CPS) Offline Diagnostic Report

- **Services**: Database
- **Release Date**: March 01, 2022

Control Plane Server (CPS) offline diagnostic report assists you in troubleshooting connectivity issues between the CPS and OCI endpoints.

It is your responsibility to maintain and troubleshoot the network infrastructure of your data center. To connect to OCI Region, Exadata Cloud@Customer Gen2 depends on your infrastructure and reliability. Exadata Cloud@Customer's connectivity from OCI Region to Exadata Cloud@Customer's Control Plane Servers (CPS) may be impacted by any changes you make to your infrastructure. Nevertheless, Oracle does not have any control over your firewall or networking.

In the event that the connection between CPS and OCI is broken, the Control Plane Server (CPS) Offline Diagnostic Report provides information that may help you in diagnosing problems in your network infrastructure.

To view the report, do the following:

1. Find the CPS IP addresses.
   For more information, see *Using the Console to View Exadata Infrastructure Network Configuration Details*.
2. From your local network, access the report over HTTP.
   To view the report in HTML format, use `http://<CPSPublicIP>:18080/report`

   To view the report in JSON format, use `http://<CPSPublicIP>:18080/report/json`

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> - You cannot enable or disable Control Plane Server (CPS) Offline Diagnostic Report if the Exadata Infrastructure is in `DISCONNECTED` mode.
> - Every hour, even if no issues are detected on the CPS, the system will still generate and save a diagnostic report in HTML and JSON formats. Whenever a connectivity issue arises between CPS and OCI endpoints, the system generates a report immediately.
> - At any given point in time, the report is available only at the primary Control Plane Server. When generating a report, if the first IP address specified for the Control Plane Servers does not work, then you may try the second IP.

For more information, see *ExaCC gen2: Troubleshooting VPN/WSS connection from Customer Side (Doc ID 2745571.1)*.

Related Topics

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Enabling or Disabling the Control Plane Server Diagnostic Offline Report
  To enable or disable CPS offline report, use this procedure.
- Using the Console to View Exadata Infrastructure Network Configuration Details
  To view network configuration details, follow these steps. Save this information for later use to troubleshoot if you face network issues.
- ExaCC gen2: Troubleshooting VPN/WSS connection from Customer Side (Doc ID 2745571.1)

Enhanced Infrastructure Maintenance Controls

- **Services**: Database
- **Release Date**: March 01, 2022

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

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

Related Topics

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- View or Edit Quarterly Maintenance Preferences
  To edit your Oracle Exadata Database Service on Cloud@Customer infrastructure quarterly maintenance preferences, be prepared to provide values for the infrastructure configuration. The changes you make will only apply to future maintenance runs, not those already scheduled.
- Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure Maintenance Controls
  Oracle Exadata Database Service on Cloud@Customer uses the same API as Oracle Cloud Infrastructure to manage infrastructure maintenance controls.
- Configure Oracle-Managed Infrastructure Maintenance
  Oracle performs the updates to all of the Oracle-managed infrastructure components on Oracle Exadata Database Service on Cloud@Customer.
- Using the Console to Configure Oracle-Managed Infrastructure Updates
  Full Exadata infrastructure software updates are scheduled on a quarterly basis. In addition, important security updates are scheduled monthly. While you cannot opt-out of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal and allows scheduling flexibility to help you plan for them.
- Exadata Infrastructure Maintenance Event Types
  Review the list of event types that Exadata Infrastructure Maintenance emits.

Manage Pluggable Databases on Exadata Cloud@Customer

- **Services**: Database
- **Release Date**: February 16, 2022

Create and manage pluggable databases (PDBs) in Oracle Exadata Cloud@Customer systems using the Console and APIs.

Related Topics

- Manage Pluggable Databases on Oracle Exadata Database Service on Cloud@Customer
  Learn to manage pluggable databases on Oracle Exadata Database Service on Cloud@Customer.

Allow Customers to Choose Data Guard Type

- **Services**: Database
- **Release Date**: February 16, 2022

Select a Data Guard type, Active Data Guard or Data Guard, based on the Oracle Database software license type you have deployed.

Related Topics

- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.
- Using the Console To View and Edit Data Guard Group Details
  You can switch between Data Guard types based on the Oracle Database software license type you have deployed.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

- **Services**: Database
- **Release Date**: February 16, 2022

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

Related Topics

- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.

Specify db_unique_name and SID for Primary and Standby Databases in Data Guard Association

- **Services**: Database
- **Release Date**: October 20, 2021

Oracle databases are identified by three important names: `db_name`, `db_unique_name`, and `instance_name` (SID). This new feature provides consistent database naming controls across primary and standby databases and allows entering the `db_unique_name` and the SID prefix on both primary and standby databases. This helps support different naming conventions to manage the Oracle Database fleet.

Related Topics

- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.
- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.

VM Cluster Node Subsetting

- **Services**: Database
- **Release Date**: October 19, 2021

> **Note**
>
> VM Cluster Node Subsetting feature is now available in all OCI commercial regions.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Overview of VM Cluster Node Subsetting
  VM Cluster Node Subsetting enables you to allocate a subset of database servers to new and existing VM clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Add VMs to a Provisioned Cluster
  To add virtual machines to a provisioned cluster, use this procedure.
- Using the Console to View a List of DB Servers on an Exadata Infrastructure
  To view a list of database server hosts on an Oracle Exadata Cloud@Customer system, use this procedure.
- Using the Console to Remove a VM from a VM Cluster
  To remove a virtual machine from a provisioned cluster, use this procedure.
- Permissions Required for Each API Operation
  Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- Permissions and API operation details for DB Servers
- VM Node Subsetting Event Types
  Review the list of event types that VM Node Subsetting emits.
- Adding a VM to a VM Cluster Fails
- CPU Offline Scaling Fails

X9M-2 System Support

- **Services**: Database
- **Release Date**: September 28, 2021

Oracle Exadata Cloud@Customer comes in different infrastructure shapes to support workloads of different sizes. In this release, the capability of Oracle Exadata Cloud@Customer has been extended to support X9M-2 system.

For more information, see:

- *System Configuration Options for Oracle Exadata Cloud@Customer*
- *Oracle Exadata X9M-2 System Model Specifications*
- *Estimating How Much Local Storage You Can Provision to Your VMs*
- *Overview of Elastic Storage Expansion*

Related Topics

- System Configuration Options for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.
- Oracle Exadata X9M-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.
- Estimating How Much Local Storage You Can Provision On Your VMs

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Overview of Elastic Storage Expansion
  With elastic storage expansion, you can dynamically increase your storage capacity to meet your growing workload requirements.

Customize SCAN Listener Port

- **Services**: Database
- **Release Date**: August 27, 2021

You can now specify a SCAN listener port (TCP/IP) within the permissible range while creating a VM cluster network resource. For more information, see:

- *Using the Console to Create a VM Cluster Network*
- *Using the Console to View SCAN Listener Port Configured*

Related Topics

- Using the Console to Create a VM Cluster Network
  To create your VM cluster network with the Console, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to View SCAN Listener Port Configured
  You can only edit a VM cluster network that is not associated with a VM cluster.

Creating DG Association/Standby Database Using Existing Database Home

- **Services**: Database
- **Release Date**: August 08, 2021

Select an existing Database Home or create a new one for the standby while enabling Data Guard association. For more information, see *Using the Console to Enable Data Guard on an Exadata Cloud@Customer System*.

Related Topics

- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.

Upgrading Oracle Grid Infrastructure on an Exadata Cloud@Customer VM Cluster

- **Services**: Database
- **Release Date**: July 20, 2021

Upgrade Oracle Grid Infrastructure (GI) on an Exadata Cloud@Customer VM cluster using the Oracle Cloud Infrastructure Console or APIs. For more information, see *Upgrading Oracle Grid Infrastructure on an Exadata Cloud@Customer VM Cluster*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Upgrading Oracle Grid Infrastructure on an Oracle Exadata Database Service on Cloud@Customer VM Cluster
  Learn to upgrade Oracle Grid Infrastructure on an Oracle Exadata Database Service on Cloud@Customer VM cluster using the Oracle Cloud Infrastructure Console or API.

Updating Guest VM Operating System

- **Services**: Database
- **Release Date**: July 20, 2021

Update the operating system image on Exadata Cloud@Customer VM cluster nodes in an automated manner from the OCI console and APIs. For more information, see *Updating Guest VM Operating System*.

Related Topics

- Updating Guest VM Operating System
  Learn to update the operating system image on Oracle Exadata Database Service on Cloud@Customer VM cluster nodes in an automated manner from the OCI console and APIs.

Upgrading Oracle Databases

- **Services**: Database
- **Release Date**: July 20, 2021

Upgrade Oracle Database 19c (Long Term Release) using the Console and the APIs. For more information, see *Upgrading Oracle Databases*.

Related Topics

- Upgrading Oracle Databases
  Learn to upgrade an Exadata database instance to Oracle Database 19c and Oracle AI Database 26ai by using the Console and the API.

Download Network Validation Report

- **Services**: Database
- **Release Date**: July 20, 2021

Validate and inspect the network validation failure report without active involvement from Oracle Cloud Ops in troubleshooting networking configuration issues. For more information, see *Using the Console to Download Network Validation Report*.

Related Topics

- Using the Console to Download Network Validation Report
  Learn to validate and inspect the network validation failure report without active involvement from Oracle Cloud Ops in troubleshooting networking configuration issues.

Elastic Storage Expansion

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Services**: Database
- **Release Date**: June 15, 2021

Expand the storage associated with your Exadata Cloud@Customer infrastructure and make them available for VM cluster allocation, both during and post infrastructure provisioning. For more information, see:

- *Overview of Elastic Storage Expansion*
- *Using the Console to Scale Infrastructure Storage*
- *Using the Console to Download Scale Infrastructure Storage Configuration File*
- *Using the Console to Activate New Storage Servers*
- *Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption*
- *Using the Console to View Details of Exadata Cloud@Customer Infrastructure with Scaled Storage Capacity*
- *Permissions Required for Each API Operation*
- *exadata-infrastructures*
- *Storage Expansion Event Types*

Related Topics

- Overview of Elastic Storage Expansion
  With elastic storage expansion, you can dynamically increase your storage capacity to meet your growing workload requirements.
- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.
- Using the Console to Download Scale Infrastructure Compute and Storage Configuration File
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.
- Using the Console to Activate New Database and Storage Servers
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.
- Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption
  To make storage capacity from the new servers for VM clusters consumption, complete this procedure.
- Using the Console to View Details of Oracle Exadata Database Service on Cloud@Customer Infrastructure with Scaled Compute and Storage Capacity
  To view the storage capacity from the new storage server, or scaled compute capacity, use this procedure.
- Permissions Required for Each API Operation
  Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- Permissions and API operation details for Exadata Infrastructures
- Storage Expansion Event Types
  Review the list of event types that storage expansion emits.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Oracle Database Software Images

- **Services**: Database
- **Release Date**: April 08, 2021

Use Database Software Image resource type to create databases and Oracle Database Homes, and to patch databases. For more information, see:

- *Oracle Database Software Images*
- *Using the Console to Create Oracle Database Home on Exadata Cloud@Customer*
- *Using the Console to Perform a Patch Operation on a Database Home*

Related Topics

- Manage Software Images
- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- Using the Console to Perform Update Operation on a Database Home
  Learn to apply patches on a Database Home.

Exadata Cloud@Customer Infrastructure Patching Automation

- **Services**: Database
- **Release Date**: December 08, 2020

You can now schedule a maintenance window for Oracle-managed Exadata Cloud@Customer infrastructure patching. For more information, see Using the Console to Configure Oracle-Managed Infrastructure Updates.

Customer Maintenance Contacts

- **Services**: Database
- **Release Date**: September 22, 2020

Maintenance contacts are required for service request based communications for hardware replacement and other maintenance events.

Add a primary maintenance contact and optionally add a maximum of nine secondary contacts. Both the primary and secondary contacts receive all notifications about hardware replacement, network issues, and software maintenance runs.

You can promote any secondary contacts as the primary anytime you want. When you promote a secondary contact to primary, the current primary contact will be demoted automatically to secondary.

For more information, see:

- *Using the Console to Create Infrastructure*
- *Managing Infrastructure Maintenance Contacts*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Defining Maintenance Contacts for Exadata Cloud@Customer
  Oracle Cloud Operations will use these contacts to communicate with customers about infrastructure-related events, including hardware failures and replacements, quarterly maintenance patches, and other operational updates or incidents

X8M-2 System Support

- **Services**: Database
- **Release Date**: August 18, 2020

Oracle Exadata Cloud@Customer comes in different infrastructure shapes to support workloads of different sizes. In this release, the capability of Oracle Exadata Cloud@Customer has been extended to support X8M-2 system.

For more information, see:

- *System Configuration Options for Oracle Exadata Cloud@Customer*
- *Oracle Exadata Cloud@Customer X8M-2 System Specifications*
- *Network Requirements for Oracle Exadata Cloud@Customer*

Related Topics

- System Configuration Options for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.
- Oracle Exadata X8M-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.
- Network Requirements for Oracle Exadata Database Service on Cloud@Customer
  To provide secure and reliable network connectivity for different application and management functions, Exadata Database Service on Cloud@Customer uses different networks.

Enable and Manage Data Guard Associations

- **Services**: Database
- **Release Date**: August 18, 2020

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

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.

For more information, see:

- *Using Oracle Data Guard with Exadata Cloud@Customer*

Copyright © 2008, 2026, Oracle and/or its affiliates.

- *Using the API To Manage Data Guard Associations on an Exadata Cloud@Customer System*
- *Permissions Required for Each API Operation*
- *Data Guard Event Types*

Related Topics

- Use Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer
  Learn to configure and manage Data Guard Groups in your VM cluster.
- Using the API to Manage Data Guard Associations on an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to use the API to manage Data Guard associations on an Oracle Exadata Database Service on Cloud@Customer system.
- Permissions Required for Each API Operation
  Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- Data Guard Event Types
  Review the list of event types that Data Gurad groups and Data Guard Associations emit.

Oracle Exadata Cloud@Customer Deployment Assistant

- **Services**: Database
- **Release Date**: August 08, 2020

Oracle Exadata Cloud@Customer Deployment Assistant is an automated installation and configuration tool that enables you to set up your Oracle Exadata Cloud@Customer machine and create an Oracle Database instance with minimal effort.

Oracle Grid Infrastructure and Oracle Database Patching

- **Services**: Database
- **Release Date**: July 28, 2020

You can now view, pre-check, and apply Oracle Grid Infrastructure and Oracle Database patches by using the Oracle Cloud Infrastructure Console, API, or CLI. This functionality includes the ability to easily patch a database by moving it to a different Database Home. Similarly, you can easily roll back the version of the database by moving it back to its original Database Home.

For information and instructions, see:

- *Patching and Updating an Exadata Cloud@Customer System*
- *Troubleshooting Exadata Cloud@Customer Systems*
- *Database and Grid Infrastructure Patching Event Types*

Related Topics

- Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to perform patching operations on Exadata database virtual machines and Database Homes by using the Console, API, or the CLI.
- Troubleshooting Oracle Exadata Database Service on Cloud@Customer Systems
  These topics cover some common issues you might run into and how to address them.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Database and Grid Infrastructure Patching Event Types
  Review the list of event types that Database and Grid Infrastructure Patching emit.

Per-Second Billing for OCPU Usage

- **Services**: Database
- **Release Date**: July 14, 2020

Oracle Exadata Database Service on Cloud@Customer Gen2 uses per-second billing for OCPUs. This means that OCPU usage is billed by the second, with a minimum usage period of 1 minute.

Shared Database Home Resource Tags

- **Services**: Database
- **Release Date**: June 23, 2020

Add, update, and remove tags applied to a shared Database Home resource.

For information and instructions, see:

- *Using the Console to Create Oracle Database Home on Exadata Cloud@Customer*
- *Using the Console to Create a Database*

Related Topics

- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.

Create and Manage Multiple Virtual Machines per Exadata System (MultiVM)

- **Services**: Database
- **Release Date**: June 13, 2020

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

For information and instructions, see:

- *Using the Console to Create a VM Cluster*
- *Using the Console to Scale the Resources on a VM Cluster*
- *VM Cluster Event Types*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Scale the Resources on a VM Cluster
  Starting in Oracle Exadata Database Service on Cloud@Customer, you can scale up or down multiple resources at the same time. You can also scale up or down resources one at a time.
- VM Cluster Event Types
  Review the list of event types that VM clusters emit.

Scale OCPUs Without Cloud Connectivity

- **Services**: Database
- **Release Date**: June 13, 2020

Oracle Exadata Database Service on Cloud@Customer is considered to be in a "disconnected" mode when there is a loss of connectivity with the Database service control plane running on Oracle Cloud Infrastructure. Scale the CPU core count up or down for a virtual machine in a VM cluster in disconnected mode.

For information and instructions, see:

- *About Using the dbaascli Utility on Exadata Cloud@Customer*
- *dbaascli cpuscale get_status*
- *dbaascli cpuscale update*
- *Exadata Infrastructure Event Types*

Related Topics

- About Using the dbaascli Utility on Oracle Exadata Database Service on Cloud@Customer
  You can use the `dbaascli` utility to perform various database lifecycle and administration operations on Oracle Exadata Database Service on Cloud@Customer such as creating an Oracle Database, patching an Oracle Database, managing pluggable databases (PDBs), scaling the CPU core count in disconnected mode, and more.
- dbaascli cpuscale get_status
  To check the status of current or last scale request performed when network connectivity between the Control Plane Server and OCI region is disrupted, use the `dbaascli cpuscale get_status` command.
- dbaascli cpuscale update
  To scale up or down the CPU core count for a virtual machine in a VM cluster when network connectivity between the Control Plane Server and OCI region is disrupted, use the `dbaascli cpuscale update` command.
- Exadata Infrastructure Event Types
  Review the list of event types that Exadata Infrastructure instances emit.

Configure Oracle Database Character Set and National Character Set

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Services**: Database
- **Release Date**: May 7, 2020

Before you create the database, decide the character set that you want to use.

After a database is created, changing its character set is usually very expensive in terms of time and resources. Such operations may require converting all character data by exporting the whole database and importing it back. Therefore, it is important that you carefully select the database character set at installation time.

For information and instructions, see *Using the Console to Create a Database*.

Related Topics

- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.

Specify a Time Zone While Provisioning Oracle Exadata Database Service on Cloud@Customer Infrastructure

- **Services**: Database
- **Release Date**: May 7, 2020

The default time zone for the Oracle Exadata Database Service on Cloud@Customer infrastructure is UTC. The time is displayed in the UTC format at the operating system and database level. You can choose a different Time Zone while provisioning your Oracle Exadata Database Service on Cloud@Customer infrastructure. However, Oracle recommends setting the database time zone to UTC (0:00), as no conversion of time zones will be required.

For information and instructions, see *Using the Console to Create Infrastructure*.

Related Topics

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.

Shared Database Homes for Oracle Exadata Database Service on Cloud@Customer

- **Services**: Database
- **Release Date**: March 31, 2020

Use a single Oracle Home for multiple databases. Besides saving space, sharing an Oracle Database Home with multiple databases provides the following benefits:

- One-off patches needed for multiple databases only need to be applied to fewer Oracle Homes, reducing patching overhead and administration.
- Space savings, for example, by having only a single Oracle Home per Oracle software version (though multiple are possible).
- Oracle Database patching through moving databases between homes instead of patching the DB Home.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Fallback to move a database to a lower DB Home version without having to rollback a patch.
- Software installation of new Homes and new versions is not interrupting operation of database.
- Patching time of a database is reduced by not having to additionally lay down Oracle binaries.

Use the Oracle Cloud Infrastructure Console, API, or CLI to create and manage shared Oracle Database Homes for your databases on an Oracle Exadata Database Service on Cloud@Customer.

For information and instructions, see *Create Oracle Database Homes on an Exadata Cloud@Customer Systems*.

Related Topics

- Create Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to create Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer.

X7-2 System Support

- **Services**: Database
- **Release Date**: March 19, 2020

Oracle Exadata Database Service on Cloud@Customer comes in different infrastructure shapes to support workloads of different sizes. In this release, the capability of Oracle Exadata Database Service on Cloud@Customer has been extended to support X7-2 system.

For more information, see *Oracle Exadata X7-2 System Model Specifications*.

Related Topics

- Oracle Exadata X7-2 System Model Specifications
  Review the technical specifications of available Exadata System Shapes.

Copyright © 2008, 2026, Oracle and/or its affiliates.

3

Preparing for Oracle Exadata Database Service on Cloud@Customer

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

- Oracle Cloud Infrastructure (OCI) Requirements for Oracle Exadata Database Service on Cloud@Customer
  Learn the basic concepts to get started using Oracle Cloud Infrastructure.
- Site Requirements for Oracle Exadata Database Service on Cloud@Customer
  Review the requirements for provisioning Oracle Exadata Database Service on Cloud@Customer at your site.
- Network Requirements for Oracle Exadata Database Service on Cloud@Customer
  Review the network requirements for provisioning Oracle Exadata Database Service on Cloud@Customer at your site.
- Storage Configuration Requirements for Oracle Exadata Database Service on Cloud@Customer
  With the introduction of Exascale technology in Exadata Database Service on Cloud@Customer, you can configure the Exadata infrastructure to use ASM, Exascale, or a combination of both. When configured to coexist with ASM, you must allocate Exascale storage capacity from the total Exadata storage available on the Infrastructure. The allocated storage capacity will be solely available for Exascale use.
- Checklists for Oracle Exadata Database Service on Cloud@Customer Deployments
  To determine your readiness for an Oracle Exadata Database Service on Cloud@Customer deployment, review the deployment checklists.

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

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

Gen 2 Exadata Cloud@Customer is managed by the Oracle Cloud Infrastructure (OCI) control plane. The Exadata Database Service on Cloud@Customer resources are deployed in your OCI Tenancy.

Before you can provision Exadata Database Service on Cloud@Customer infrastructure, your Oracle Cloud Infrastructure tenancy must be enabled to use Oracle Exadata Database Service on Cloud@Customer. 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

After Oracle creates your tenancy in OCI, an administrator at your company will need to perform some set up tasks and establish an organization plan for your cloud resources and users. The information in this topic will help you get started.

- Managing Regions
  This topic describes the basics of managing your region subscriptions.
- Managing Compartments
  This topic describes the basics of working with compartments.
- Managing Users
  This topic describes the basics of working with users.
- Managing Groups
  This topic describes the basics of working with groups.
- Required IAM Policy for Oracle Exadata Database Service on Cloud@Customer
  Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Cloud@Customer 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 Cloud@Customer

Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Cloud@Customer 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 administrator the type of access you've been granted, and which compartment you should work in.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

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

Related Topics

- Let database admins manage DB systems
- Getting Started with Policies
- Common Policies
- Details for the Database Service
- Policy Details for Oracle Exadata Database Service on Cloud@Customer
  Learn to write policies to control access to Oracle Exadata Database Service on Cloud@Customer resources.

Site Requirements for Oracle Exadata Database Service on Cloud@Customer

Review the requirements for provisioning Oracle Exadata Database Service on Cloud@Customer at your site.

- Space Requirements for Oracle Exadata Database Service on Cloud@Customer
  Review the space requirements for each Exadata Database Service on Cloud@Customer Rack.
- Weight of Oracle Exadata Database Service on Cloud@Customer Racks
  Review and prepare to manage the weight of each Exadata Database Service on Cloud@Customer Rack.
- Receiving, Unpacking, and Access for Oracle Exadata Database Service on Cloud@Customer Racks
  Review and prepare the receiving area that is large enough for the Exadata Rack package.
- Flooring for Oracle Exadata Database Service on Cloud@Customer Racks
  Ensure that the Exadata Database Service on Cloud@Customer system is installed on raised flooring that is capable of supporting the Exadata Rack.
- Electrical Power for Oracle Exadata Database Service on Cloud@Customer Racks
  Exadata Database Service on Cloud@Customer can operate effectively over a wide range of voltages and frequencies.
- Temperature and Humidity Ranges for Oracle Exadata Database Service on Cloud@Customer
  Excessive internal temperatures can result in full or partial shutdown of Exadata Database Service on Cloud@Customer system components.
- Ventilation for Oracle Exadata Database Service on Cloud@Customer Racks
  Always provide adequate space in front and behind the rack for proper ventilation.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Space Requirements for Oracle Exadata Database Service on Cloud@Customer

Review the space requirements for each Exadata Database Service on Cloud@Customer Rack.

Table 3-1 Space Requirements for Oracle Exadata

Description	Millimeters (mm)	Inches (”)
Height	2000 mm	78.74”
Width	601 mm	23.66"
Depth	1237 mm	48.7"

Weight of Oracle Exadata Database Service on Cloud@Customer Racks

Review and prepare to manage the weight of each Exadata Database Service on Cloud@Customer Rack.

Model	Kilograms (kg)	Pounds (lbs)
X11M Rack min	470 kg	1036 lbs
X11M Rack max	900 kg	1985 lbs
X11M Expansion Rack min	416 kg	917 lbs
X11M Expansion Rack max	964 kg	2125 lbs
X10M Base Rack min	454 kg	999 lbs
X10M Elastic Rack max	910 kg	2006 lbs
X10M Expansion Rack min	416 kg	917 lbs
X10M Expansion Rack max	964 kg	2125 lbs
X9M Base Rack	437 kg	963 lbs
X9M Full Rack	897 kg	1976 lbs
X9M Expansion Rack min	375 kg	826 lbs
X9M Expansion Rack max	905 kg	1995 lbs

For a specific configuration please refer to Oracle Exadata Configuration Assistant

Receiving, Unpacking, and Access for Oracle Exadata Database Service on Cloud@Customer Racks

Review and prepare the receiving area that is large enough for the Exadata Rack package.

Description	Millimeters (mm)	Inches (”)
Shipping Height	2159 mm	85 inches
Shipping Width	1219 mm	48 inches
Shipping Depth	1575 mm	62 inches

If your loading dock meets the height and ramp requirements for a standard freight carrier truck, then you can use a pallet jack to unload the rack. If the loading dock does not meet the

Copyright © 2008, 2026, Oracle and/or its affiliates.

requirements, then you must provide a standard forklift, or other means to unload the rack. You can also request that the rack is shipped in a truck with a lift gate.

Use a conditioned space to remove the packaging material to reduce particles before entering the data center. Allow enough space for unpacking it from its shipping cartons.

Use the information in the following table to ensure that there is a clear pathway for moving the Exadata Database Service on Cloud@Customer rack. Also, the entire access route to the installation site should be free of raised-pattern flooring that can cause vibration.

Access Route Item	With Shipping Pallet	Without Shipping Pallet
Minimum door height	2184 mm (86 inches)	2040 mm (80.32 inches)
Minimum door width	1270 (50 inches)	640 mm (25.19 inches)
Minimum elevator depth	1625.6 mm (64 inches)	1240 mm (48.82 inches)
Maximum incline	6 degrees	6 degrees
Minimum elevator, pallet jack, and floor loading capacity	1134 kg (2500 lbs)	1134 kg (2500 lbs)

Flooring for Oracle Exadata Database Service on Cloud@Customer Racks

Ensure that the Exadata Database Service on Cloud@Customer system is installed on raised flooring that is capable of supporting the Exadata Rack.

The site floor and the raised flooring must be able to support the total weight of the Exadata Database Service on Cloud@Customer rack that you have selected. Review specifications accordingly.

Electrical Power for Oracle Exadata Database Service on Cloud@Customer Racks

Exadata Database Service on Cloud@Customer can operate effectively over a wide range of voltages and frequencies.

Reliability of Power Sources

Each rack must have a reliable power source. Damage can occur if the voltage ranges are exceeded. Electrical disturbances such as the following can damage Exadata Database Service on Cloud@Customer:

- Fluctuations caused by brownouts
- Wide and rapid variations in input voltage levels or in input power frequency
- Electrical storms
- Faults in the distribution system, such as defective wiring

To protect Exadata Database Service on Cloud@Customer from such disturbances, you should have a dedicated power distribution system, power-conditioning equipment, and lightning arresters or power cables to protect from electrical storms.

Power Distribution Unit Specifications

Each rack has two pre-installed power distribution units (PDUs). The PDUs accept different power sources. You must choose the type of PDU that is correct for your data center and the Exadata Database Service on Cloud@Customer rack. For a specific configuration please refer to Oracle Exadata Configuration Assistant

Copyright © 2008, 2026, Oracle and/or its affiliates.

The following list outlines the available PDUs for Exadata Database Service on Cloud@Customer, depending on your region. Follow each of the links to access detailed specifications for each PDU type:

- Americas, Japan, and Taiwan
  - Low-Voltage 15 kVA Single-Phase
  - Low-Voltage 15 kVA Three-Phase
  - Low-Voltage 22 kVA Single-Phase
  - Low-Voltage 24 kVA Three-Phase
- Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan
  - High-Voltage 15 kVA Three-Phase
  - High-Voltage 22 kVA Single-Phase
  - High-Voltage 24 kVA Three-Phase

Facility Power Requirements

To prevent catastrophic failures, design the input power sources to ensure that adequate power is provided to the PDUs.

Use dedicated AC breaker panels for all power circuits that supply power to the PDU. When planning for power distribution requirements, balance the power load between available AC supply branch circuits. In the United States of America and Canada, ensure that the overall system AC input current load does not exceed 80 percent of the branch circuit AC current rating.

> **Note**
>
> Electrical work and installations must comply with applicable local, state, or national electrical codes.
>
> PDU power cords are 4 meters (13.12 feet) long, and 1–1.5 meters (3.3–4.9 feet) of the cord is routed within the rack cabinet. The installation site AC power receptacle must be within 2 meters (6.6 feet) of the rack.

Circuit Breaker Requirements

If computer equipment is subjected to repeated power interruptions and fluctuations, then it is susceptible to a higher rate of component failure.

You are responsible for supplying the circuit breakers. One circuit breaker is required for each power cord. In addition to circuit breakers, provide a stable power source, such as an uninterruptible power supply (UPS) to reduce the possibility of component failures.

Use dedicated AC breaker panels for all power circuits that supply power to the server. Servers require grounded electrical circuits.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Electrical work and installations must comply with applicable local, state, or national electrical codes.

Electrical Grounding Guidelines

The cabinets for Oracle Exadata Rack are shipped with grounding-type power cords.

- Always connect the cords to grounded power outlets.
- Check the grounding type, because different grounding methods can be used, depending on your location.
- Refer to documentation such as IEC documents for the correct grounding method.
- Ensure that the facility administrator or qualified electrical engineer verifies the grounding method for the building, and performs the grounding work.

Temperature and Humidity Ranges for Oracle Exadata Database Service on Cloud@Customer

Excessive internal temperatures can result in full or partial shutdown of Exadata Database Service on Cloud@Customer system components.

Temperature and Humidity Ranges

Condition	Operating Requirement	Non-operating Requirement	Optimal Requirement
Temperature	5–32 degrees Celsius (41–89.6 degrees Fahrenheit)	-40–70 degrees Celsius (-40–158 degrees Fahrenheit)	21–23 degrees Celsius (70–74 degrees Fahrenheit)
Relative Humidity	10–90 percent relative humidity, non-condensing	Up to 93 percent relative humidity	45–50 percent, non-condensing
Altitude	3048 meters (10000 feet) maximum	12,000 meters (40000 feet) maximum	Maximum ambient temperature is reduced by 1 degree Celsius for every 300 meters of altitude over 900 meters above sea level.

Temperature and Humidity Guidelines

To minimize the chance of downtime because of component failure, set conditions to the optimal temperature and humidity ranges. Maintaining an Exadata Database Service on Cloud@Customer system for extended periods at or near the operating limits can significantly increase the potential for hardware component failure.

The ambient temperature range of 21–23 degrees Celsius (70–74 degrees Fahrenheit) is optimal for server reliability and operator comfort. Most computer equipment can operate in a wide temperature range, but near 22 degrees Celsius (72 degrees Fahrenheit) is desirable because it is easier to maintain safe humidity levels. Operating in this temperature range provides a safety buffer in case the air conditioning system fails for some time.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The ambient relative humidity range of 45–50 percent is suitable for safe data processing operations. Most computer equipment can operate in a wide range (20–80 percent), but the range of 45–50 percent is recommended for the following reasons:

- The optimal range helps protect computer systems from corrosion problems associated with high humidity levels..
- The optimal range provides the greatest operating time buffer in case the air conditioning system fails for some time.
- The optimal range avoids failures or temporary malfunctions caused by interference from static discharges that can occur when relative humidity is too low. Electrostatic discharge (ESD) is easily generated, and hard to dissipate in areas of low relative humidity, such as below 35 percent. ESD becomes critical when humidity drops below 30 percent

Ventilation for Oracle Exadata Database Service on Cloud@Customer Racks

Always provide adequate space in front and behind the rack for proper ventilation.

Do not obstruct the front or rear of the rack with equipment or objects that might prevent air from flowing through the rack. Each Exadata Database Service on Cloud@Customer rack draws cool air in through the front of the rack, and discharges warm air out the rear of the rack. There is no air flow requirement for the left and right sides, because of front-to-back cooling.

Each Exadata Database Service on Cloud@Customer rack is designed to function while installed in a natural convection air flow. To ensure adequate air flow, allow a minimum clearance of 1219.2 mm (48 inches) at the front of the server, and 914 mm (36 inches) at the rear of the server for ventilation.

Use perforated tiles, approximately 400 CFM/tile, in front of the rack for cold air intake. The tiles can be arranged in any order in front of the rack, as long as cold air from the tiles can flow into the rack. Inadequate cold air flow could result in a higher inlet temperature in the servers because of exhaust air recirculation. The following is the recommended number of floor tiles:

- Four floor tiles for an Exadata Database Service on Cloud@Customer Full Rack.
- Three floor tiles for an Exadata Database Service on Cloud@Customer Half Rack.
- One floor tile for an Exadata Database Service on Cloud@Customer Quarter Rack or Base System.

Network Requirements for Oracle Exadata Database Service on Cloud@Customer

Review the network requirements for provisioning Oracle Exadata Database Service on Cloud@Customer at your site.

- Network Requirements for Oracle Exadata Database Service on Cloud@Customer
  To provide secure and reliable network connectivity for different application and management functions, Exadata Database Service on Cloud@Customer uses different networks.
- Data Center Network Services for Oracle Exadata Database Service on Cloud@Customer
  Before you deploy Oracle Exadata Database Service on Cloud@Customer, ensure that your data center network meets requirements.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- IP Addresses and Subnets for Oracle Exadata Database Service on Cloud@Customer
  You must allocate a range of IP addresses to the administration network, and another range of IP addresses to the RoCE private network.
- Uplinks for Oracle Exadata Database Service on Cloud@Customer
  Ensure that your Oracle Exadata Database Service on Cloud@Customer system meets control plane server and database server uplink requirements.
- Network Cabling for Oracle Exadata Database Service on Cloud@Customer
  You can choose to use the supplied network equipment, or you can build your own SFP network.
- Establish a Secure Connection Between Your CPS and OCI Using OCI's FastConnect Service
  Consider the solution outlined below, which leverages OCI's FastConnect service, if you want additional isolation for the connection between your CPS and OCI in addition to the default TLS tunnel approach.

Network Requirements for Oracle Exadata Database Service on Cloud@Customer

To provide secure and reliable network connectivity for different application and management functions, Exadata Database Service on Cloud@Customer uses different networks.

The following list outlines the minimum network requirements to install an Exadata Database Service on Cloud@Customer system:

- **Exadata Cloud@Customer Service Network:** These network will be set up to Oracle specifications and should not be modified by customer without Oracle agreement.
  - **Control Plane Network**
    This virtual private network (VPN) connects the two control plane servers that are located in the Exadata Database Service on Cloud@Customer rack to Oracle Cloud Infrastructure. The VPN facilitates secure customer-initiated operations using the Oracle Cloud Infrastructure Console and APIs. It also facilitates secure monitoring and administration of the Oracle-managed infrastructure components in Exadata Database Service on Cloud@Customer.
    * **Control Plane Connectivity Considerations**
      In order for the control plane to function, the control plane server must be able to connect to certain Oracle Cloud Infrastructure (OCI) addresses. You must enable TCP port 443 outbound access to the endpoints in a specific OCI region as follows:

Table 3-2 Ports to Open for Control Plane Connectivity

Description / Purpose	Open Port	Location
Outgoing Tunnel Service for Cloud Automation delivery	443 outbound	Use this URL format, replacing *oci_region* with your region: `https://wss.exacc.`*`oci_region`*`.oci.oraclecloud.com`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 3-2 (Cont.) Ports to Open for Control Plane Connectivity

Description / Purpose	Open Port	Location

Secure Tunnel Service for remote Oracle operator access	443 outbound	Use these URL formats, replacing *oci_region* with your region:
		
		`https://mgmthe1.exacc.oci_region.oci.oraclecloud.com`
		
		`https://mgmthe2.exacc.oci_region.oci.oraclecloud.com`

Object Storage Service to retrieve system updates, infrastructure monitoring, and log collection	443 outbound	Use this URL format, replacing *oci_region* with your region:
		
		`https://objectstorage.oci_region.oraclecloud.com`
		
		`https://swiftobjectstorage.oci_region.oraclecloud.com`
		
		`https://*.objectstorage.oci_region.oci.customer-oci.com`

Monitoring Service to record and process Infrastructure Monitoring Metrics (IMM)	443 outbound	Use this URL format, replacing *oci_region* with your region: `https://telemetry-ingestion.oci_region.oraclecloud.com`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 3-2 (Cont.) Ports to Open for Control Plane Connectivity

Description / Purpose	Open Port	Location

Identity Service for Authorization and Authentication	443 outbound	Use this URL format, replacing *oci_region* with your region:
		
		`https://identity.oci_region.oraclecloud.com`
		
		`https://auth.oci_region.oraclecloud.com`

Outgoing Tunnel Service for Cloud Automation delivery	443 outbound	Use this URL format, replacing *oci_region* with your region: `https://wsshe.adbd-exacc.oci_region.oci.oraclecloud.com`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 3-2 (Cont.) Ports to Open for Control Plane Connectivity

Description / Purpose	Open Port	Location

Logging Service	443 outbound	Use this URL format replacing *oci_region* with your region:
		
		`https://frontend.logging.ad1.`*`oci_region`*`.oracleiaas.com`
		
		`https://frontend.logging.ad2.`*`oci_region`*`.oracleiaas.com`
		
		`https://frontend.logging.ad3.`*`oci_region`*`.oracleiaas.com`
		
		`https://controlplane.logging.ad1.`*`oci_region`*`.oracleiaas.com`
		
		`https://controlplane.logging.ad2.`*`oci_region`*`.oracleiaas.com`
		
		`https://controlplane.logging.ad3.`*`oci_region`*`.oracleiaas.com`

Resource Principal based authentication and Autonomous AI Database service delivery	443 outbound	Use this URL format, replacing *oci_region* with your region: `https://database.`*`oci_region`*`.oraclecloud.com`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 3-2 (Cont.) Ports to Open for Control Plane Connectivity

Description / Purpose	Open Port	Location

VM console	443 outbound	Use this URL format, replacing *oci_region* with your region:
		
		`https://console1.exacc.oci_region.oci.oraclecloud.com`
		
		`https://console2.exacc.oci_region.oci.oraclecloud.com`

Temporary Secure Tunnel Service for remote Oracle operator access for ADB-D resources	443 outbound	Use this URL format, replacing *region* with your region: `https://mgmthe.adbd-exacc.region.oci.oraclecloud.com`
Monitoring Service to record and process Infrastructure Monitoring Metrics (IMM) resources	443 outbound	Use this URL format, replacing *region* with your region: `https://ingestion.logging.region.oci.oraclecloud.com`

Metering and Monitoring	443 outbound	Use this URL format, replacing `oci_region` with your region:
		
		`https://*.oci_region.functions.oci.oraclecloud.com`
		
		`https://*.functions.oci_region.oci.oraclecloud.com`

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> ExaDB-C@C infrastructure access to the stated service endpoints in Table 3-2 Ports to Open for Control Plane Connectivity is mandatory for full service functionality. Failing to permit all the mandatory URIs may result in a service reduction or features not working as designed.

Note that the Control Plane Server must be able to establish TCP Port 443 outbound access only. While outbound connections on Port 443 must be allowed, TCP Port 443 inbound access is not required, and it may be desirable from a security standpoint to block inbound connections. (Functionally, bi-directional traffic is still possible over the connection once the secure outbound connection is established.)

The Control Plane Server requires customer DNS and NTP services to be functional. Minimum bandwidth requirements for the Control Plane Server internet connection to OCI are 50/10 mbs download/upload.

Some customers have security policies requiring the use of proxies for all internet connections to IT infrastructure. The customer should use an HTTP Open Proxy server, which is a passive/corporate proxy supported for the Control Plane Server connection to OCI. The customer should not use a Reverse Proxy server, or a challenge proxy, and not perform traffic inspection as these are not supported.

If you are using IP address filtering based firewall rules, due to the dynamic nature of cloud interfaces, you must allow traffic with all the relevant IP CIDR ranges associated with your OCI region as identified by https://docs.oracle.com/en-us/iaas/tools/public_ip_ranges.json.

- **Administration Network**
  This network connects Exadata Database Service on Cloud@Customer servers and switches to the two control plane servers that are located in the Exadata Database Service on Cloud@Customer rack. It facilitates customer-initiated operations using the Oracle Cloud Infrastructure Console and APIs. It also facilitates monitoring and administration of the Oracle-managed infrastructure components in Exadata Database Service on Cloud@Customer.

  This network is fully contained within the Exadata Database Service on Cloud@Customer rack, and does not connect to your corporate network. However, the Exadata infrastructure is indirectly connected to your corporate network through the control plane servers. This connection is required to provide Domain Name System (DNS) and Network Time Protocol (NTP) services to the Exadata infrastructure. Therefore, the IP addresses that are allocated to the administration network must not exist elsewhere in your corporate network.

  Each Oracle Database server and Exadata Storage Server has two network interfaces connected to the administration network. One provides management access to the server through one of the embedded Ethernet ports (`NET0`). The other provides access to the Integrated Lights-Out Management (ILOM) subsystem through a dedicated ILOM Ethernet port. Exadata Database Service on Cloud@Customer is delivered with the ILOM and `NET0` ports connected to the Ethernet switch in the rack. Cabling or configuration changes to these interfaces are not permitted.

- **InfiniBand or RDMA Over Converged Ethernet (ROCE) Network**
  This network connects the database servers, Exadata Storage Servers, and control plane servers using the InfiniBand or ROCE switches on the rack. Each server contains two InfiniBand network interfaces (`IB0` and `IB1`) or ROCE interface (`re0` and

Copyright © 2008, 2026, Oracle and/or its affiliates.

re1) that are connected to separate InfiniBand or ROCE switches in the rack. Primarily, Oracle Database uses this network for Oracle RAC cluster interconnect traffic, and for accessing data on Exadata Storage Servers.

This non-routable network is fully contained within the Exadata Cloud@Customer rack, and does not connect to your corporate network. However, because the control plane servers are connected to the InfiniBand or ROCE network and to your corporate network, the IP addresses that are allocated to the InfiniBand or ROCE network must not exist elsewhere in your corporate network.

- **Customer Network:** Customer-owned and managed networks required for the Exadata Cloud@Customer data plane to access related systems.
  - **Client Network**
    This network connects the Exadata Cloud@Customer database servers to your existing client network and is used for client access to the virtual machines. Applications access databases on Exadata Database Service on Cloud@Customer through this network using Single Client Access Name (SCAN) and Oracle Real Application Clusters (Oracle RAC) Virtual IP (VIP) interfaces.

    The client access network uses a pair of network interfaces on each database server, which are connected to the customer network.

    > **Note**
    >
    > When you enable Data Guard, replication of data happens over the client network by default.

  - **Backup Network**
    This network is similar to the client access network, as it connects the Exadata Database Service on Cloud@Customer Oracle Database servers to your existing network. It can be used for access to the virtual machines for various purposes, including backups and bulk data transfers.

    Like the client network, the backup network uses a pair of network interfaces on each database server, which are connected to the customer network. Physically connecting the backup network to a customer network is required.

    If the customer's on-premises storage (NFS or ZDLRA) is to be used exclusively as a backup destination for databases, then no external connectivity to OCI is required for the backup network.

    Exadata Cloud@Customer also offers an Oracle-managed cloud object storage backup destination. If Oracle's Object Storage Service is to be leveraged as a backup destination for database backups, then ensure that the backup network can reach the Object Storage Service through external connection. You must enable TCP port 443 outbound access for the backup network as follows:

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 3-3 Ports to Open for Backup Network

Description / Purpose	Open Port	Location

Object Storage Service for cloud-based database backups (Optional)	443 outbound	Use this URL format, replacing *oci_region* with your region:
		`https://objectstorage.oci_region.oraclecloud.com`
		`https://swiftobjectstorage.oci_region.oraclecloud.com`

- **Disaster Recovery Network (authorized customers only)**

  > **Note**
  >
  > This additional network is available only to certain authorized customers under special circumstances. Your system may not be equipped with it.

  The Disaster Recovery network is configured similarly to the client access network but is dedicated solely to transmitting Data Guard traffic. If your system includes this network, Data Guard traffic will be routed through it instead of the client network, utilizing Single Client Access Name (SCAN) and Oracle Real Application Clusters (Oracle RAC) Virtual IP (VIP) interfaces. If you encounter any issues, see Data Guard Association Creation Fails When Using DR Network.

  Only the Data Guard prechecks require the client network to be open. Once the precheck stage is complete, Data Guard provisioning automatically registers the database with the Disaster Recovery network and uses it for both datafile restoration and redo transport. If opening the client network is not an option, you must manually register the primary database with the Disaster Recovery network listener before initiating Data Guard provisioning.

  The Disaster Recovery network consists of a pair of network interfaces on each database server, which are directly connected to the customer network.

Related Topics

- Object Storage Service API
- Monitoring API
- Identity and Access Management Service API

Copyright © 2008, 2026, Oracle and/or its affiliates.

Data Center Network Services for Oracle Exadata Database Service on Cloud@Customer

Before you deploy Oracle Exadata Database Service on Cloud@Customer, ensure that your data center network meets requirements.

Domain Name System (DNS) Requirements

As part of the deployment process, you must decide on the host names and IP addresses to be used for various Oracle Exadata Database Service on Cloud@Customer network interfaces. Oracle requires that you register the host names and IP addresses for the Oracle Exadata Database Service on Cloud@Customer client and backup network interfaces in your corporate DNS. At least one reliable DNS server is required, which must be accessible to the control plane servers and to all of the servers on the client network. Up to three DNS servers can be registered in Oracle Exadata Database Service on Cloud@Customer to ensure coverage in case a server is unavailable.

Network Time Protocol (NTP) Services Requirements

Oracle Exadata Database Service on Cloud@Customer uses NTP to ensure that all system components are synchronized to the same time. At least one reliable NTP server is required, which must be accessible to the control plane servers and to all of the servers on the client network. Up to three NTP servers can be registered in Oracle Exadata Database Service on Cloud@Customer to ensure coverage in case a server is unavailable.

IP Addresses and Subnets for Oracle Exadata Database Service on Cloud@Customer

You must allocate a range of IP addresses to the administration network, and another range of IP addresses to the RoCE private network.

Administration and Private Network Requirements for Oracle Exadata Database Service on Cloud@Customer

No overlap is permitted between the address ranges for the administration network and the RoCE private network, and all IP addresses should be unique within your corporate network. You must also allocate IP addresses from your corporate network to the control plane servers. These network configuration details are specified when you create the Exadata infrastructure.

When you create the Exadata infrastructure, the console pre-populates default values for the administration network CIDR block and the InifinBand network CIDR block. You can use the suggested CIDR blocks if there is no overlap with existing IP addresses in your corporate network.

Review IP address requirements for each of these networks. The table specifies the maximum and minimum CIDR block prefix length that are allowed for each network. The maximum CIDR block prefix length defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Oracle Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Network Type	IP Address Requirements

Administration network	Maximum CIDR block prefix length:
	`/23`
	Minimum CIDR block prefix length:
	`/16`

Private network	Maximum CIDR block prefix length:
	`/22`
	Minimum CIDR block prefix length:
	`/19`

Control plane network	2 IP addresses, 1 for each control plane server

For more information about Administration and Private Network CIDR requirements, see *Table 5-5 X9M CIDR Requirements* and *Table 5-6 X8M CIDR Requirements* in Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure.

Host Name and IP Address Requirements for Oracle Exadata Database Service on Cloud@Customer

To connect to your corporate network, Oracle Exadata Database Service on Cloud@Customer requires several host names and IP addresses for network interfaces on the client network and the backup network. The precise number of IP addresses depends on the Exadata system shape. These network configuration details, including host names and IP addresses, are specified when you create a VM cluster network. All IP addresses must be statically assigned IP addresses, not dynamically assigned (DHCP) addresses. The client network and the backup network require separate subnets.

The following table outlines the IP address requirements for the client and backup networks. The table specifies the maximum and recommended CIDR block prefix length for each network. The maximum CIDR block prefix length defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Oracle Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Network Type	IP Address Requirements for Base System, Quarter Rack, or Half Rack	IP Address Requirements for Full Rack
Client network	Maximum CIDR block prefix length: /28 Recommended CIDR block prefix length: /27	Maximum CIDR block prefix length: /27 Recommended CIDR block prefix length: /26
Backup network	Maximum CIDR block prefix length: /29 Recommended CIDR block prefix length: /28	Maximum CIDR block prefix length: /28 Recommended CIDR block prefix length: /27
(For authorized customers only) Disaster Recovery network	Maximum CIDR block prefix length: /28 Recommended CIDR block prefix length: /27	Maximum CIDR block prefix length: /27 Recommended CIDR block prefix length: /26

Uplinks for Oracle Exadata Database Service on Cloud@Customer

Ensure that your Oracle Exadata Database Service on Cloud@Customer system meets control plane server and database server uplink requirements.

Control Plane Servers

Four uplinks (2 per server) are required to connect the Control Plane Servers to your corporate network to support the outbound secure network connections to the OCI services required by Oracle Exadata Database Service on Cloud@Customer.

Database Server Connections

Typically, four uplinks are required for each database server to connect to your corporate network. Using this configuration, two uplinks support the client network and the other two uplinks support the backup network.

Copyright © 2008, 2026, Oracle and/or its affiliates.

On Quarter Rack, Half Rack, or Full Rack systems, you can choose to use 10 Gbps RJ45 copper, 10 Gbps SFP+ fiber, or 25Gbps SFP28 fiber network connections to your corporate network. However, you cannot mix copper and fiber networks on the same physical server. For example, you cannot use fiber for the client network and copper for the backup network.

On Base System configurations, the options are more limited because of the physical network interfaces that are available on each database server. On Base Systems, you can choose to use copper or fiber network connections only for the client network, while the backup network uses a fiber connection.

You can also use shared network interfaces on the Base System for the client network and the backup network, which reduces the uplink requirement to two uplinks for each database server. Using shared network interfaces also enables you to use copper network connections to support both the client and backup networks on Base System configurations. However, in general, Oracle recommends that you do not use shared network interfaces, because sharing networks compromises the bandwidth and availability of both networks. Shared network interfaces are not supported for Quarter, Half, and Full Rack configurations.

Network Cabling for Oracle Exadata Database Service on Cloud@Customer

You can choose to use the supplied network equipment, or you can build your own SFP network.

Supplied Network Equipment Option

Every Oracle Exadata Database Service on Cloud@Customer rack is shipped with all of the network equipment and cables that are required to interconnect all hardware in the Oracle Exadata Database Service on Cloud@Customer rack.

Small Form-Factory Pluggable Network Option

Oracle supplies small form-factor pluggable (SFP) network interfaces to enable connectivity to your corporate network. However, if you choose to configure an SFP network, then you are responsible to provide the required cabling to connect the Exadata Database Servers and Control Plane Servers to your corporate network.

Establish a Secure Connection Between Your CPS and OCI Using OCI's FastConnect Service

Consider the solution outlined below, which leverages OCI's FastConnect service, if you want additional isolation for the connection between your CPS and OCI in addition to the default TLS tunnel approach.

For more information, see *Oracle Cloud Infrastructure FastConnect*.

Oracle Exadata Database Service on Cloud@Customer service supports the public or private peering connectivity model of FastConnect.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Figure 3-1 Oracle Exadata Database Service on Cloud@Customer FastConnect connectivity to OCI through public peering

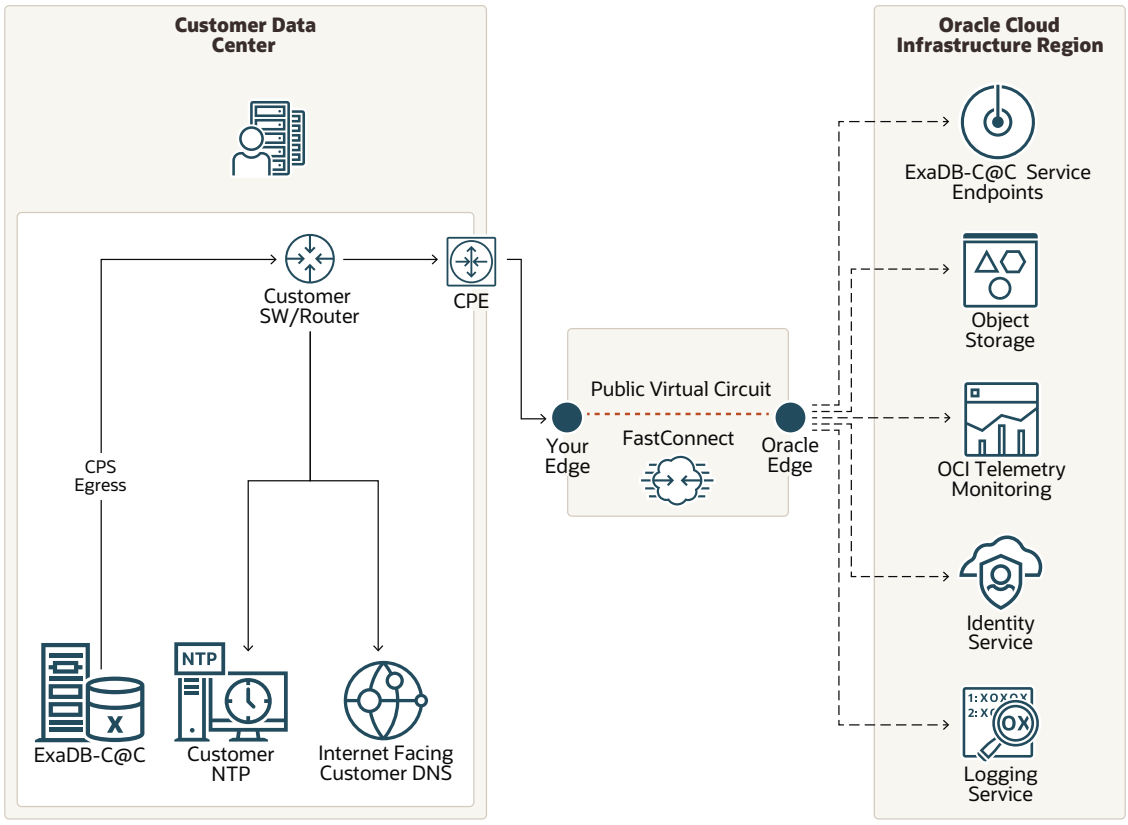

As shown in the figure, the Oracle Exadata Database Service on Cloud@Customer Control Plane network egresses to FastConnect provider and to Oracle edge. Customers who may already have existing FastConnect connectivity can use it to connect Oracle Exadata Database Service on Cloud@Customer to the OCI region using public peering.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Figure 3-2 Oracle Exadata Database Service on Cloud@Customer FastConnect connectivity to OCI through private peering

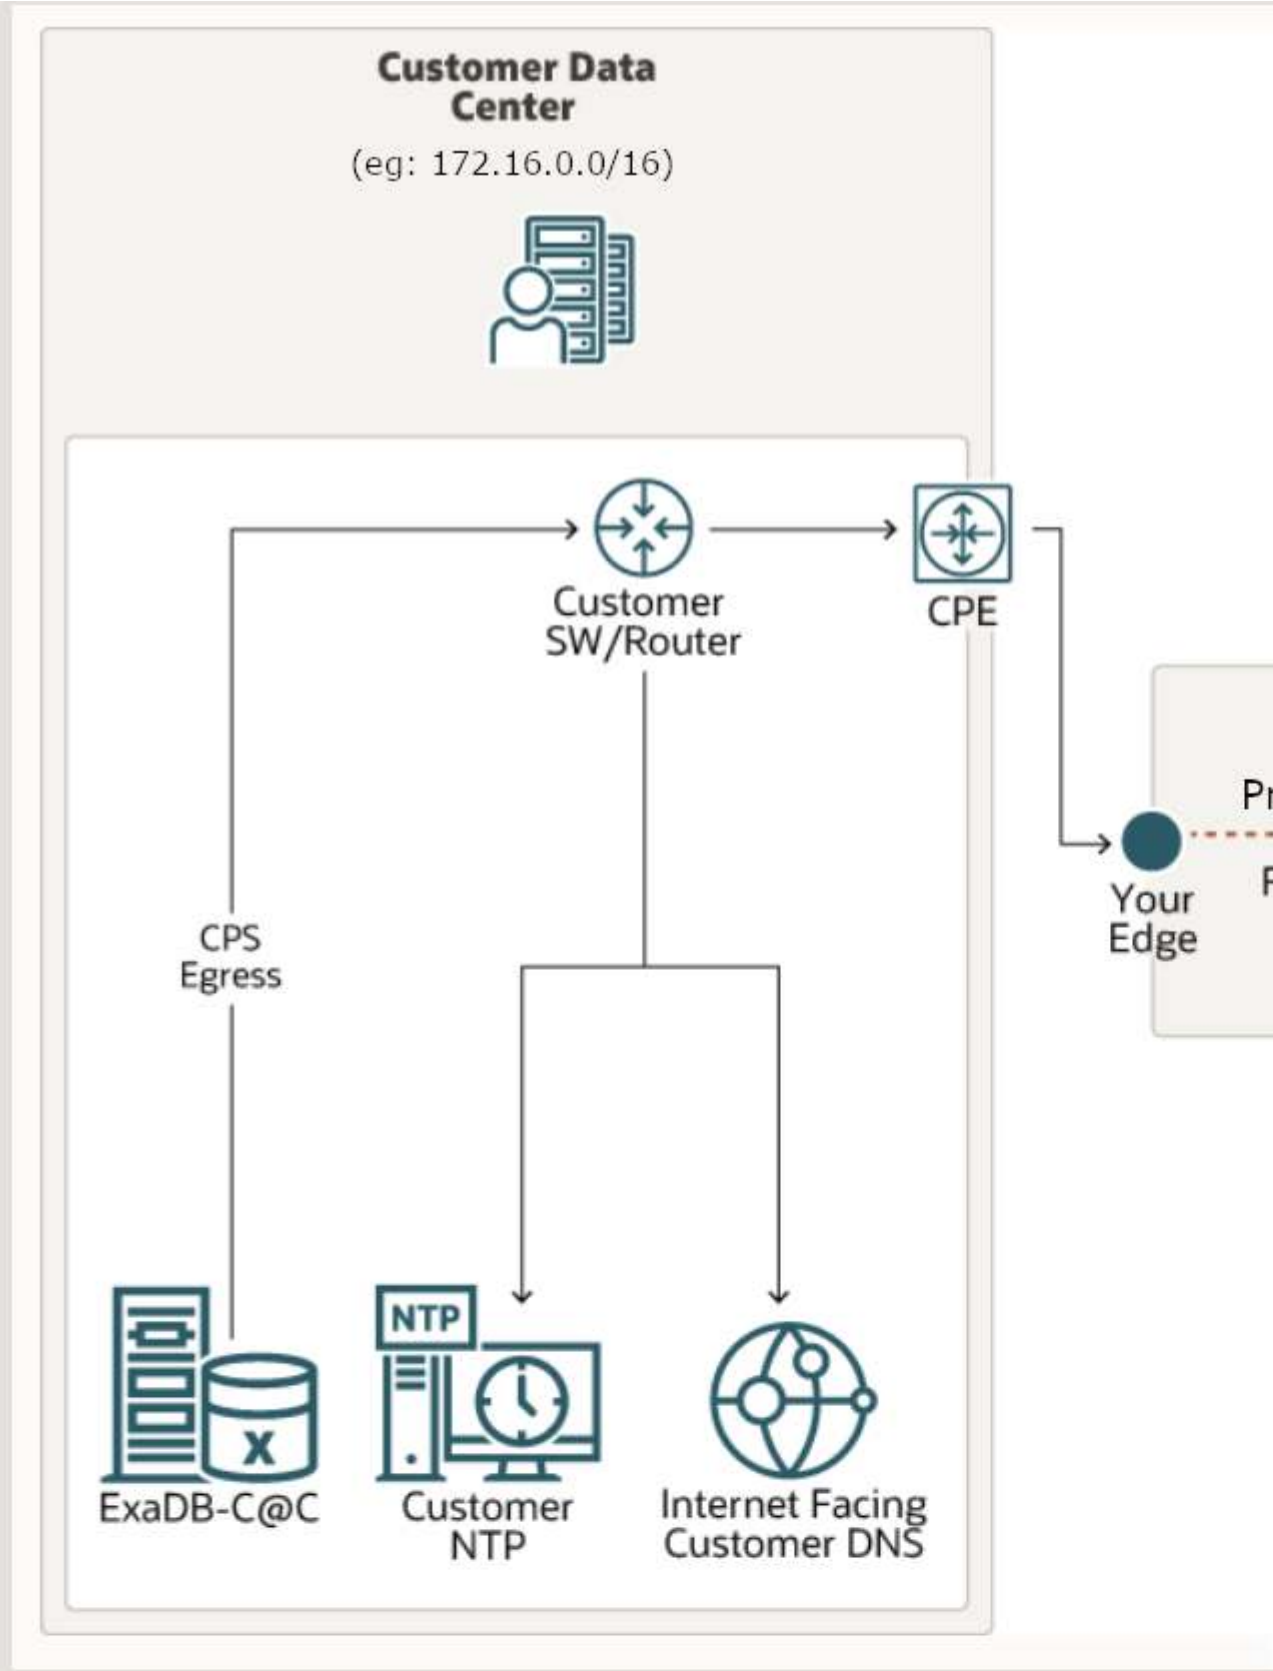

Copyright © 2008, 2026, Oracle and/or its affiliates.

Configuring FastConnect for Oracle Exadata Database Service on Cloud@Customer

You can set up and configure Oracle Cloud Infrastructure FastConnect either before or after deploying Oracle Exadata Database Service on Cloud@Customer.

Oracle Exadata Database Service on Cloud@Customer rack egress network configuration for FastConnect

- As shown in the figures, it is important to set the Oracle Exadata Database Service on Cloud@Customer Control Plane Server network egress rules to forward traffic over FastConnect, and also have a route to an internet-facing customer DNS. This "Customer DNS" is used by Oracle Exadata Database Service on Cloud@Customer infrastructure to resolve OCI public endpoints.
- Corporate HTTP proxy between Oracle Exadata Database Service on Cloud@Customer Control Plane Servers and OCI region is not recommended with FastConnect as it is already a dedicated network. If a corporate proxy is highly desired, then the proxy needs to have additional routing to ensure that Oracle Exadata Database Service on Cloud@Customer network traffic is sent over FastConnect.
- If you are using private peering, then ensure that you configure transit routing at your VCN side. For more information, see *Transit Routing to the Oracle Services Network*.

Related Topics

- Oracle Cloud Infrastructure FastConnect
- Transit Routing to the Oracle Services Network

Storage Configuration Requirements for Oracle Exadata Database Service on Cloud@Customer

With the introduction of Exascale technology in Exadata Database Service on Cloud@Customer, you can configure the Exadata infrastructure to use ASM, Exascale, or a combination of both. When configured to coexist with ASM, you must allocate Exascale storage capacity from the total Exadata storage available on the Infrastructure. The allocated storage capacity will be solely available for Exascale use.

Review the storage requirements for ASM, Exascale, and VM file systems to plan and optimize storage allocation based on your enterprise needs.

> **Note**
>
> If you are looking for information to plan the local VM storage available on your VM Clusters refer to Estimating How Much Local Storage You Can Provision On Your VMs

- Configuring Exadata ASM Storage
- Configuring Exadata Exascale Storage

Configuring Exadata ASM Storage

Copyright © 2008, 2026, Oracle and/or its affiliates.

- About ASM Storage Configuration for Oracle Exadata Database Service on Cloud@Customer
  As part of configuring each Oracle Exadata Database Service on Cloud@Customer VM cluster, the storage space inside the Exadata Storage Servers is configured for use by Oracle Automatic Storage Management (ASM).
- Allocation of Storage Space Options on Oracle Exadata Storage Servers
  Select the storage option that best meets your planned use case on your Oracle Exadata Storage Servers.
- Allocation Proportions for DATA, RECO and SPARSE Disk Groups
  Determine the storage allocation between the `DATA`, `RECO`, and `SPARSE` disk groups for Oracle Exadata Storage Servers.
- Virtual Machine File System Structure for Oracle Exadata Database Service on Cloud@Customer

About ASM Storage Configuration for Oracle Exadata Database Service on Cloud@Customer

As part of configuring each Oracle Exadata Database Service on Cloud@Customer VM cluster, the storage space inside the Exadata Storage Servers is configured for use by Oracle Automatic Storage Management (ASM).

By default, the following ASM disk groups are created:

- The `DATA` disk group is primarily intended for the storage of Oracle Database data files. Also, a small amount of space is allocated from the `DATA` disk group to support the shared file systems that are used to store software binaries (and patches) and files associated with the cloud-specific tooling. You should not store your own data, including Oracle Database data files, backups, trace files, and so on, inside the system-related ACFS file systems.
- The `RECO` disk group is primarily used for storing the Fast Recovery Area (FRA), which can be used to provide a local store for files related to backup and recovery. By default, the FRA is used to store archived redo log files and the backup control file. If you configure your VM cluster with the option to allocate storage for local backups, then you can use the FRA as a database backup destination. Finally, if you enable flashback features on a database, then the FRA is used to store the flashback logs.

In addition, you can choose to create the `SPARSE` disk group. The `SPARSE` disk group is required to support Exadata snapshot functionality. Exadata snapshots enable space-efficient clones of Oracle databases that can be created and destroyed very quickly and easily. Snapshot clones are often used for development, testing, or other purposes that require a transient database. For more information about Exadata snapshot functionality, see *Setting Up Oracle Exadata Storage Snapshots* in the *Oracle Exadata System Software User's Guide*.

Related Topics

- Setting Up Oracle Exadata Storage Snapshots

Allocation of Storage Space Options on Oracle Exadata Storage Servers

Select the storage option that best meets your planned use case on your Oracle Exadata Storage Servers.

As an input to the virtual machine (VM) cluster creation process, you must choose options that determine how storage space in the Oracle Exadata Storage Servers is allocated to the Oracle

Copyright © 2008, 2026, Oracle and/or its affiliates.

ASM disk groups. Your choices profoundly affect how storage space in the Exadata Storage Servers is allocated to the ASM disk groups. Consider which option best meets your needs:

- Allocate Storage for Exadata Snapshots

  If you select this option, then the SPARSE disk group is created, and less space is allocated to the DATA and RECO disk groups. If you do not select this option, then the SPARSE disk group is not created, and you cannot use Exadata snapshot functionality.

- Allocate Storage for Local Backups

  If you select this option, then more space is allocated to the RECO disk group to accommodate local backups to Oracle Exadata storage. If you do not select this option, then more space is allocated to the DATA disk group, but you cannot use local Oracle Exadata storage as a backup destination for any databases in the VM cluster.

Allocation Proportions for DATA, RECO and SPARSE Disk Groups

Determine the storage allocation between the DATA, RECO, and SPARSE disk groups for Oracle Exadata Storage Servers.

Exadata Storage Server Configuration Allocation With No Exadata Snapshot Storage or Local Backup

When you select **Allocate Storage for Exadata Snapshots: No** and **Enable Backups on Local Exadata Storage: No**, then storage allocation is as follows:

- DATA Disk Group: 80%
- RECO Disk Group: 20%
- SPARSE Disk Group 0% (The SPARSE disk group is not created.)

Exadata Storage Server Configuration Allocation With No Exadata Snapshot Storage and Local Backup Enabled

When you select **Allocate Storage for Exadata Snapshots: No** and **Enable Backups on Local Exadata Storage: Yes**, so that backups are enabled on local storage, then storage allocation is as follows:

- DATA Disk Group: 40%
- RECO Disk Group: 60%
- SPARSE Disk Group 0% (The SPARSE disk group is not created.)

Exadata Storage Server Configuration Allocation With Exadata Snapshot Storage and No Local Backup

When you select **Allocate Storage for Exadata Snapshots: Yes** and **Enable Backups on Local Exadata Storage: No**, so that storage is allocated for Exadata snapshots, then storage allocation is as follows:

- DATA Disk Group: 60%
- RECO Disk Group: 20%
- SPARSE Disk Group 20%

Copyright © 2008, 2026, Oracle and/or its affiliates.

Exadata Storage Server Configuration Allocation With Both Exadata Snapshot Storage and Local Backup Enabled

When you select **Allocate Storage for Exadata Snapshots: Yes** and **Enable Backups on Local Exadata Storage: Yes**, so that storage is allocated for Exadata snapshots, and storage is allocated for local backups, then storage allocation is as follows:

- `DATA` Disk Group: 35%
- `RECO` Disk Group: 50%
- `SPARSE` Disk Group 15%

Virtual Machine File System Structure for Oracle Exadata Database Service on Cloud@Customer

Exadata Cloud@Customer use the following file system organization on the virtual machines. To plan for local storage allocation on the virtual machines please refer to *Estimating How Much Local Storage You Can Provision to Your VMs*.

Filesystem	Mounted On
`devtmpfs`	`/dev`
`tmpfs`	`/dev/shm`
`tmpfs`	`/run`
`tmpfs`	`/sys/fs/cgroup`
`tmpfs`	`/run/user/0`
`/dev/mapper/VGExaDb-LVDbSys1`	`/`
`/dev/mapper/VGExaDb-LVDbOra1`	`/u01`
`/dev/mapper/VGExaDb-LVDbTmp`	`/tmp`
`/dev/mapper/VGExaDb-LVDbVar1`	`/var`
`/dev/mapper/VGExaDb-LVDbVarLog`	`/var/log`
`/dev/mapper/VGExaDb-LVDbHome`	`/home`
`/dev/mapper/VGExaDbDisk.u02_extra.img-LVDBDisk`	`/u02`
`/dev/mapper/VGExaDb-LVDbVarLogAudit`	`/var/log/audit`
`/dev/sda1`	`/boot`
`/dev/mapper/VGExaDbDisk.grid19.0.0.0.200414.img-LVDBDisk`	`/u01/app/19.0.0.0/grid`
`/dev/asm/acfsvol01-142`	`/acfs01`

Related Topics

- Estimating How Much Local Storage You Can Provision On Your VMs

Configuring Exadata Exascale Storage

Copyright © 2008, 2026, Oracle and/or its affiliates.

- About Exascale Storage Configuration Oracle Exadata Database Service on Cloud@Customer
  To use Exascale storage during VM cluster provisioning, you must first configure it from the total Exadata storage available on your Exadata Infrastructure.
- Using the Console to Configure Exascale Storage on Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To configure Exascale storage on your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the infrastructure configuration.
- Using the Console to Create an Exascale Storage Vault
  To create an Exascale storage vault, be prepared to provide values for the Exascale storage vault configuration.
- Using the Console to Scale an Exascale Storage Vault
  To scale an Exascale storage vault, be prepared to provide values for the Exascale storage vault configuration.
- Using the Console to Move an Exascale Storage Vault to Another Compartment
  To move an Exascale storage vault to another compartment, use this procedure.
- Using the Console to Delete an Exascale Storage Vault
  To delete an Exascale storage vault, use this procedure.
- Using the Console to Configure Exascale VM Storage on Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To configure Exascale VM storage on your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the infrastructure configuration.
- Configure All Future VM Backups to Use Exascale Storage on a VM Cluster
  If a VM Cluster (ASM or Exascale) is using local database server storage for Oracle-managed VM backups, This option allows you to migrate all future VM backups to Exascale storage.
- Create an Advanced Cluster File System (ACFS) Using Exascale Block Volume
  Before you begin, ensure that the VM cluster is configured to use Exascale for database storage and that sufficient storage capacity is available in the Exascale storage vault associated with the cluster.
- Deconfigure Exascale Storage on the Exascale Infrastructure
  Before you begin, ensure that the VM cluster is configured to use Exascale for database storage and that sufficient storage capacity is available in the Exascale storage vault associated with the cluster.

About Exascale Storage Configuration Oracle Exadata Database Service on Cloud@Customer

To use Exascale storage during VM cluster provisioning, you must first configure it from the total Exadata storage available on your Exadata Infrastructure.

> **Note**
>
> During the initial Exascale storage configuration, all database servers on the Exadata Infrastructure will be restarted in a rolling reboot manner.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Specify the storage capacity you want to allocate for Exascale usage on the infrastructure. A minimum of 2TB of Exadata storage must be available to configure Exascale storage. Configuring Exascale storage requires a rolling reboot of the data nodes. After the initial configuration, you can scale the Exascale storage capacity online as needed. Note that scaling storage may trigger a storage rebalancing process, which could temporarily affect performance during the operation.

If sufficient storage is not available on the Infrastructure to configure the Exascale storage or to scale the existing Exascale storage capacity, additional storage servers can be added to the Infrastructure using the Scale Infrastructure option as follows:

- Add new storage servers to the infrastructure.
- Make the new storage capacity available to all ASM and Exascale VM clusters.
- Adding new storage will proportionally increase the Exascale storage capacity.
- Expand Exascale storage vaults as needed to utilize the additional capacity.

Alternatively, customers can shrink the existing ASM or Exascale clusters to free up storage and allocate it to Exascale to meet the additional capacity demand.

Using the Console to Configure Exascale Storage on Oracle Exadata Database Service on Cloud@Customer Infrastructure

To configure Exascale storage on your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the infrastructure configuration.

> **Note**
>
> - If an Exadata Infrastructure has been updated to Exadata system software release 25.1.9 (and later), or 25.2.2 (and later), a banner with the message "`You can now configure Exascale storage on this infrastructure`" will be displayed to indicate that the infrastructure is ready for Exascale configuration. This banner will not be displayed for infrastructures that are already configured to use Exascale.
> - The minimum Exascale storage that can be configured on Exadata Infrastructure is 2 TB.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where you want to configure Exascale storage.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Actions**, and then select **Configure Exascale storage**.
6. On the resulting Configure Exascale storage window, enter the storage capacity you want to allocate for Exascale.
7. Click **Submit**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The Exadata storage section on the Exadata Infrastructure Details page displays the storage details allocated for ASM and Exascale.

Using the Console to Create an Exascale Storage Vault

To create an Exascale storage vault, be prepared to provide values for the Exascale storage vault configuration.

An Exascale vault is a logical storage container that uses the physical resources provided by Exascale storage pools. Each vault is associated with at least one storage pool.

> **Note**
>
> - The minimum configurable or resizable size for an Exascale Storage Vault is 2 TB.
> - To successfully create a vault, ensure that the storage pool size exceeds the requested vault size. A portion of the storage pool is reserved for system overhead and is not available for vault allocation. For example, creating a 2 TB vault requires a storage pool larger than 2 TB (typically at least 2.2 TB).
> - A vault can be deleted and moved to another compartment.
> - A vault can be shared between two or more Exascale VM clusters.

> **Warning**
>
> Attempting to create a vault that is the same size as the total storage pool (for example, creating a 2 TB vault on a 2 TB storage pool) will fail due to insufficient available space. The system reserves a portion of the storage pool for overhead, so the full nominal capacity is not available for vault creation. Always allocate additional storage capacity beyond the requested vault size to avoid failures.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where you want to create an Exascale storage vault.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Exascale storage vaults**.
6. Click **Create Exascale storage vault**.
7. On the resulting Create Exascale storage vault window, enter the following:
   - Compartment: Select a compartment where you want this resource to be created.
   - Name: Enter a descriptive name for the vault.
   - Storage capacity for the databases: Enter a reasonable storage capacity within the minimum and maximum values displayed on the screen.

Copyright © 2008, 2026, Oracle and/or its affiliates.

8. Click **Create**.
   The Exascale Storage Vault Details page provides key information, including the allocated storage and a list of VM Clusters associated with the vault.
9. Alternatively, you can create a new vault on the fly using the Create new storage vault option while provisioning the Exascale VM cluster.

Using the Console to Scale an Exascale Storage Vault

To scale an Exascale storage vault, be prepared to provide values for the Exascale storage vault configuration.

> **Note**
>
> The minimum configurable or resizable size for an Exascale Storage Vault is 2 TB.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where the Exascale storage vault you want to scale reside.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Exascale storage vaults**.
6. Click the name of the **Exascale storage vault** you want to scale.
7. On the resulting Exascale storage vault Details page, click **Scale storage vault**.
8. On the resulting Scale storage vault window, enter the storage capacity for the databases within the minimum and maximum values displayed on the screen.
9. Click **Save**.

Using the Console to Move an Exascale Storage Vault to Another Compartment

To move an Exascale storage vault to another compartment, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where the Exascale storage vault you want to scale reside.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Exascale storage vaults**.
6. Click the name of the **Exascale storage vault** you want to move.

Copyright © 2008, 2026, Oracle and/or its affiliates.

7. On the resulting **Exascale storage vault Details** page, click **Actions**, and then select **Move resource**.
8. On the resulting dialog, choose the new compartment for the Exascale storage vault, and click **Move**.

Using the Console to Delete an Exascale Storage Vault

To delete an Exascale storage vault, use this procedure.

> **Note**
>
> An Exascale storage vault can only be deleted when it is not associated with any Exascale VM clusters. To delete an active vault, you must first terminate all VM clusters and the underlying databases that are using the vault.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where the Exascale storage vault you want to scale reside.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Exascale storage vaults**.
6. Click the name of the **Exascale storage vault** you want to move.
7. On the resulting **Exascale storage vault Details** page, click **Delete**.
8. On the resulting dialog, enter the name of the Exascale storage vault, and click **Terminate Exascale storage vault** to confirm the action.

Using the Console to Configure Exascale VM Storage on Oracle Exadata Database Service on Cloud@Customer Infrastructure

To configure Exascale VM storage on your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the infrastructure configuration.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where you want to configure Exascale storage.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Actions**, and then select **Configure Exascale storage**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

6. On the Configure Exascale storage panel, specify the storage capacity you want to allocate for Exascale VM storage.
7. Click **Submit**.

Configure All Future VM Backups to Use Exascale Storage on a VM Cluster

If a VM Cluster (ASM or Exascale) is using local database server storage for Oracle-managed VM backups, This option allows you to migrate all future VM backups to Exascale storage.

> **Note**
>
> This migration from database server storage to Exascale storage is a one-time operation and cannot be reversed.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select the region that contains your Exadata infrastructure.
3. Click **Exadata VM Clusters**.
4. From the list, select the VM cluster for which you want to change the storage type for Oracle-managed VM backups.
5. On the VM cluster details page, in the Exadata VM storage section, an Edit button will appear next to Backup storage type field. Click **Edit**.
6. On the Edit VM backup storage setting panel, review the message and enter the **VM cluster name** to confirm.
7. Click **Save**.

After you save the changes, the backup storage type is updated to Exascale, and all future Oracle-managed VM backups are stored on Exascale storage.

Create an Advanced Cluster File System (ACFS) Using Exascale Block Volume

Before you begin, ensure that the VM cluster is configured to use Exascale for database storage and that sufficient storage capacity is available in the Exascale storage vault associated with the cluster.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Click **Exadata VM Clusters**.
3. From the list, select the VM cluster for which you want to create the Advanced Cluster File System (ACFS).
4. On the VM cluster details page, select the **Advanced Cluster File System** tab.
5. Click **Create Advanced Cluster File System**.
6. On the Create Advanced Cluster File System page, enter a name and specify the desired storage capacity for the ACFS.
7. Click **Save**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The VM cluster transitions to the Updating state. Once the operation completes, the state changes back to Available (green), and the new ACFS metadata is displayed in the Advanced Cluster File System section.

The system creates a new Advanced Cluster File System with the specified size and mounts it across all nodes in the VM cluster, providing shared storage access to the cluster.

> **Note**
>
> You can also perform mount, unmount, scale ACFS storage, and terminate operations on ACFS.

Deconfigure Exascale Storage on the Exascale Infrastructure

Before you begin, ensure that the VM cluster is configured to use Exascale for database storage and that sufficient storage capacity is available in the Exascale storage vault associated with the cluster.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Click **Exadata Infrastructure**.
3. From the list, select the Exadata Infrastructure where you want to deconfigure Exascale storage.
4. On the Exadata Infrastructure details page, in the Exadata Storage section, click **Configure**.
5. On the Configure Exascale Storage page, toggle **Deconfigure Exascale Storage** to On, and then click **Submit**.

The Exadata Infrastructure resource enters the Updating state. When the update is complete, the state changes back to Available (green), and the Exadata Storage section displays Exascale Storage Capacity: 0.

Once deconfiguration is complete, Exascale is removed from the Exadata Infrastructure, and the previously allocated Exascale storage is returned to the available Exadata storage pool.

> **Note**
>
> If Exascale storage is currently being used for VM or Database storage in the infrastructure, you must first terminate all VMs and Exascale storage vaults utilizing the Exascale storage before de-configuring it on the Exadata Infrastructure.

Checklists for Oracle Exadata Database Service on Cloud@Customer Deployments

To determine your readiness for an Oracle Exadata Database Service on Cloud@Customer deployment, review the deployment checklists.

- System Components Checklist for Oracle Exadata Database Service on Cloud@Customer
  Use this checklist to ensure that the system component considerations are addressed.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Data Center Room Checklist for Oracle Exadata Database Service on Cloud@Customer
  Use this checklist to ensure that the data center room requirements are addressed.
- Data Center Environment Checklist for Oracle Exadata Database Service on Cloud@Customer
  Use this checklist to ensure that the data center environment requirements are addressed.
- Access Route Checklist for Oracle Exadata Database Service on Cloud@Customer
  Use this checklist to ensure that the access route requirements are addressed.
- Facility Power Checklist for Oracle Exadata Database Service on Cloud@Customer
  Use this checklist to ensure that the facility power requirements are addressed.
- Safety Checklist for Oracle Exadata Database Service on Cloud@Customer
  Use this checklist to ensure that safety requirements are addressed.
- Logistics Checklist for Oracle Exadata Database Service on Cloud@Customer
  Use this checklist to ensure that the logistics requirements are addressed.
- Network Configuration Checklist for Oracle Exadata Database Service on Cloud@Customer
  Use this checklist to ensure that the network configuration requirements are addressed.
- Reracking Checklist for Oracle Exadata Database Service on Cloud@Customer
  Use this checklist to determine your readiness for reracking.

System Components Checklist for Oracle Exadata Database Service on Cloud@Customer

Use this checklist to ensure that the system component considerations are addressed.

- How many racks do you plan to install?
- Will additional equipment be attached to or installed in the rack?
  If additional equipment is attached, then ensure that the additional equipment meets Oracle guidelines, and there is sufficient power and cooling.

Data Center Room Checklist for Oracle Exadata Database Service on Cloud@Customer

Use this checklist to ensure that the data center room requirements are addressed.

Answer yes, no, not applicable, or add your comments. Or let the site survey team fill in the requested information.

- Has the rack location been allocated and is it vacant?
- Does the floor layout meet the equipment maintenance access requirements?
- Will the rack be positioned so that the exhaust air of one rack does not enter the air inlet of another rack?
- Have cabinet stabilization measures been considered?
- If the data center has a raised floor:
  - Does the raised floor satisfy the weight requirements for the rack?
  - Is permission required to remove floor tiles for cabling and servicing below the floor?
- Will the rack location require any non-standard cable lengths?
- Is the floor-to-ceiling height a minimum of 2914 mm (114.72 inches)?

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Is the depth of the raised floor a minimum of 46 cm (18 inches)?

Data Center Environment Checklist for Oracle Exadata Database Service on Cloud@Customer

Use this checklist to ensure that the data center environment requirements are addressed.

Answer yes, no, not applicable, or add your comments. Or, let the site survey team fill in the requested information.

- Does the computer room air conditioning meet temperature and humidity requirements?
- Does the installation floor layout satisfy the ventilation requirements?
- If the room cooling is from a raised floor:
  - Are the perforated floor tiles each rated at 400 CFM or greater?
  - Can additional perforated floor tiles be obtained if required for additional cooling?
- Does the data center air conditioning provide sufficient front-to-back airflow?
- Is airflow adequate to prevent hot spots?
- Can the data center continuously satisfy the environmental requirements?

Access Route Checklist for Oracle Exadata Database Service on Cloud@Customer

Use this checklist to ensure that the access route requirements are addressed.

Answer yes, no, not applicable, or add your comments. Or, let the site survey team fill in the requested information.

- Has the access route been checked for clearance of the rack, including the minimum width and height requirements for all doors on the route?
- Are there any stairs, ramps, or thresholds that are of concern?
  If yes, then provide details.
- Are all access route incline angles within the permitted range (under 6 degrees)?
- Are all the surfaces acceptable for rolling the new unpacked and packed equipment?
- If a pallet jack is to be used:
  - Can the pallet jack support the weight of the rack?
  - Are the pallet jack tines compatible with the shipping pallet?
- If there are stairs, is a loading elevator available for the equipment?
- If an elevator is to be used:
  - Does the elevator car meet the height, width, and depth requirements for carrying the rack?
  - Do the elevator doors meet the height and width requirements for moving the rack?
  - Does the elevator meet the weight requirements for transporting the rack?
- Can the complete access route support the weight of the rack?
- Is the access route onto the raised floor rated for dynamic loading of the rack?

Copyright © 2008, 2026, Oracle and/or its affiliates.

Facility Power Checklist for Oracle Exadata Database Service on Cloud@Customer

Use this checklist to ensure that the facility power requirements are addressed.

Answer yes, no, not applicable, or add your comments. Or, let the site survey team fill in the requested information.

- Have the operating voltage and electric current requirements been reviewed?
- What type of power supply will be used?
  - Single-phase or 3-phase.
  - Low-voltage or High-voltage.
- Are enough power outlets provided within 2 meters for each rack?
- Do the power outlets have appropriate socket receptacles for the planned Power Distribution Units (PDUs)?
- Will optional ground cables be attached to the rack?
- Are the electrical circuits suitable in terms of voltage and current-carrying capacities?
- Does the power frequency meet the equipment specifications?
- Are power outlets available for the new equipment at the designated location?
- Will system power be delivered from two separate grids?
- Is there a UPS to power the equipment?
- Are the minimum required power sources available to support the power load (kW or kVA) for the new hardware?

Safety Checklist for Oracle Exadata Database Service on Cloud@Customer

Use this checklist to ensure that safety requirements are addressed.

Answer yes, no, not applicable, or add your comments. Or, let the site survey team fill in the requested information.

- Is there an emergency power shut off?
- Is there a fire protection system in the data center room?
- Is the computer room adequately equipped to extinguish a fire?
- Is antistatic flooring installed?
- Is the area below the raised floor free of obstacles and blockages?

Logistics Checklist for Oracle Exadata Database Service on Cloud@Customer

Use this checklist to ensure that the logistics requirements are addressed.

Answer yes, no, not applicable, or add your comments. Or, let the site survey team fill in the requested information.

- Is contact information for the data center personnel available?

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Is there security or access control for the data center?
- Are there any security background checks or security clearances required for Oracle personnel to access the data center?
  If yes, then provide the process for Oracle to follow.
- How many days in advance must background checks be completed?
- Are there any additional security access issues?
- Is computer room access available for installation personnel?
- Are laptops allowed in the data center?
- Are cell phones allowed in the data center?
- Are cameras allowed in the data center?
- Does the building have a delivery dock?
- Is there a delivery / unpacking / staging area?
- Is inside delivery planned (direct to the final rack location in the data center room)?
- If the delivery is not inside, then is the site prepared for uncrating?
- Is the delivery / unpacking / staging area protected from the elements?
- Does the building have adequate receiving space?
- Is the unpacking area air-conditioned to avoid thermal shock for various hardware components?
- Will sufficient moving personnel be available to transport the rack?
- Is union labor required for any part of the delivery or installation?
- Is the site prepared for uncrating and packaging removal?
  Package removal should take place outside the data center room.
- Is uncrating of cabinet and packaging removal required?
- Are there any restrictions on delivery truck length, width, or height?
- Is there storage space (cabinet) for the ride along spares?
  If not, does the customer allow cardboard boxes and other packing material in the computer room, since the spares are packed in cardboard boxes?
- Is there a time constraint on dock access?
  If yes, provide time constraints.
- Is a tail or side lift required on the delivery carrier to unload the equipment at the delivery dock?
- Will any special equipment be required to place the rack in the data center room?
  For example:
  - Stair walkers
  - Lifters
  - Ramps
  - Steel plates
  - Floor covers
- Does the delivery carrier require any special equipment, such as non-floor damaging rollers, transport dollies, pallet jacks, or fork lifts?

Copyright © 2008, 2026, Oracle and/or its affiliates.

Network Configuration Checklist for Oracle Exadata Database Service on Cloud@Customer

Use this checklist to ensure that the network configuration requirements are addressed.

Answer yes, no, not applicable, or add your comments. Or, let the site survey team fill in the requested information.

- Are the required network cables laid from the network equipment to the location where the Oracle Exadata Rack will be installed?
- Are the network cables that will connect the Oracle Exadata Rack labeled?
- Will the 10 GbE or 25 GbE interfaces be used for the client and backup access networks? If so, did the customer install the appropriate cables to their switch?

Reracking Checklist for Oracle Exadata Database Service on Cloud@Customer

Use this checklist to determine your readiness for reracking.

Use this checklist to ensure that the reracking requirements are addressed.

> **Note**
>
> - Reracking requires prior approval. The checklist below provides high level guidance on re-racking requirements. Oracle maintains a detailed internal checklist that must be approved prior to performing reracking.
> - You must purchase the Oracle Reracking service.
> - Oracle does not provide technical support for customer-supplied equipment.

Answer yes, no, not applicable, or add your comments. Or, let the site survey team to fill in the requested information.

- Have you purchased the Oracle Reracking Service?
- Is there a cart capable of carrying the weight of the servers to move the components and associated cabling from the supplied rack to the rack that you supply?
- Is the target rack empty?
- Attach pictures of the target rack (inside and outside).
- Does the target rack meet the following requirements?
  - Height: 42 RU
  - Width: 600 mm (23.62 inches)
  - Depth: 1112 mm (43.78 inches) without front and rear doors

  If the rack is less than 42 RU tall, then the rack must be at least 30 RU tall, and you must provide compatible PDUs to install in the target rack.
- Is the distance between the front and rear mounting planes between the minimum of 610 mm and the maximum 915 mm (24–36 inches)?

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Is the clearance depth in the front of the front mounting plane (distance to the front cabinet door) at least 25.4 mm (1 inch)?
- Does the target rack meet the following minimum load capacity?
  - 19 kg (41.89 lb) per RU
  - 785 kg (1730.63 lb) total
- Is the rack a four-post rack (mounting at both front and rear)?

> **Note**
>
> Two-post racks are not compatible.

- Does the target rack's horizontal opening and unit vertical pitch conform to ANSI/EIA 310-D-1992 or IEC 60297 standards?
- Does the target rack have RETMA rail support?

> **Note**
>
> Oracle Exadata rack requires 19 inches (483 mm) for RETMA rail spacing width. The minimum rack width of 600 mm (23.63 inches) is recommended to accommodate the PDU and cable harnesses on the side. If the rack is less than 600 mm wide, then it must have additional depth to accommodate mounting behind the server cable management arms.

- Does the target rack support Oracle cable management arms?
- Does the target rack support installation of Oracle vented and solid filler panels?
- Can the target rack provide tie-downs along the left rear side of the rack (as viewed from the front of the rack) to support the lnfiniBand cables?
- Can the target rack provide tie-downs for the Ethernet wiring harness?
- Is there sufficient space for the cable harnesses and the PDUs in the target rack?
- Can a label with the Oracle Exadata Rack serial number be printed and attached to the target rack?
- Does the target rack support installation of standard Oracle PDUs?
  If not, then complete the following checklist items:
  - Can you provide provide an equivalent pair of PDUs?
  - Can you provide provide two PDUs, each with a capacity of 10 kVA?
  - Can you provide provide at least 17 x lOA C13 plugs per PDU?
  - Can you provide a single PDU and its circuits to support the Oracle Exadata Rack power requirements in case one PDU fails?
  - Can you ensure that power loads are evenly distributed across all circuits of a single PDU?
  - Can you provide appropriate power drops for the PDUs?

Copyright © 2008, 2026, Oracle and/or its affiliates.

4

Getting Started with Oracle Exadata Database Service on Cloud@Customer Deployment

After completing the preparation tasks in Preparing for Oracle Exadata Database Service on Cloud@Customer, get started with deploying your Oracle Exadata Database Service on Cloud@Customer system following these procedures.

- About Provisioning Oracle Exadata Database Service on Cloud@Customer Systems
  To provision an Oracle Exadata Database Service on Cloud@Customer system, you must work with Oracle to set up and configure the system.
- Tagging Oracle Exadata Database Service on Cloud@Customer Resources
  Tagging is a powerful foundational service for Oracle Cloud Infrastructure (OCI) that enables users to search, control access, and do bulk actions on a set of resources based on the tag.
- Overview of Elastic Storage Expansion
  With elastic storage expansion, you can dynamically increase your storage capacity to meet your growing workload requirements.
- Using the Console to Provision Oracle Exadata Database Service on Cloud@Customer
  Learn how to provision an Oracle Exadata Database Service on Cloud@Customer system.
- Creating First VM Cluster Network on Oracle Exadata Database Service on Cloud@Customer
  Follow the steps in the following sections to create your first VM Cluster Network.
- Provisioning the First VM Cluster on an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to create the first VM cluster on your Oracle Exadata Database Service on Cloud@Customer system.
- Creating Database Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer provides a backup facility, which you can configure individually on each database.
- Creating First Database Home on an Oracle Exadata Database Service on Cloud@Customer System
  After provisioning a VM Cluster, create your first Oracle Database Home on Exadata Cloud@Customer.
- Creating First Database on an Oracle Exadata Database Service on Cloud@Customer System
  After provisioning an Oracle Database Home and any needed backup destinations, you are ready to create your first database on Oracle Exadata Database Service on Cloud@Customer.
- Connecting to an Oracle Exadata Database Service on Cloud@Customer System
  After deploying your Oracle Exadata Database Service on Cloud@Customer system with a VM cluster, Oracle Database Home, and Oracle Database, learn how to connect to your VM Cluster virtual machine using SSH and to connect to an Oracle Exadata Database Service on Cloud@Customer database using Oracle Net Services (SQL*Net).

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Create API Access Control
  Use this procedure to create API Access Control for your ExaDB-C@C resources—Exadata Infrastructure, VM Cluster, Database, and Pluggable Database.

Using the Console to Create API Access Control

Use this procedure to create API Access Control for your ExaDB-C@C resources—Exadata Infrastructure, VM Cluster, Database, and Pluggable Database.

1. Go to the details page of the respective resource.
2. In the **General Information** section, locate **API Access Control**, and click **Create**.
   - If API Access Control has not been created yet, this option will be available, and the status will be shown as **Disabled**.
   - After creation, the status changes to **Enabled**, and the **Create** option is replaced with the **OCID** of the newly created API Access Control.
3. Click the OCID to view the API Access Control details.
4. You can disable the API Access Control at any time, if needed.

About Provisioning Oracle Exadata Database Service on Cloud@Customer Systems

To provision an Oracle Exadata Database Service on Cloud@Customer system, you must work with Oracle to set up and configure the system.

Provisioning an Oracle Exadata Database Service on Cloud@Customer system is a collaborative process. The process is performed in the following sequence:

1. You create the Oracle Exadata Database Service on Cloud@Customer infrastructure.
2. You generate a file containing the infrastructure configuration details, and provide it to Oracle.
3. The Oracle Exadata Database Service on Cloud@Customer system is physically installed in your data center.
4. Oracle uses the infrastructure configuration file to perform initial system configuration. At the end of this task, Oracle supplies you with an activation file.
5. You activate the Exadata Database Service on Cloud@Customer infrastructure by using the supplied activation file.

When the provisioning process is complete, the Oracle Exadata Database Service on Cloud@Customer system is ready for you to use. You can then create a virtual machine (VM) cluster, and later create some databases.

> ⚠ **Caution**
>
> Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, the APIs, or the command-line interface.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Tagging Oracle Exadata Database Service on Cloud@Customer Resources

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

Importance of Tagging

Using the Oracle Cloud Infrastructure (OCI) tagging system, you can tag resources per your organizational scheme allowing you to group resources, manage costs, and give insights into usage. Tags also help you build a governance model around security and Maximum Availability Architecture (MAA). As your organization expands its cloud footprint, it can become challenging to keep track of the deployment architectures, security best practices, MAA, application tier, and so on. Using metadata tags to identify workload attributes can help keep up with the security and availability of your tenancy without cost overruns.

To enable customers to manage OCI resources securely and cost-effectively, Oracle provides a set of pre-defined tags in line with best practices for tagging resources. These tags are grouped into two namespaces, the `oracleStandard` namespace, and the `OracleApplicationName` namespace. You can think of a tag namespace as a container for your tag keys.

Consider a scenario where your organization has multiple cloud resources such as Exadata Infrastructure, VM Cluster, DB Home, Oracle Database, Autonomous Exadata VM Cluster, Autonomous Container Database, Autonomous AI Database, and VM Cluster Networks across multiple compartments in your tenancy. across multiple compartments in your tenancy. Suppose you wish to track these cloud resources for specific purposes, report on them, or take bulk actions. In that case, you will need a system that lets you group these resources based on different criteria such as environment, criticality, target users, application, etc. You can achieve this by applying appropriate tags to these resources.

For example, you may tag all resources in your development stack with `Oracle-Standard.Environment=Dev` or for a business-critical application stack set `Oracle-Standard.Criticality=High` or `Extreme`. In the event of service disruptions due to various reasons, you would then be able to quickly identify all OCI resources associated with an application or business function or be able to separate critical and non-critical workloads.

Tagging can also help you deploy optimized configurations based on workload attributes identified via tags. For example, database deployments for the PeopleSoft application require a specific configuration. Setting the `ApplicationName` and `AppMajorVersion` tags while deploying an Oracle Database or an Oracle Autonomous AI Database can ensure that the database is configured and ready for the particular application, for example, PeopleSoft out of the box.

Moreover, integration with the Cloud Advisor OCI service can provide you with direct, deep insight into how well your cloud services adhere to the corporate guidelines and help your management govern with a vision. See *Cloud Advisor Overview* for more details.

Adding Tags

You can tag resources using the Oracle Cloud Infrastructure (OCI) console, command-line interface, or SDK.

There are many cloud resources that can be tagged in an Oracle Exadata Database Service on Cloud@Customer deployment. Exadata Infrastructure, VM Cluster, DB Home, Oracle Database, Autonomous Exadata VM Cluster, Autonomous Container Database, Autonomous AI Database, and VM Cluster Networks are some of them. Tags can either be applied while

Copyright © 2008, 2026, Oracle and/or its affiliates.

creating the resources or modified later. For example, you can apply tags to an Autonomous Container Database (ACD) while provisioning the ACD or add them later from its **Details** page.

See *How Tagging Works* for more details on using tags. Tagging integrates with Oracle Cloud Infrastructure authorization system. You can use IAM policy controls to enable delegation or restriction of tag manipulation. See *Authentication and Authorization* to learn about the permissions required to work with defined and free-form tags. (Required) Enter introductory text here, including the definition and purpose of the concept.

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

Your tenancies come with a library of standard tags that would apply to most resources. These tags are currently available as a set of Tag Namespaces that your governance administrators can deploy. OCI best practices recommend applying these tags to all resources a standard tag can be applied to. Besides reporting and governance, OCI service automation can deliver workload-specific optimizations based on standard tag values.

For example, database deployments for the PeopleSoft application require a specific configuration. By setting the appropriate application tag key in the `Oracle-ApplicationName` tag namespace while deploying an Autonomous AI Database, can ensure that the database is configured ready for the particular application, for example, PeopleSoft out of the box.

Figure 4-1 Tagging Example

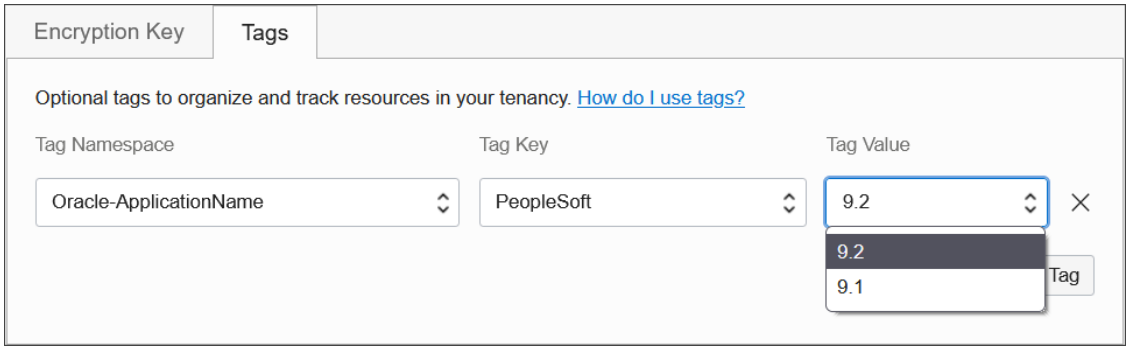

Oracle Standard Tags

Your tenancy governance administrators can deploy the standard tags at the tenancy level and may also mark certain tags as required, thereby enforcing tags on resources in those compartments. The following are the standard tags defined in the namespace called `OracleStandard`. For more information about importing standard tags, see *To import standard tags* under the *Managing Tag Namespaces* section.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 4-1 Oracle Standard 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.
`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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 4-1 (Cont.) Oracle Standard Tags

Tag Key	Tag Value Options	Description
`OracleStandard.EndUserCount`	• 1 • 10 • 100 • 1000 • 10000 • 100000 • 1000000 • 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@acme.com* or *app_support_grp@acme.com*	Denotes the email address of the resource owner.
`OracleStandard.Org`	• HR • Finance • Marketing • Sales • Legal • R&D • Customer Suppport • 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.
`OracleStandard.RecoveryTimeObjectiveMinutes`	0-10080	Time in minutes. Denotes the maximum time within which the resource is required to recover from a failure.
`OracleStandard.RecoveryPointObjectiveMinutes`	0-1440	Time in minutes. Maximum data loss tolerance for a data store resource such as a database or a storage device.

List of Compliance Regulations

Table 4-2 List of Compliance Regulations

Regulation	Description
PCI DSS	Payment Card Industry Data Security Standard
HIPAA	Health Insurance Portability and Accountability Act
ISO	International Standards Organization
SOC1	System and Organization Controls 1
SOC 2	System and Organization Controls 2
FedRamp	Federal Risk and Authorization Management Program

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 4-2 (Cont.) List of Compliance Regulations

Regulation	Description
GLBA	Gramm–Leach–Bliley Act
CCPA	California Consumer Privacy Act
SOX	Sarbanes Oxley
NIST	National Institute of Standards and Technology - Cyber Security
FISMA	Federal Information Security Management
HITECH	Health Information Technology for Economic and Clinical Health Act
FERPA	Family Educational Rights and Privacy Act (Student privacy)
FACTA	Fair and Accurate Credit Transaction Act
Texas HB300	Texas Medical Records Privacy Act
CIS	Center for Internet Security
CJIS	Criminal Justice Information Services Security Policy
C-TPAT	Customs-Trade Partnership Against Terrorism
COPPA	Children's Online Privacy Protection Act
PIPED Act, or PIPEDA	Personal Information Protection and Electronic Documents Act
GDPR	General Data Protection Regulation
PIPL	Personal Information Protection Law

Oracle Application Name Tags

Table 4-3 Oracle Application Name Tags

Tag Key	Tag Value Options	Description
Hyperion	• 11.2 • 11.1	Denotes the version of the Hyperion application.
JD Edwards	• 9.2 • 9.1 • 9.0	Denotes the version of the JD Edwards application.
Oracle_E-Business_Suite	• 12.2 • 12.1 • 12.1 • 11i	Denotes the version of the Oracle E-Business Suite application.
PeopleSoft	• 9.2 • 9.1	Denotes the version of the PeopleSoft application.
Siebel	• 8.2 • 8.1	Denotes the version of the Siebel application.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 4-3 (Cont.) Oracle Application Name Tags

Tag Key	Tag Value Options	Description
Other_Oracle_Application	Free form tag in string format.	Can be used to denote any application other than those listed above. You can enter the application name as a string value.

Related Topics

- Cloud Advisor Overview
- Oracle Autonomous AI Database Dedicated for Fleet Administrators Workshop
- How Tagging Works
- Authentication and Authorization
- Managing Tag Namespaces
- Using Predefined Values

Overview of Elastic Storage Expansion

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

Expand the storage capacity on-demand by scaling up the infrastructure with additional storage servers. You can allocate additional storage capacity available from the newly added storage servers to the already deployed VM Cluster without disrupting the current running workloads. Additional storage capacity from newly added storage servers is also available for provisioning new VM Clusters on the infrastructure.

With the elastic storage expansion capability, you can:

- Provision new Exadata Infrastructure with custom storage capacity.
- Start with a smaller storage footprint for the Exadata Infrastructure at install time.
- Expand the storage capacity on existing deployed Exadata Infrastructure on-demand in an automated, elastic fashion.
- Allocate additional storage capacity available from newly added storage servers to already deployed VM clusters and/or use them for provisioning new VM clusters on the infrastructure.

Table 4-4 Exadata Storage Server X7-2 key additional resources

Specification	Exadata Base System Storage Server X7-2	Exadata Storage Server X7-2
Additional Raw Flash Storage Capacity	6.4 TB	25.6 TB
Additional Raw Disk Storage Capacity	48 TB	120 TB
Additional Usable Storage Capacity	14 TB	35.3 TB

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 4-5 Exadata Storage Server X8-2 key additional resources

Specification	Exadata Base System Storage Server X8-2	Exadata Storage Server X8-2
Additional Raw Flash Storage Capacity	12.8 TB	25.6 TB
Additional Raw Disk Storage Capacity	84 TB	168 TB
Additional Usable Storage Capacity	24.6 TB	49.6 TB

Table 4-6 Exadata Storage Server X8M-2 key additional resources

Specification	Exadata Base System Storage Server X8M-2	Exadata Storage Server X8M-2
Additional Raw Flash Storage Capacity	12.8 TB	25.6 TB
Additional Raw Disk Storage Capacity	84 TB	168 TB
Additional Usable Storage Capacity	24.6 TB	49.6 TB
Additional Persistent Memory	-	1.5 TB

Table 4-7 Exadata Storage Server X9M-2 key additional resources

Specification	Exadata Base System Storage Server X9M-2	Exadata Storage Server X9M-2
Additional Raw Flash Storage Capacity	12.8 TB	25.6 TB
Additional Raw Disk Storage Capacity	84 TB	216 TB
Additional Usable Storage Capacity	24.6 TB	63.6 TB
Additional Persistent Memory	-	1.5 TB

Table 4-8 Exadata Storage Server X10M key additional resources

Specification	Exadata Storage Server X10M
Additional Raw Flash Storage Capacity	27.2 TB
Additional Raw Disk Storage Capacity	264 TB
Additional Usable Storage Capacity	80 TB
Additional Exadata RDMA Memory	1.25 TB

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 4-9 Exadata Storage Server X11M key additional resources

Specification	Exadata Storage Server X11M-BASE	Exadata Storage Server X11M-HC
Additional Raw Flash Storage Capacity	13.6 TB	27.2 TB
Additional Raw Disk Storage Capacity	132 TB	264 TB
Additional Usable Storage Capacity	35.6 TB	80 TB
Additional Exadata RDMA Memory	-	1.25 TB

Elastic scaling of Exadata Storage Servers is subject to the following conditions:

- The Exadata Cloud@Customer system configuration must be based on Oracle Exadata X7 hardware, Oracle Exadata X8 hardware, Oracle Exadata X8M hardware, or Oracle Exadata X9M hardware.
- Each Exadata Cloud@Customer system configuration can have an absolute maximum of:
  - X7, X8: 12 Exadata Storage Servers
  - X8M, X9M, X10M, X11M: 64 Storage Servers
- Exadata Infrastructure deployed with base configuration shape can only be expanded using base expansion SKU storage servers.
- Exadata Infrastructures deployed with X7 generation at install time can be scaled with X8 generation storage servers. X8 storage servers used to scale X7 infrastructure will only present the same total usable capacity as all other X7 storage servers that are already part of the infrastructure.
- Exadata Infrastructures deployed with X8M generation at install time can only be scaled with X8M or higher generation storage servers. X9M and X10M storage servers used to scale X8M infrastructure will only present the same total usable capacity as other X8M storage servers that are already part of the infrastructure.
- Exadata Infrastructures deployed with X9M generation at install time can only be scaled with X9M or higher generation storage servers. X10M storage servers used to scale X9M infrastructure will only present the same total usable capacity as other X9M storage servers that are already part of the infrastructure.
- Exadata Infrastructures deployed with X10M generation at install time can only be scaled with X10M or higher generation storage servers. X11M storage servers used to scale X10M infrastructure will only present the same total usable capacity as other X10M storage servers that are already part of the infrastructure.

Exadata Infrastructures deployed with additional storage servers will be configured as an Elastic shape with the total number of storage servers and usable capacity clearly called out for the given infrastructure.

Before you can scale the number of Exadata storage servers, review the site and network requirements, and the checklists to prepare and deploy Exadata Cloud@Customer. Ensure that you have worked with sales and followed the procurement process. The following figure provides you with an overview of the order and deployment process.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Figure 4-2 Overview of Order and Deployment Process

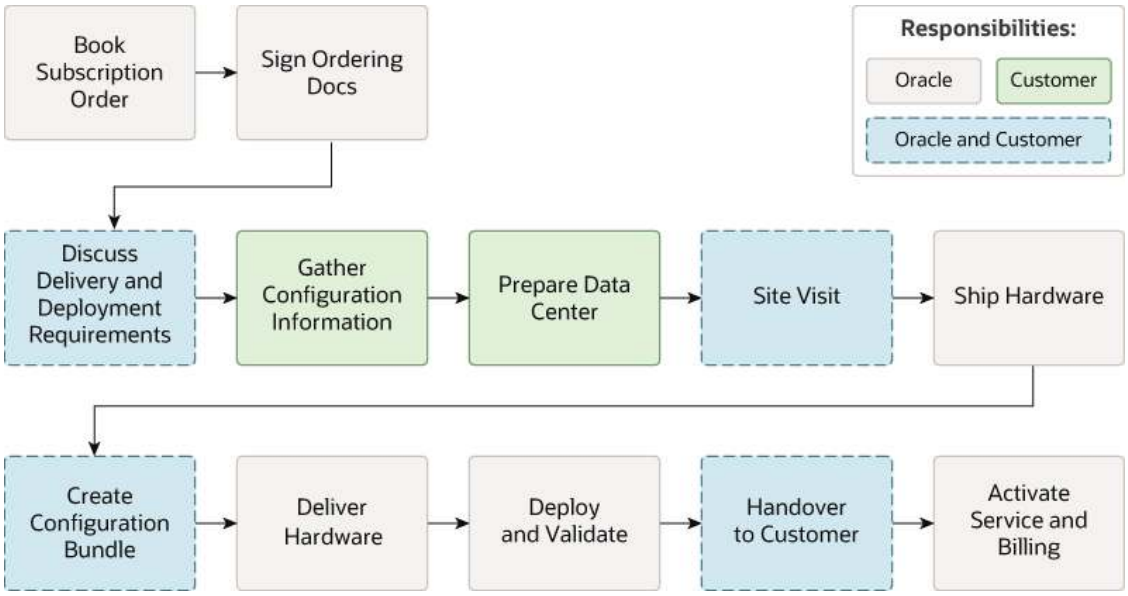

Multi-Rack Expansion

- **Rack Minimum Configuration**
  - **X11M/X10M**
    - Maximum 16 total servers in the initial rack
    - Maximum 18 total servers in the expansion rack
  - **X8M/X9M**
    - Base system expandable to a single rack
    - Maximum 8 database and 12 storage servers in the first rack
    - Maximum 11 database and 18 storage servers in the expansion racks
- **Minimum Expansion Rack Configuration**
  - Start with one storage server and add at least 3 other servers (storage and compute).
- **Maximum Overall Configuration**
  - **Maximum total database servers:** 32
  - **Maximum total storage servers:** 64
  - **Maximum total racks (including initial):** 6

Related Topics

- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.
- Using the Console to Download Scale Infrastructure Compute and Storage Configuration File
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Activate New Database and Storage Servers
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.
- Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption
  To make storage capacity from the new servers for VM clusters consumption, complete this procedure.
- Using the Console to View Details of Oracle Exadata Database Service on Cloud@Customer Infrastructure with Scaled Compute and Storage Capacity
  To view the storage capacity from the new storage server, or scaled compute capacity, use this procedure.

Using the Console to Provision Oracle Exadata Database Service on Cloud@Customer

Learn how to provision an Oracle Exadata Database Service on Cloud@Customer system.

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Enabling or Disabling the Control Plane Server Diagnostic Offline Report
  To enable or disable CPS offline report, use this procedure.
- Using the Console to View Exadata Infrastructure Network Configuration Details
  To view network configuration details, follow these steps. Save this information for later use to troubleshoot if you face network issues.
- Using the Console to Edit Oracle Exadata Database Service on Cloud@Customer Infrastructure Networking Configuration
  To edit your Oracle Exadata Database Service on Cloud@Customer infrastructure networking configuration, be prepared to provide values for the infrastructure configuration.
- Using the Console to Download a File Containing Configuration Data
  To download an Oracle Exadata Database Service on Cloud@Customer configuration file, complete this procedure.
- Using the Console to Activate Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To activate Oracle Exadata Database Service on Cloud@Customer infrastructure, ensure that you meet the prerequisites, and complete this procedure.
- Using the Console to Check the Status of Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To find the status of your Oracle Exadata Database Service on Cloud@Customer infrastructure, use this procedure to check the Infrastructure Details page.

Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure

To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

2. Under **Region**, select the region that you want to associate with the Oracle Exadata infrastructure.
   The region that is associated with your Oracle Exadata infrastructure cannot be changed after the Oracle Exadata infrastructure is created. Therefore, ensure that you select the most appropriate region for your infrastructure. Consider the following factors:
   - Consider any business policies or regulations that preclude the use of a particular region. For example, you can be required to maintain all operations within national boundaries.
   - Consider the physical proximity of the region to your data center. Needless extra physical separation adds unnecessary latency to network communications between Oracle Cloud Infrastructure and your corporate data center.
3. Click **Exadata Infrastructure**.
4. Click **Create Exadata Infrastructure**.
5. In the Create Exadata Infrastructure page, provide the requested information:
   - **Provide the display name:** The display name is a user-friendly name that you can use to identify the Exadata infrastructure. The name doesn't need to be unique, because an Oracle Cloud Identifier (OCID) uniquely identifies the Oracle Exadata infrastructure.
   - **Choose a compartment:** From the list of available compartments, choose the compartment that you want to contain the Oracle Exadata infrastructure.
     For more information, see *Understanding Compartments*.
   - **Oracle Cloud Infrastructure region:** The region that is associated with your Oracle Exadata infrastructure cannot be changed after the Oracle Exadata infrastructure is created. Therefore, check the displayed region to ensure that you are using the most appropriate region for your infrastructure.
     See step 2 (earlier in this procedure) for further considerations. To switch regions now, use the **Region** menu at the top of the console.
   - **Select the Exadata system model:** From the list, choose an Oracle Exadata hardware.
     The Oracle Exadata system model and system shape combine to define the amount of CPU, memory, and storage resources that are available in the Exadata infrastructure.
     For more information, see *System Configuration*.
   - **Compute and storage configuration:**
     - **Database server type (displayed if Exadata Cloud@Customer X11M is selected):**
       - **Choices:** X11M-BASE, X11M, X11M-L, and X11M-XL
     - **Storage server type (displayed if Exadata Cloud@Customer X11M is selected):**
       - **Choices:** X11M-BASE and X11M-HC
     - **Database servers:**
       - **X11M/X10M:** You can expand the database and storage servers in the initial rack up to 16 total servers.
       - **X8M/X9M:** You can add a minimum of 2 and extend up to a maximum of 8 database servers.
     - **Storage servers:**
       - **X11M/X10M:** You can expand the database and storage servers in the initial rack up to 16 total servers.

Copyright © 2008, 2026, Oracle and/or its affiliates.

    - **X8M/X9M:** You can add a minimum of 3 and extend up to a maximum of 12 storage servers. For each storage server you add, the storage capacity that will be added is displayed on the right.

  - **Multi-rack deployment:** If you scale beyond the maximum limits, Database servers: 8 and Storage servers: 12, then the multi-rack deployment toggle is automatically enabled. You can also enable that option if the database servers count and storage server count are within the minimum and maximum bounds; however, you want the infrastructure to span multiple racks.
    A new field, **Upload multi-rack configuration file**, appears when you enable this toggle button. You can drag and drop the JSON file or select one from the local filesystem. For multi-rack deployments, the Control Plane cannot generate the configuration bundle correctly without the JSON file. The JSON file will be used to generate host names and IP addresses for the additional components. Contact your Oracle Field Engineer for more information.

    If you fail to identify an infrastructure as multi-rack and subsequently determine it is multi-rack, then you will have to delete that infrastructure and recreate it.

> **Note**
>
> Resource allocation will display ECPUs for X11M database servers and OCPUs for X10M and older generation Exadata infrastructure.
> **ECPU:** An ECPU is an abstracted measure of compute resources. ECPUs are based on the number of cores elastically allocated from a pool of compute and storage servers.

- **Site network details**
  - **Time Zone:** The default time zone for the Exadata Infrastructure is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

> **Note**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, then select the **Select another time zone** option, select a **Region** or **country**, and then select the corresponding **Time zone**.
>
> If you do not see the region or country you want, then select **Miscellaneous**, and then select an appropriate **Time zone**.

- **Configure the control plane server network**
  Each Oracle Exadata Database Service on Cloud@Customer system contains two control plane servers, which enable connectivity to Oracle Cloud Infrastructure. The control plane servers are connected to the control plane network, which is a subnet on your corporate network. The following settings define the network parameters:
  - **Control Plane Server 1 IP Address:** Provide the IP address for the first control plane server. This IP address is for the network interface that connects the first control plane server to your corporate network using the control plane network.
  - **Control Plane Server 2 IP Address:** Provide the IP address for the second control plane server. This IP address is for the network interface that connects the second control plane server to your corporate network using the control plane network.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Netmask:** Specify the IP netmask for the control plane network.
- **Gateway:** Specify the IP address of the control plane network gateway.
- **HTTP Proxy:** (Optional) You can choose to use this field to specify your corporate HTTP proxy. The expected format is as follows, where *server* is the server name, *domain* is the domain name, and *port* is the assigned port:

  `http://server.domain:port`

  For example:

  `http://proxy.example.com:80`

  For enhanced security, when possible, Oracle recommends that you use an HTTP proxy.
- **Enable Control Plane Server Offline Report:** Enabling the Control Plane Server (CPS) offline report helps in diagnosing connectivity issues between the CPS and OCI endpoints, should they arise.
  To view the report, do the following:

  a. Find the CPS IP addresses.
  For more information, see Using the Console to View Exadata Infrastructure Network Configuration Details.

  b. From your local network, access the report over HTTP.
  To view the report in HTML format, use `http://<CPSPublicIP>:18080/report`

  To view the report in JSON format, use `http://<CPSPublicIP>:18080/report/json`

  For more information, see *ExaCC gen2: Troubleshooting VPN/WSS connection from Customer Side*.

- **Configure the Oracle Exadata system networks**
  Each Oracle Exadata Database Service on Cloud@Customer system contains two system networks, which are not connected to your corporate network. The following settings define IP address allocations for these networks:

  - **Administration Network CIDR Block:** Specifies the IP address range for the administration network using CIDR notation. The administration network provides connectivity that enables Oracle to administer the Exadata system components, such as the Exadata compute servers, storage servers, network switches, and power distribution units. You can accept the suggested default, or specify a custom value.
    The minimum CIDR block is /23. The maximum CIDR block is /21. Because this CIDR block cannot be changed in the future, it is recommended to use a /21 CIDR block to maximize flexibility for future possible expansion. However, the below table can be used for further guidance.
  - **Secure Fabric Network:** Specifies the IP address range for the Exadata secure fabric network using CIDR notation. The Exadata secure fabric network provides the high-speed low-latency interconnect used by Exadata software for internal communications between various system components. You can accept the suggested default, or specify a custom value.
    The minimum CIDR block is /22. The maximum CIDR block is /19. Because this CIDR block cannot be changed in the future, it is recommended to use a /19 CIDR block to maximize flexibility for future possible expansion. However, the below table can be used for further guidance.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 4-10 X11M and X10M (Quarter Rack X10M, Quarter Rack X10M-L, and Quarter Rack X10M-XL) CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
6	64	/23	/22
7	45		
8	26		
9	9		
7	64	/22	/22
8	55		
9	30		
10	5		
15	64	/22	/21
16	60		
17	45		
18	30		
19	15		
17	64	/21	/21
18	61		
19	36		
20	11		
32	64	/21	/20

Table 4-11 X9M CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
6	64	/23	/22
7	48		
8	30		
9	9		
7	64	/22	/22
8	55		
9	30		
10	5		
15	64	/22	/21
16	62		
17	45		
18	30		
19	15		
17	64	/21	/21
18	61		
19	36		
20	11		
32	64	/21	/20

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 4-12 X8M CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
8	64	/23	/22
9	61		
10	47		
11	37		
12	28		
13	13		
11	64	/22	/22
12	51		
13	34		
14	17		
23	64	/22	/21
24	56		
25	45		
26	38		
27	21		
26	64	/21	/21
27	52		
28	35		
29	18		
32	64	/21	/20

> **Note**
>
> You cannot change this value in the future. Consider all future expansions when making this choice. Otherwise, your future expandability may be limited.

- **Configure DNS and NTP services**
  Each Exadata Database Service on Cloud@Customer system requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. The following settings specify the servers that provide these services to the Exadata infrastructure:
  - **DNS Servers:** Provide the IP address of a DNS server that is accessible using the control plane network. You may specify up to three DNS servers.
  - **NTP Servers:** Provide the IP address of an NTP server that is accessible using the control plane network. You may specify up to three NTP servers.
- **Network bonding mode**

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> The use of LACP requires both the server and the switch(es) have compatible settings to allow LACP to work properly. To use LACP for ExaDB-C@C, you must configure your network switches such that they are compatible with the following parameter from the Linux `ifcfg-bondethx` configuration file:
>
> ```
> BONDING_OPTS="mode=802.3ad miimon=100 downdelay=200 updelay=200
> lacp_rate=1 xmit_hash_policy=layer3+4"
> ```
>
> Note that the `BONDING_OPTS` line in the host operating system cannot be changed, so customer switch settings must be compatible with the parameters stated above without alteration.

The client and backup networks are configured to use active-backup bonding (mode 1) by default. Since the setting gets applied to each network interface individually, you have the flexibility to configure client and backup network interfaces independently of each other. For example, you can configure the backup network to LACP and the client network to active-backup, if desired.

- **Enable LACP for Client Network**: Select to configure the client network with LACP mode 4 (mode=802.3ad) active/active dynamic link.
- **Enable LACP for Backup Network**: Select to configure the backup network with LACP mode 4 (mode=802.3ad) active/active dynamic link.

You can check the bonding mode of the client and backup networks under the Network section on the Exadata Infrastructure Details page.

- **Active/Backup** indicates that the network interfaces are configured to use the default mode 1.
- **LACP** indicates that the network interfaces are mode 4 enabled.

- **Provide maintenance details**
  - Click **Configure maintenance**.
    On the resulting Configure maintenance page, configure the following:
    - **Maintenance scheduling preference: Oracle managed schedule**
      - **Choose a maintenance method:**
        - **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
        - **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
      - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The

Copyright © 2008, 2026, Oracle and/or its affiliates.

maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.

- **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

> **Note**
>
> Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 30 minutes

Maximum: 120 minutes

- **Maintenance scheduling preference: Customer managed schedule**
  - **Maintenance schedule:** Define maintenance preferences for this infrastructure
    - **Configure maintenance preference:** Define maintenance time preferences for each quarter. If more than one preference is defined for a quarter, Oracle automation will select one of them to perform maintenance on all components in your infrastructure.
      Select at least one month every two quarters.
    - **Specify a schedule:** Choose your preferred week, weekday, start time, and lead time for infrastructure maintenance.
      - Optional. Under **Week of the month**, specify which week of the month, maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days of the month, and have a duration of 7 days. Weeks start and end based on calendar dates, not days of the week. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days. If you do not specify a week of the month, Oracle will run the maintenance update in a week to minimize disruption.
      - Optional. Under **Day of the week**, specify the day of the week on which the maintenance will occur. If you do not specify a day of the week, Oracle will run the maintenance update on a weekend day to minimize disruption.
      - Optional. Under **Hour of the day**, specify the hour during which the maintenance run will begin. If you do not specify a start hour, Oracle will pick the least disruptive time to run the maintenance update.
      - Under **Notification Lead Time**, specify the minimum number of weeks ahead of the maintenance event you would like to receive a notification message. Your lead time ensures that a newly

Copyright © 2008, 2026, Oracle and/or its affiliates.

released maintenance update is scheduled to account for your required minimum period of advanced notification.

* **Choose a maintenance method:**
  * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
  * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
* **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
  * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

    > **Note**
    >
    > Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

    Default: 30 minutes

    Maximum: 120 minutes
* **Show advanced options:**
  * Enable monthly security infrastructure maintenance: Select this check box to perform monthly security infrastructure maintenance.

* **Maintenance schedule:** Use maintenance window preferences from a scheduling policy. During infrastructure provisioning, after the scheduling policy is selected, Oracle generates a recommended maintenance scheduling plan to apply updates to all the components in your infrastructure. The recommended plan schedules all DB Servers, followed by Storage Servers and Network Switches, into the maintenance windows from your policy based on duration. After provisioning the infrastructure,

Copyright © 2008, 2026, Oracle and/or its affiliates.

you can update the scheduling plan by editing the 'Maintenance Scheduling Plan' resource and customize the update to specific components to align with different windows in your scheduling policy.

* Click **Select policy**.
* In the resulting Select maintenance scheduling policy window, choose a compartment and a policy.
  You can also create a maintenance scheduling policy and use it. For more information, see Create a Maintenance Scheduling Policy. Note that you can add additional maintenance windows to the policy after creating it. For more information, see Add Additional Maintenance Windows to a Maintenance Scheduling Policy.
* Click **Save**.

Click **Save**.

If you switch from rolling to non-rolling maintenance method, then **Confirm Non-rolling Maintenance Method** dialog is displayed.

Enter the name of the infrastructure in the field provided to confirm the changes.

Click **Save Changes**.

> **Note**
>
> After creating the infrastructure, you can find the maintenance method, maintenance schedule, DB Server version, and Storage Server version details under the **Maintenance** and **Version** sections on the **Infrastructure Details** page.

- **Show Advanced Options**

  You have the option to configure advanced options.

  – **Tags:** (Optional) You can choose to apply tags. If you have permission to create a resource, then you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, then skip this option (you can apply tags later) or ask your administrator.

6. Click **Create**.
   If all of your inputs are valid, then the Infrastructure Details page appears. The page outlines the next steps in the provisioning process. Initially, after creation, the state of the Oracle Exadata infrastructure is **Requires-Activation**.

Related Topics

- Understanding Compartments
- System Configuration Options for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.
- ExaCC gen2: Troubleshooting VPN/WSS connection from Customer Side
- Resource Tags

Copyright © 2008, 2026, Oracle and/or its affiliates.

Enabling or Disabling the Control Plane Server Diagnostic Offline Report

To enable or disable CPS offline report, use this procedure.

> **Note**
>
> - You cannot enable or disable Control Plane Server diagnostic offline report if the Exadata Infrastructure is in `DISCONNECTED` mode.
> - At any given point in time, the report is available only at the primary Control Plane Server. When generating a report, if the first IP address specified for the Control Plane Servers does not work, then you may try the second IP.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Region**, select the region that you want to associate with the Oracle Exadata infrastructure.
3. Click **Exadata Infrastructure**.
4. From the list of infrastructures, click the name of the infrastructure that you're interested in.
5. On the Infrastructure Details page, find the Control Plane Server Offline Report details under the **Network** section.
6. Click the **Enable** or **Disable** link as needed.
   - If you click **Enable**, then the Enable Control Plane Server Offline Report window is displayed.
     Review the information provided on the popup window, and then click **Enable**.
   - If you click **Disable**, then Disable Control Plane Server Offline Report window is displayed.
     Review the information provided on the popup window, and then click **Disable**.

Viewing the Control Plane Server Diagnostic Offline Report

1. Find the CPS IP addresses.
   For more information, see *Using the Console to View Exadata Infrastructure Network Configuration Details*.
2. From your local network, access the report over HTTP.
   To view the report in HTML format, use `http://<CPSPublicIP>:18080/report`

   To view the report in JSON format, use `http://<CPSPublicIP>:18080/report/json`

Related Topics

- Using the Console to View Exadata Infrastructure Network Configuration Details
  To view network configuration details, follow these steps. Save this information for later use to troubleshoot if you face network issues.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to View Exadata Infrastructure Network Configuration Details

To view network configuration details, follow these steps. Save this information for later use to troubleshoot if you face network issues.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**.
2. Under **Region**, select the region that you want to associate with the Oracle Exadata infrastructure.
3. Click **Exadata Infrastructure**.
4. From the list of infrastructures, click the name of the infrastructure that you're interested in. Note that the infrastructure must be in **Active** state.
5. On the **Infrastructure Details** page, find the network configuration details under the **Network** section.

Using the Console to Edit Oracle Exadata Database Service on Cloud@Customer Infrastructure Networking Configuration

To edit your Oracle Exadata Database Service on Cloud@Customer infrastructure networking configuration, be prepared to provide values for the infrastructure configuration.

You can only edit Oracle Exadata Database Service on Cloud@Customer infrastructure networking configuration only if the current state of the Oracle Exadata infrastructure is **Requires Activation**. Also, ensure that you do not edit the Exadata infrastructure after you download the configuration file and provide it to Oracle.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Edit Infrastructure Networking**.
6. Use the Edit Infrastructure Networking dialog to edit the Oracle Exadata infrastructure networking:
   a. **Time zone:** The default time zone for the Exadata Infrastructure is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, then select the **Select another time zone** option, select a **Region** or **country**, and then select the corresponding **Time zone**.
>
> If you do not see the region or country you want, then select **Miscellaneous**, and then select an appropriate **Time zone**.

b. **Configure the cloud control plane network**

Each Oracle Exadata Database Service on Cloud@Customer system contains two Control Plane Servers, which enable connectivity to Oracle Cloud Infrastructure. The Control Plane Servers are connected to the control plane network, which is a subnet on your corporate network. The following settings define the network parameters:

- **Control Plane Server 1 IP Address:** Provide the IP address for the first control plane server. This IP address is for the network interface that connects the first Control Plane Server to your corporate network using the control plane network.
- **Control Plane Server 2 IP Address:** Provide the IP address for the second control plane server. This IP address is for the network interface that connects the second Control Plane Server to your corporate network using the control plane network.
- **Netmask:** Specify the IP netmask for the control plane network.
- **Gateway:** Specify the IP address of the control plane network gateway.
- **HTTP Proxy:** Optionally, you can use this field to specify your corporate HTTP proxy to use for the HTTPS connection from the Control Plane Server to Oracle Cloud Infrastructure. The expected format is:

  `http://server.domain:port`

  For example:

  `http://proxy.example.com:80`

  For enhanced security, when possible, Oracle recommends that you use an HTTP proxy.

c. **Configure the Exadata system networks**

Each Oracle Exadata Database Service on Cloud@Customer system contains two system networks, which are not connected to your corporate network. The following settings define IP address allocations for these networks:

- **Administration Network CIDR Block:** Specifies the IP address range for the administration network using CIDR notation. The administration network provides connectivity that enables Oracle to administer the Exadata system components, such as the Exadata compute servers, storage servers, network switches, and power distribution units.

  The maximum CIDR block prefix length is `/23`, which defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Oracle Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth. The minimum CIDR block prefix length is `/16`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Ensure that the IP address range does not conflict with other hosts your corporate network, and does not overlap with the InfiniBand network CIDR block.

- **Secure Fabric Network CIDR Block:** Specifies the IP address range for the Exadata InfiniBand network using CIDR notation. The Exadata InfiniBand network provides the high-speed low-latency interconnect used by Exadata software for internal communications between various system components.

  The maximum CIDR block prefix length is `/22`, which defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Oracle Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth. The minimum CIDR block prefix length is `/19`.

  Ensure that the IP address range does not conflict with other hosts your corporate network, and does not overlap with the administration network CIDR block.

d. **Configure DNS and NTP services**
Each Oracle Exadata Database Service on Cloud@Customer system requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. The following settings specify the servers that provide these services to the Exadata infrastructure:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the control plane network. You can specify up to three DNS servers.
- **NTP Servers:** Provide the IP address of an NTP server that is accessible using the control plane network. You may specify up to three NTP servers.

7. Click **Save Changes**.

Using the Console to Download a File Containing Configuration Data

To download an Oracle Exadata Database Service on Cloud@Customer configuration file, complete this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Oracle Exadata infrastructure for which you want to download a file containing the infrastructure configuration details.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure for which you want to download a file containing the infrastructure configuration details.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Show guide** on the banner displayed.
6. Click **Download Configuration**.

   Your browser downloads a file containing the infrastructure configuration details.

The generated configuration file includes all the relevant configuration details for the additional storage servers included as part of the create infrastructure flow.

When you provide the generated infrastructure configuration file to Oracle, ensure that it has not been altered in any way. Also, ensure that you do not edit the Oracle Exadata infrastructure after you download the configuration file and provide it to Oracle.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Activate Oracle Exadata Database Service on Cloud@Customer Infrastructure

To activate Oracle Exadata Database Service on Cloud@Customer infrastructure, ensure that you meet the prerequisites, and complete this procedure.

- Ensure that you have added a primary contact. You cannot activate your infrastructure without adding a primary maintenance contact.
- Locate the activation file. This file is supplied to you by Oracle after installation and initial configuration of your Oracle Exadata Database Service on Cloud@Customer system.
- Ensure that the current state of your infrastructure is **Requires Activation**. You can only activate Oracle Exadata if its state is **Requires Activation**.

1. Download the activation file.
2. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
3. Choose **Region** and **Compartment**, and select the region and compartment that contains the Oracle Exadata infrastructure that you want to activate.
4. Click **Exadata Infrastructure**.
5. Click the name of the Oracle Exadata infrastructure that you want to activate.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
6. Click **Show activation guide** on the banner displayed.

   The **Show activation guide** button is only available if the Oracle Exadata infrastructure requires activation. You cannot activate Oracle Exadata infrastructure multiple times.
7. Use the Activate Exadata Infrastructure dialog to upload the activation file, and then click **Activate**.

   The activation file includes all the relevant details for the additional storage servers included as part of the create infrastructure flow.

   After activation, the state of the Oracle Exadata infrastructure changes to **Active**.

Using the Console to Check the Status of Oracle Exadata Database Service on Cloud@Customer Infrastructure

To find the status of your Oracle Exadata Database Service on Cloud@Customer infrastructure, use this procedure to check the Infrastructure Details page.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Oracle Exadata infrastructure that you are interested in.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure that you are interested in.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. Check the icon on the Infrastructure Details page. The color of the icon and the text below it indicates the status of the Oracle Exadata infrastructure.
   - **Creating**: Yellow icon. The Oracle Exadata infrastructure definition is being created in the control plane.
   - **Requires Activation**: Yellow icon. The Oracle Exadata infrastructure is defined in the control plane, but it must be provisioned and activated before it can be used.
   - **Active**: Green icon. The Oracle Exadata infrastructure is successfully provisioned and activated.
   - **Deleting**: Gray icon. The Oracle Exadata infrastructure is being deleted by using the Console or API.
   - **Deleted**: Gray icon. The Oracle Exadata infrastructure is deleted, and is no longer available. This state is transitory. It is displayed for a short time, after which the Oracle Exadata infrastructure is no longer displayed.
   - **Activation Failed**: Red icon. An error condition currently prevents the activation of the Oracle Exadata infrastructure. Typically, this state is auto-correcting, and does not require user intervention.

Creating First VM Cluster Network on Oracle Exadata Database Service on Cloud@Customer

Follow the steps in the following sections to create your first VM Cluster Network.

- About Managing VM Cluster Networks on Oracle Exadata Database Service on Cloud@Customer
  The VM cluster provides a link between your Oracle Exadata Database Service on Cloud@Customer infrastructure and Oracle Databases you deploy.
- Using the Console to Create a VM Cluster Network
  To create your VM cluster network with the Console, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Edit a VM Cluster Network
  To edit a VM Cluster Network, use this procedure. You can only edit a VM cluster network that is not associated with a VM cluster.
- Using the Console to Download a File Containing the VM Cluster Network Configuration Details
  To provide VM cluster network information to your network administrator, you can download and supply a file containing the network configuration.
- Using the Console to Validate a VM Cluster Network
  You can only validate a VM cluster network if its current state is **Requires Validation**, and if the underlying Exadata infrastructure is activated.

About Managing VM Cluster Networks on Oracle Exadata Database Service on Cloud@Customer

The VM cluster provides a link between your Oracle Exadata Database Service on Cloud@Customer infrastructure and Oracle Databases you deploy.

Before you can create any databases on your Exadata Cloud@Customer infrastructure, you must create a VM cluster network, and you must associate it with a VM cluster.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The VM cluster network specifies network resources, such as IP addresses and host names, that reside in your corporate data center and are allocated to Exadata Cloud@Customer. The VM cluster network includes definitions for the Exadata client network and the Exadata backup network. The client network and backup network contain the network interfaces that you use to connect to the VM cluster virtual machines, and ultimately the databases that reside on those virtual machines.

> **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.

Using the Console to Create a VM Cluster Network

To create your VM cluster network with the Console, be prepared to provide values for the fields required for configuring the infrastructure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure for which you want to create a VM cluster network.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure for which you want to create a VM cluster network.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click **Create VM Cluster Network**.
6. Provide the requested information on the Data Center Network Details page:
   a. Provide the display name.

      The display name is a user-friendly name that you can use to identify the VM cluster network. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the VM cluster network.

   b. Assign IPs to DB servers.

      By default, all DB servers are assigned IP addresses to enable easy addition and removal of VMs to the cluster in the future. Note that while DB Servers can be added to or removed from the VM Cluster Network, addresses cannot be changed in the future. It is advisable to plan ahead and reserve enough addresses with your network team to accommodate any and all future growth of your VM Cluster to avoid a database migration scenario in the future. If needed, then edit the list of DB servers from assigning IP addresses. If needed, then edit the list of DB servers from assigning IP addresses.

      i. Click **Edit**.
      Edit DB Server Selection dialog is displayed.

      ii. Clear the checkboxes to omit one or more DB servers.

      iii. Click **Save**.
      Assign IPs to DB servers section displays the list of selected DB servers.

Copyright © 2008, 2026, Oracle and/or its affiliates.

c. Provide client network details.

The client network is the primary channel for application connectivity to Exadata Database Service on Cloud@Customer resources. The following settings define the required network parameters:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the client network between `1` and `4094`, inclusive. To specify no VLAN tagging, enter "`1`". (This is equivalent to a "`NULL`" VLAN ID tag value.)

> **Note**
>
> The values "`0`" and "`4095`" are reserved and cannot be entered.

- **Database servers starting IP address:** Using CIDR notation, provide the IP address range for the client network.
- **Netmask:** Specify the IP netmask for the client network.
- **Gateway:** Specify the IP address of the client network gateway.
- **Hostname Prefix:** Specify the prefix that is used to generate the hostnames in the client network.
- **Domain Name:** Specify the domain name for the client network.

d. Provide backup network details.

The backup network is the secondary channel for connectivity to Exadata Database Service on Cloud@Customer resources. It is typically used to segregate application connections on the client network from other network traffic. The following settings define the required network parameters:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the backup network between `1` and `4094`, inclusive. To specify no VLAN tagging, enter "`1`". (This is equivalent to a "`NULL`" VLAN ID tag value.)

> **Note**
>
> The values "`0`" and "`4095`" are reserved, and cannot be entered.

- **Database servers starting IP address:** Using CIDR notation, provide the IP address range for the backup network.
- **Netmask:** Specify the IP netmask for the backup network.
- **Gateway:** Specify the IP address of the backup network gateway.
- **Hostname Prefix:** Specify the prefix that is used to generate the hostnames in the backup network.
- **Domain Name:** Specify the domain name for the backup network.

e. (**For authorized customers only**) Optionally, provide disaster recovery network details. The disaster recovery network is typically used to segregate Data Guard network traffic. The following settings define the required network parameters:

- Select the **Enable disaster recovery network** checkbox to configure the third network.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the disaster recovery between `1` and `4094`, inclusive. To specify no VLAN tagging, enter "1". (This is equivalent to a "`NULL`" VLAN ID tag value.)

> **Note**
>
> The values "0" and "4095" are reserved, and cannot be entered.

- **Database servers starting IP address:** Using CIDR notation, provide the IP address range for the backup network.
- **Netmask:** Specify the IP netmask for the backup network.
- **Gateway:** Specify the IP address of the backup network gateway.
- **Hostname Prefix:** Specify the prefix that is used to generate the hostnames in the backup network.
- **Domain Name:** Specify the domain name for the backup network.

f. Provide DNS and NTP server details.

The VM cluster network requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. The following settings specify the servers that provide these services:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the client network. You may specify up to three DNS servers.
- **NTP Servers:** Provide the IP address of an NTP server that is accessible using the client network. You may specify up to three NTP servers.

g. Configure Advanced Options.

Network:

- (Optional) Assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. Default is 1521.
- (Optional) Assign a SCAN listener port (TCPS) in the range between 1024 and 8999. Default is 2484.

Tags: (Optional) You can choose to apply tags. If you have permission to create a resource, then you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, refer to information about 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.

7. Click **Review Configuration**.

The Review Configuration page displays detailed information about the VM cluster network, including the hostname and IP address allocations. These allocations are initially system-generated and are based on your inputs.

8. (Optional) You can choose to adjust the system-generated network definitions on the Review Configuration page.

a. Click **Edit**.

b. Use the **Edit** dialog to adjust the system-generated network definitions to meet your requirements.

c. Click **Save Changes**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

9. Click **Create VM Cluster Network**.

   The VM Cluster Network Details page is now displayed. Initially, after creation, the state of the VM cluster network is **Requires Validation**.

Maximum and recommended CIDR block prefix lengths for each Exadata system shape
The following table specifies the maximum and recommended CIDR block prefix lengths for each Exadata system shape. The maximum CIDR block prefix length defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth.

Max # Database Servers	Client and DR Subnet Size	Backup Subnet Size
5	/28 (16 IPs)	/29 (8 IPs)
13	/27 (32 IPs)	/28 (16 IPs)
29	/26 (64 IPs)	/27 (32 IPs)
32	/25 (128 IPs)	/26 (64 IPs)

Related Topics

- Resource Tags

Using the Console to Edit a VM Cluster Network

To edit a VM Cluster Network, use this procedure. You can only edit a VM cluster network that is not associated with a VM cluster.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to edit.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network that you want to edit.

   The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Edit VM Cluster Network**.
7. Use the **Edit** dialog to edit the VM cluster network attributes:
   a. **Client Network**

      The client network is the primary channel for application connectivity to Exadata Database Service on Cloud@Customer resources. You can edit the following client network settings:

      - **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the client network between `1` and `4094`, inclusive. To specify no VLAN tagging, enter "`1`". (This is equivalent to a "`NULL`" VLAN ID tag value.)

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> The values "0" and "4095" are reserved and cannot be entered.

- **Netmask:** Specify the IP netmask for the client network.
- **Gateway:** Specify the IP address of the client network gateway.
- **Domain Name:** Specify the domain name of the client network.
- **DB Server:** Edit the host names and IP addresses of the DB servers if applicable.

b. **SCAN Addresses:** Edit the SCAN IP addresses if applicable.

c. **Backup Network**

The backup network is the secondary channel for connectivity to Exadata Database Service on Cloud@Customer resources. It is typically used to segregate application connections on the client network from other network traffic. You can edit the following backup network settings:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the backup network between 1 and 4094, inclusive. To specify no VLAN tagging, enter "1". (This is equivalent to a "NULL" VLAN ID tag value.)

> **Note**
>
> The values "0" and "4095" are reserved and cannot be entered.

- **Netmask:** Specify the IP netmask for the disaster recovery network.
- **Domain Name:** Specify the domain name of the backup network.
- **DB Server:** Edit the host names and IP addresses of the DB servers if applicable.

d. **Disaster Recovery Network**

If so equipped (for authorized customers only), you can edit the following disaster recovery network settings:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the backup network between 1 and 4094, inclusive. To specify no VLAN tagging, enter "1". (This is equivalent to a "NULL" VLAN ID tag value.)

> **Note**
>
> The values "0" and "4095" are reserved and cannot be entered.

- **Netmask:** Specify the IP netmask for the disaster recovery network.
- **Domain Name:** Specify the domain name of the backup network.

e. **Configure DNS and NTP Servers**

The VM cluster network requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. You can edit the following settings:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the client network. You may specify up to three DNS servers.
- **NTP Servers:** Provide the IP address of an NTP server that is accessible using the client network. You may specify up to three NTP servers.

Copyright © 2008, 2026, Oracle and/or its affiliates.

8. Click **Save Changes**.

   After editing, the state of the VM cluster network is **Requires Validation**.

Using the Console to Download a File Containing the VM Cluster Network Configuration Details

To provide VM cluster network information to your network administrator, you can download and supply a file containing the network configuration.

Use this procedure to download a configuration file that you can supply to your network administrator. The file contains the information needed to configure your corporate DNS and other network devices to work along with Exadata Database Service on Cloud@Customer.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you are interested in.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network for which you want to download a file containing the VM cluster network configuration details.

   The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Download Network Configuration**.

   Your browser downloads a file containing the VM cluster network configuration details.

Using the Console to Validate a VM Cluster Network

You can only validate a VM cluster network if its current state is **Requires Validation**, and if the underlying Exadata infrastructure is activated.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to validate.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network that you want to validate.

   The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Validate VM Cluster Network**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Validation performs a series of automated checks on the VM cluster network. The Validate VM Cluster Network button is only available if the VM cluster network requires validation.

7. In the resulting dialog, click **Validate** to confirm the action.

   After successful validation, the state of the VM cluster network changes to **Validated** and the VM cluster network is ready to use. If validation fails for any reason, examine the error message and resolve the issue before repeating validation.

   If network validation fails, then a banner with the **Download Report** button is displayed. Download the report, edit the network resources assignment, and then try validating again.

Provisioning the First VM Cluster on an Oracle Exadata Database Service on Cloud@Customer System

Learn how to create the first VM cluster on your Oracle Exadata Database Service on Cloud@Customer system.

- About Managing VM Clusters on Oracle Exadata Database Service on Cloud@Customer
  The VM cluster provides a link between your Oracle Exadata Database Service on Cloud@Customer infrastructure and Oracle Databases you deploy.
- Prerequisites for VM Clusters on Oracle Exadata Database Service on Cloud@Customer
  To connect to the VM cluster virtual machine, you use an SSH public key.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create an Exascale VM Cluster
  To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

About Managing VM Clusters on Oracle Exadata Database Service on Cloud@Customer

The VM cluster provides a link between your Oracle Exadata Database Service on Cloud@Customer infrastructure and Oracle Databases you deploy.

The VM cluster contains an installation of Oracle Clusterware, which supports databases in the cluster. In the VM cluster definition, you also specify the number of enabled CPU cores, which determines the amount of CPU resources that are available to your databases

Before you can create any databases on your Exadata Cloud@Customer infrastructure, you must create a VM cluster network, and you must associate it with a VM cluster.

> **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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Prerequisites for VM Clusters on Oracle Exadata Database Service on Cloud@Customer

To connect to the VM cluster virtual machine, you use an SSH public key.

The public key is in OpenSSH format, from the key pair that you plan to use for connecting to the VM cluster virtual machines through SSH. The following shows an example of a public key, which is abbreviated for readability.

```
ssh-rsa AAAAB3NzaC1yc2EAAAABJQAA....lo/gKMLVM2xzc1xJr/Hc26biw3TXWGEakrK1OQ==
rsa-key-20160304
```

Related Topics

- Managing Key Pairs on Linux Instances

Using the Console to Create an ASM VM Cluster

To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Exadata VM Clusters**.
4. Click **Create Exadata VM Cluster**.
5. Provide the requested information on the Create Exadata VM Cluster page:
   a. **Provide the display name:** The display name is a user-friendly name that you can use to identify the VM cluster. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the VM cluster.
   b. **Select a compartment:** From the list of available compartments, choose the compartment that you want to contain the VM cluster.
   c. **Select an Exadata Infrastructure:** From the list, choose the Exadata infrastructure to host the VM cluster. You are not able to create a VM cluster without available and active Exadata infrastructure.
   d. **Select a VM Cluster Network:** From the list, choose a VM cluster network definition to use for the VM cluster. You must have an available and validated VM cluster network before you can create a VM cluster.
   e. **VM Cluster Type:**

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

   - **Exadata Database:** Standard Database VM with no restrictions, suitable for all workloads.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Exadata Database-Developer:** Developer Database VM with restrictions, suitable for application development only.

f. **Configure:**

- **DB Servers:**
  - Click **Change DB Servers** for VM placement to allocate VM resources.
  - On the **Change DB Servers** dialog:
    VM Cluster Type - Exadata Database: Select a minimum of one database server for VM placement. If you require a high availability database service that remains available during maintenance and unplanned outages, select at least two database servers. Maximum resources available for allocation per VM are based on the number of database servers selected.

    VM Cluster Type - Exadata Database-Developer: Select one database server for VM placement. Only one database server may be selected.

    > **Note**
    >
    > * DB Servers, which already have 8 VMs running on them are not available for selection.
    > * When calculating maximum local storage resources across selected DB Servers, the reserved local storage needed by the system to host a VM is deducted from the DB Server with the least resources.
    > For example, if the local storage available across selected DB servers is 823 GB for DB Server 3 and 813 GB for DB Server 4, then the minimum across selected servers is 813 GB and the maximum available for resource allocation is 813 GB - 184 GB (reserved local storage for hosting VM on X8M DB servers) = 629 GB.
    >
    > For more information, see *Estimating How Much Local Storage You Can Provision to Your VMs*.

  - Click **Save**.
- **Specify the OCPU (ECPUs for X11M) count per VM:** Specify the OCPU (ECPUs for X11M) count to be provisioned for each VM in this cluster. The minimum value is 2 OCPUs per VM or 8 ECPUs per VM for X11M (for a live VM condition), unless you are specifying zero OCPUs or 0 ECPUs for X11M (for a shutdown VM condition).
  If you specify a value of zero, then the VM cluster virtual machines are all shut down at the end of the cluster creation process. In this case, you can later start the virtual machines by scaling the OCPU (ECPUs for X11M) resources. See *Using the Console to Scale the Resources on a VM Cluster*.

  The OCPU (ECPUs for X11M) count for the whole VM Cluster will be calculated automatically based on the per-VM OCPU (ECPUs for X11M) count you have specified and the number of physical Database Servers configured for the VM Cluster.

  OCPU: An Oracle Compute Unit (OCPU) provides CPU capacity equivalent of one physical core of an Intel Xeon processor with hyperthreading enabled. Each OCPU corresponds to two hardware execution threads, known as vCPUs.

Copyright © 2008, 2026, Oracle and/or its affiliates.

See, *Oracle Platform as a Service and Infrastructure as a Service – Public Cloud Service DescriptionsMetered & Non-Metered*.

ECPU: An ECPU is an abstracted measure of compute resources. ECPUs are based on the number of cores elastically allocated from a pool of compute and storage servers.

- **Requested OCPU (ECPUs for X11M) count for the VM Cluster:** Displays the total number of CPU cores allocated to the VM cluster based on the value you specified in the **Specify the OCPU (ECPUs for X11M) count per VM** field. This field is not editable.
- **Specify the memory per VM (GB):** Specify the memory for each individual VM. The value must be a multiple of 1 GB and is limited by the available memory on the Exadata infrastructure.
- **Requested memory for the VM Cluster (GB):** Displays the total amount of memory allocated to the VM cluster based on the value you specified in the **Specify the memory per VM (GB)** field. This field is not editable.

g. **VM Storage:** This section lets you specify storage options for the VM image and VM backups. The VM image is the total storage needed by all file systems during cluster provisioning. VM backups are managed by Oracle and store the configured VM's backup data. Previously, both were stored on local database server storage. Now, you can select Exascale storage for the VM image, VM backups, or both.

> **Note**
>
> VM clusters using ASM for databases can also use Exascale for VM storage.

- **VM file system storage:** Specify the storage for the VM Image (VM file system storage).
  - **Database server (default):** Uses local database server storage for all VM file systems.
  - **Exascale storage:** Uses Exascale storage for all VM file systems.
- **VM file system storage capacity per VM (GB):** Displays the total storage required to provision the VM. This field is not editable. This value is calculated by adding the storage allocated to each guest VM file systems, such as `/u01`, `/u02`, `/(root)`, `/var` etc. To view current allocation details, click **Change file system configuration**. You can update the storage allocation for any file system by entering the new value within the displayed storage limits. Click **Update**. The VM file system storage capacity per VM (GB) field now displays the updated value.
- **Total VM file system storage capacity across VM cluster (GB):** Displays the total storage required to provision the VM cluster. This field is not editable. This value is calculated by multiplying the number of nodes by the VM file system storage required per VM.
- **VM backup storage:** Specify the storage for the Oracle-managed VM backups.
  - **Database server (default):** Uses local database server storage for VM backups.
  - **Exascale:** Uses Exascale storage for VM backups.
- **Backup storage per VM (GB):** Displays the total storage required for a single VM backup. This field is hidden, If you choose Exascale as the VM backup storage.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Backup storage across VM cluster:** Total backup storage across VM cluster (GB): If you choose Database server for VM backups, it displays the total storage required across the VM cluster for one VM backup whereas if you choose Exascale for VM backups, it displays the total storage required across the VM cluster for two VM backups.

h. **Exadata Storage:** The following settings define how the Exadata storage is configured for use with the VM cluster. The storage type once selected cannot be changed later on once the VM cluster is provisioned with the desired storage type. You have two options to choose: Automatic storage type (ASM) and Exascale. For more information about Exascale storage type, see *Using the Console to Create an Exascale VM Cluster*.

Automatic Storage Management (ASM)

- **Usable Exadata Storage (TB):** Specify the size for each individual VM. The minimum recommended size is 2 TB.
- **Storage allocation to ASM disk groups:** You have following options to choose:
  - **Allocate Oracle chosen default storage sizes:** Selecting this option allocates Oracle-defined default storage sizes to the DATA, RECO, and SPARSE ASM disk groups.
    * **Allocate storage for Exadata snapshots:** Check this option to create a sparse disk group, which is required to support Exadata snapshot functionality. Exadata snapshots enable space-efficient clones of Oracle databases that can be created and destroyed very quickly and easily.
    * **Allocate Storage for local backups:** Check this option to configure the Exadata storage to enable local database backups. If you select this option, more space is allocated to the RECO disk group to accommodate the backups. If you do not select this option, you cannot use local Exadata storage as a backup destination for any databases in the VM cluster.

Table 4-13 Storage Allocation

Storage Allocation	DATA Disk Group	RECO Disk Group	SPARSE Disk Group
Exadata Snapshots: **No** Enable Backups on Local Exadata Storage: **No**	80%	20%	0% (The SPARSE disk group is not created.)
Exadata Snapshots: **No** Enable Backups on Local Exadata Storage: **Yes**	40%	60%	0% (The SPARSE disk group is not created.)
Allocate Storage for Exadata Snapshots: **Yes** Enable Backups on Local Exadata Storage: **No**	60%	20%	20%

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 4-13 (Cont.) Storage Allocation

Storage Allocation	DATA Disk Group	RECO Disk Group	SPARSE Disk Group
Allocate Storage for Exadata Snapshots: **Yes** Enable Backups on Local Exadata Storage: **Yes**	35%	50%	15%

- **Allocate custom storage sizes:** Select this option to specify custom storage sizes for the DATA, RECO, and SPARSE ASM disk groups, subject to system-defined limits. These values override the Oracle-selected default allocations. The following minimum disk group size requirements apply:
  - DATA disk group: Minimum of 10% of total storage.
  - RECO disk group: Minimum of 10% of total storage.
  - SPARSE disk group: Minimum of 0%. If you specify a non-zero value, the minimum is 10%. The maximum is subject to Exadata system limits.

> **Note**
>
> After provisioning, you cannot reduce a disk group's size below its current used space plus a reserved capacity equal to 15% of the disk group's total size. This restriction applies in addition to the requirements described above.

i. **Select version:**

- **Choose the Oracle Grid Infrastructure release:** From the drop-down menu, choose either 19c or 26ai as the Grid Infrastructure release.

> **Note**
>
> The selected Oracle Grid Infrastructure release determines which Oracle Database versions are supported on the VM cluster.
>
> You cannot provision an Oracle Database version that is newer than the selected Grid Infrastructure version. For example, to provision a 26ai database, you must choose the 26ai Grid Infrastructure release.

> **Note**
>
> To provision a VM cluster with Grid Infrastructure 26ai, the following minimum software versions are required:
>
> - Exadata Guest VM running Exadata System Software 23.1.8
> - Exadata Infrastructure running Exadata System Software 23.1.x

Copyright © 2008, 2026, Oracle and/or its affiliates.

- The **Display All Available Grid Infrastructure Versions** switch is OFF by default. When it is OFF, the system automatically uses the latest Grid Infrastructure version during VM cluster provisioning.
  When toggled ON, the drop-down menu displays the latest four supported Grid Infrastructure versions, allowing you to select your preferred version from the list.
- **Choose an Exadata guest version:**
  - **Exadata infrastructure with Oracle Linux 7 and Exadata image version 22.1.10.0.0.230422:**
    - The **Change image** button is not enabled.
    - The Oracle Grid Infrastructure version defaults to 19.0.0.0.0.
    - The Exadata guest version will be the same as that of the host OS.
  - **Exadata infrastructure with Oracle Linux 8 and Exadata image version 23.1.3.0.0.230613:**
    - The Exadata guest version defaults to the latest (23.1.3.0).
    - The Oracle Grid Infrastructure version defaults to 19.0.0.0.0
    - The **Change image** button is enabled.
    - Click **Change image**.
      The resulting Change image panel displays the list of available major versions of Exadata image (23.1.3.0 and 22.1.3.0).

      The most recent release for each major version is indicated by "(latest)".
    - Slide **Display all available versions**.
      Six past versions including the latest versions of Exadata images 23.1.3.0 and 22.1.3.0 are displayed.
    - Choose a version.
    - Click **Save Changes**.

j. **Add SSH Key:** Specify the public key portion of an SSH key pair that you want to use to access the VM cluster virtual machines. You can upload a file containing the key, or paste the SSH key string.

   To provide multiple keys, upload multiple key files or paste each key into a separate field. For pasted keys, ensure that each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.

k. **Choose a license type:**

- **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.

> **Note**
>
> BYOL is not available for Exadata Database-Developer VM Cluster type.

- **License Included:** Select this option to subscribe to Oracle Database software licenses as part of Exadata Database Service on Cloud@Customer.

l. **Diagnostics Collection:**

   By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and

Copyright © 2008, 2026, Oracle and/or its affiliates.

effectively. Subscribe to Events to get notified about resource state changes. For more information, see *Getting Started with Events*.

> **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.

- **Enable Diagnostic Events:** Allow Oracle to collect and publish critical, warning, error, and information events to me.
- **Enable Health Monitoring:** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
- **Enable Incident Logs and Trace Collection:** Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

  All three checkboxes are selected by default. You can leave the default settings as is or clear the checkboxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information** >> **Diagnostics Collection**.

  - **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
  - **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).
  - **Partially Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).

m. **Show Advanced Options:**

- **Time zone:** The default time zone for the Exadata Infrastructure is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

  > **Note**
  >
  > If you want to set a time zone other than UTC or the browser-detected time zone, then select the **Select another time zone** option, select a **Region** or **country**, and then select the corresponding **Time zone**.
  >
  > If you do not see the region or country you want, then select **Miscellaneous**, and then select an appropriate **Time zone**.

- **Cloud Automation Update:** Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster. Set the start time for cloud automation updates.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable early access for cloud tools update: VM clusters designated for early access receive updates 1-2 weeks before they are available to other systems. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period: Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

> **Note**
>
> – The freeze period can extend for a maximum of 45 days from the start date.
>
> – Oracle automation will automatically apply updates with critical security fixes (CVSS >= 9) even during a configured freeze period.

- **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.

6. Optionally, you can save the resource configuration as a stack.
   - **To save the resource configuration as a Stack:**
     a. Click **Save as Stack**.
     b. In the resulting Save as Stack dialog, provide the following details:
        i. **Name:** (Optional) Provide an easy to remember descriptive name.
        ii. **Description:** (Optional) Enter a short description.
        iii. **Compartment:** Select a compartment where this Stack will reside.
        iv. **Tags:** Add tags.
     c. Click **Save**.
        After saving the Stack, the system displays a banner with a link to the saved Stack.
     d. Click the link to open the Stack in the Resource Manager Service console. See, *Resource Manager and Terraform*.
   - **To view the details of a Stack:**
     a. Open the navigation menu. Under **Developer Services**, click **Resource Manager**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

b. Click **Stacks**.

c. Click the name of the Stack that you want to view details.
Or, click the Actions menu (three dots), and select the **View stack details** option.

7. Click **Create VM Cluster**.

The VM Cluster Details page is now displayed. While the creation process is running, the state of the VM cluster is **Pending**. When the VM cluster creation process completes, the state of the VM cluster changes to **Available**.

The Exadata Database Storage section on the VM Cluster Details page shows the type of storage configured, which, in this case, is ASM.

Related Topics

- Oracle Exadata Database Service on Cloud@Customer Service Description
  Learn how you can leverage the combined capabilities of Oracle Exadata and Oracle Cloud Infrastructure with Oracle Exadata Database Service on Cloud@Customer
- Using the Console to Scale the Resources on a VM Cluster
  Starting in Oracle Exadata Database Service on Cloud@Customer, you can scale up or down multiple resources at the same time. You can also scale up or down resources one at a time.
- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.
- Estimating How Much Local Storage You Can Provision On Your VMs
- Resource Tags
- Oracle PaaS/IaaS Cloud Service Description documents
- Oracle Platform as a Service and Infrastructure as a Service – Public Cloud Service DescriptionsMetered & Non-Metered
- Getting Started with Events
- Overview of Database Service Events
- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Resource Manager and Terraform

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Create an Exascale VM Cluster

To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

To create an Exascale VM cluster, ensure that you that have:

- Active Exadata infrastructure available to host the VM cluster.
- A validated VM cluster network available for the VM cluster to use.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Exadata VM Clusters**.
4. Click **Create Exadata VM Cluster**.
5. Provide the requested information on the Create Exadata VM Cluster page:
   a. **Select a compartment:** From the list of available compartments, choose the compartment that you want to contain the VM cluster.
   b. **Provide the display name:** The display name is a user-friendly name that you can use to identify the VM cluster. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the VM cluster.
   c. **Select Exadata Infrastructure:** From the list, choose the Exadata infrastructure to host the VM cluster. You are not able to create a VM cluster without available and active Exadata infrastructure.
   d. **Select a VM Cluster Network:** From the list, choose a VM cluster network definition to use for the VM cluster. You must have an available and validated VM cluster network before you can create a VM cluster.
   e. **VM Cluster Type**

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

      - **Exadata Database:** Standard Database VM with no restrictions, suitable for all workloads.
      - **Exadata Database-Developer:** Developer Database VM with restrictions, suitable for application development only.

   f. **Configure:**
      - **DB Servers:**
        – Click **Change DB Servers** for VM placement to allocate VM resources.
        – On the **Change DB Servers** dialog:
        VM Cluster Type - Exadata Database: Select a minimum of one database server for VM placement. If you require a high availability database service that remains available during maintenance and unplanned outages, select at

Copyright © 2008, 2026, Oracle and/or its affiliates.

least two database servers. Maximum resources available for allocation per VM are based on the number of database servers selected.

VM Cluster Type - Exadata Database-Developer: Select one database server for VM placement. Only a single database server may be selected.

> **Note**
>
> * DB Servers, which already have 8 VMs running on them are not available for selection.
> * When calculating maximum local storage resources across selected DB Servers, the reserved local storage needed by the system to host a VM is deducted from the DB Server with the least resources.
> For example, if the local storage available across selected DB servers is 823 GB for DB Server 3 and 813 GB for DB Server 4, then the minimum across selected servers is 813 GB and the maximum available for resource allocation is 813 GB - 184 GB (reserved local storage for hosting VM on X8M DB servers) = 629 GB.
>
> For more information, see *Estimating How Much Local Storage You Can Provision to Your VMs*.

– Click **Save**.

- **Specify the OCPU (ECPUs for X11M) count per VM:** Specify the OCPU (ECPUs for X11M) count to be provisioned for each VM in this cluster. The minimum value is 2 OCPUs per VM or 8 ECPUs per VM for X11M (for a live VM condition), unless you are specifying zero OCPUs or 0 ECPUs for X11M (for a shutdown VM condition).
If you specify a value of zero, then the VM cluster virtual machines are all shut down at the end of the cluster creation process. In this case, you can later start the virtual machines by scaling the OCPU (ECPUs for X11M) resources. See *Using the Console to Scale the Resources on a VM Cluster*.

  The OCPU (ECPUs for X11M) count for the whole VM Cluster will be calculated automatically based on the per-VM OCPU (ECPUs for X11M) count you have specified and the number of physical Database Servers configured for the VM Cluster.

  OCPU: An Oracle Compute Unit (OCPU) provides CPU capacity equivalent of one physical core of an Intel Xeon processor with hyperthreading enabled. Each OCPU corresponds to two hardware execution threads, known as vCPUs.

  See, *Oracle Platform as a Service and Infrastructure as a Service – Public Cloud Service DescriptionsMetered & Non-Metered*.

  ECPU: An ECPU is an abstracted measure of compute resources. ECPUs are based on the number of cores elastically allocated from a pool of compute and storage servers.

- **Requested OCPU (ECPUs for X11M) count for the VM Cluster:** Displays the total number of CPU cores allocated to the VM cluster based on the value you specified in the **Specify the OCPU (ECPUs for X11M) count per VM** field. This field is not editable.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Specify the memory per VM (GB):** Specify the memory for each individual VM. The value must be a multiple of 1 GB and is limited by the available memory on the Exadata infrastructure.
- **Requested memory for the VM Cluster (GB):** Displays the total amount of memory allocated to the VM cluster based on the value you specified in the **Specify the memory per VM (GB)** field. This field is not editable.

g. **VM Storage:** This section lets you specify storage options for the VM image and VM backups. The VM image is the total storage needed by all file systems during cluster provisioning. VM backups are managed by Oracle and store the configured VM's backup data. Previously, both were stored on local database server storage. Now, you can select Exascale storage for the VM image, VM backups, or both.

> **Note**
>
> VM clusters using ASM for databases can also use Exascale for VM storage.

- **VM file system storage:** Specify the storage for the VM Image (VM file system storage).
  - **Database server (default):** Uses local database server storage for all VM file systems.
  - **Exascale storage:** Uses Exascale storage for all VM file systems.
- **VM file system storage capacity per VM (GB):** Displays the total storage required to provision the VM. This field is not editable. This value is calculated by adding the storage allocated to each guest VM file systems, such as `/u01`, `/u02`, `/(root)`, `/var` etc. To view current allocation details, click **Change file system configuration**. You can update the storage allocation for any file system by entering the new value within the displayed storage limits. Click **Update**. The VM file system storage capacity per VM (GB) field now displays the updated value.
- **Total VM file system storage capacity across VM cluster (GB):** Displays the total storage required to provision the VM cluster. This field is not editable. This value is calculated by multiplying the number of nodes by the VM file system storage required per VM.
- **VM backup storage:** Specify the storage for the Oracle-managed VM backups.
  - **Database server (default):** Uses local database server storage for VM backups.
  - **Exascale:** Uses Exascale storage for VM backups.
- **Backup storage per VM (GB):** Displays the total storage required for a single VM backup. This field is hidden, If you choose Exascale as the VM backup storage.
- **Backup storage across VM cluster:** Total backup storage across VM cluster (GB): If you choose Database server for VM backups, it displays the total storage required across the VM cluster for one VM backup whereas if you choose Exascale for VM backups, it displays the total storage required across the VM cluster for two VM backups.

h. **Storage type:** When creating a VM cluster, you must choose storage type for the database. You can select one of two options: Automatic Storage Management (ASM) or Exascale. This choice is permanent and cannot be changed after the VM is provisioned. For more information about ASM storage type, see *Using the Console to Create an ASM VM Cluster*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Minimum Requirements for Configuring Exascale Storage
>
> - Supported on Exadata Infrastructure Model X8M or later.
> - Requires Exadata System Software release 24.1 or later.
> - Requires Oracle Grid Infrastructure version 26ai (24.3)
> - Supports Oracle Database versions 26ai (23.4) or later.
> - Storage servers (cells) must be running release 25.1.9 or later, or 25.2.2 or later.
>   If storage servers are running release 25.2.1, the Exascale feature must be disabled.
>
> Exascale option will be disabled If the minimum requirement is not met.

Exascale database storage vault:

- **Create new storage vault:** Choose this option to create a new Exascale database storage vault during VM cluster provisioning.
  - **Storage vault name:** Enter a descriptive name for the vault. Click the **change compartment** link and choose a compartment if you want to create this vault in a different compartment.
  - **Storage capacity for databases:** Enter the storage capacity for the databases within the minimum and maximum values displayed on the screen.

    > **Note**
    >
    > If additional space is needed beyond the maximum shown, the Exascale capacity must be increased. For more information, see *Using the Console to Scale an Exascale Storage Vault*.

- **Select existing storage vault:** Select a vault that resides in the compartment of your choice.

i. **Select version:**

> **Note**
>
> Only Oracle Database 26ai can be provisioned on the Exascale VM cluster.

- **Choose the Oracle Grid Infrastructure release:** The Oracle Grid Infrastructure release defaults to 26ai.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> The selected Oracle Grid Infrastructure release determines which Oracle Database versions are supported on the VM cluster.
>
> You cannot provision an Oracle Database version that is newer than the selected Grid Infrastructure version. For example, to provision a 26ai database, you must choose the 26ai Grid Infrastructure release.

> **Note**
>
> To provision a VM cluster with Grid Infrastructure 26ai, the following minimum software versions are required:
>
> – Exadata Guest VM running Exadata System Software 23.1.8
> – Exadata Infrastructure running Exadata System Software 23.1.x

- The **Display All Available Grid Infrastructure Versions** switch is OFF by default. When it is OFF, the system automatically uses the latest Grid Infrastructure version during VM cluster provisioning.
  When toggled ON, the drop-down menu displays the latest four supported Grid Infrastructure versions, allowing you to select your preferred version from the list.
- **Choose an Exadata guest version:**
  - The Exadata guest version defaults to the latest (24.1.6.0)
  - The Oracle Grid Infrastructure version defaults to 26ai
  - The **Change image** button is enabled.
  - Click **Change image**.
    The resulting Change image panel displays the list of available major versions of Exadata image (24.1.6.0 and later).

    The most recent release for each major version is indicated by "(latest)".
  - Slide **Display all available versions**.
    Six past versions including the latest versions of Exadata images 24.1.6.0 and later are displayed.
  - Choose a version.
  - Click **Save Changes**.

j. **Add SSH Key:** Specify the public key portion of an SSH key pair that you want to use to access the VM cluster virtual machines. You can upload a file containing the key, or paste the SSH key string.

   To provide multiple keys, upload multiple key files or paste each key into a separate field. For pasted keys, ensure that each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.

k. **Choose a license type:**

   - **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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> BYOL is not available with Exadata Database-Developer VM Cluster type.

- **License Included:** Select this option to subscribe to Oracle Database software licenses as part of Exadata Database Service on Cloud@Customer.

l. **Diagnostics Collection:**

By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes. For more information, see *Getting Started with Events*.

> **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.

- **Enable Diagnostic Events:** Allow Oracle to collect and publish critical, warning, error, and information events to me.
- **Enable Health Monitoring:** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
- **Enable Incident Logs and Trace Collection:** Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

  All three checkboxes are selected by default. You can leave the default settings as is or clear the checkboxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information** >> **Diagnostics Collection**.

  - **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
  - **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).
  - **Partially Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).

m. **Show Advanced Options:**

- **Time zone:** The default time zone for the Exadata Infrastructure is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, then select the **Select another time zone** option, select a **Region** or **country**, and then select the corresponding **Time zone**.
>
> If you do not see the region or country you want, then select **Miscellaneous**, and then select an appropriate **Time zone**.

- **Cloud Automation Update:** Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster. Set the start time for cloud automation updates.

> **Note**
>
> Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable early access for cloud tools update: VM clusters designated for early access receive updates 1-2 weeks before they are available to other systems. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period: Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

> **Note**
>
> – The freeze period can extend for a maximum of 45 days from the start date.
>
> – Oracle automation will automatically apply updates with critical security fixes (CVSS >= 9) even during a configured freeze period.

- **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.

6. Optionally, you can save the resource configuration as a stack.
   - **To save the resource configuration as a Stack:**
     a. Click **Save as Stack**.
     b. In the resulting Save as Stack dialog, provide the following details:

Copyright © 2008, 2026, Oracle and/or its affiliates.

i. **Name:** (Optional) Provide an easy to remember descriptive name.

ii. **Description:** (Optional) Enter a short description.

iii. **Compartment:** Select a compartment where this Stack will reside.

iv. **Tags:** Add tags.

c. Click **Save**.
After saving the Stack, the system displays a banner with a link to the saved Stack.

d. Click the link to open the Stack in the Resource Manager Service console.
See, *Resource Manager and Terraform*.

- **To view the details of a Stack:**

a. Open the navigation menu. Under **Developer Services**, click **Resource Manager**.

b. Click **Stacks**.

c. Click the name of the Stack that you want to view details.
Or, click the Actions menu (three dots), and select the **View stack details** option.

7. Click **Create VM Cluster**.

The VM Cluster Details page is now displayed. While the creation process is running, the state of the VM cluster is **Pending**. When the VM cluster creation process completes, the state of the VM cluster changes to **Available**.

The Exadata Database Storage section on the VM Cluster Details page shows the type of storage configured, which, in this case, is Exascale.

Related Topics

- Oracle Exadata Database Service on Cloud@Customer Service Description
  Learn how you can leverage the combined capabilities of Oracle Exadata and Oracle Cloud Infrastructure with Oracle Exadata Database Service on Cloud@Customer
- Using the Console to Scale the Resources on a VM Cluster
  Starting in Oracle Exadata Database Service on Cloud@Customer, you can scale up or down multiple resources at the same time. You can also scale up or down resources one at a time.
- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.
- Estimating How Much Local Storage You Can Provision On Your VMs
- Resource Tags
- Oracle PaaS/IaaS Cloud Service Description documents
- Oracle Platform as a Service and Infrastructure as a Service – Public Cloud Service DescriptionsMetered & Non-Metered
- Getting Started with Events
- Overview of Database Service Events
- 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 © 2008, 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 Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Resource Manager and Terraform

Creating Database Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

Oracle Exadata Database Service on Cloud@Customer provides a backup facility, which you can configure individually on each database.

To store database backups on a Recovery Appliance or on a network file storage (NFS) location that you manage, then you must first create a backup destination.

- About Managing Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  For backups, you can either use the Oracle Exadata Database Service on Cloud@Customer backup facility, or you can configure a backup location on a location you manage.
- Prerequisites for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  To configure backup destinations on a Zero Data Loss Recovery Appliance location, or an NFS backup location, review the prerequisites.
- Using the Console for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  Learn how to use the console to create, edit, move, and terminate a backup destination for your infrastructure for Oracle Exadata Database Service on Cloud@Customer.

Related Topics

- Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer
  Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Cloud@Customer.

About Managing Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

For backups, you can either use the Oracle Exadata Database Service on Cloud@Customer backup facility, or you can configure a backup location on a location you manage.

Oracle Exadata Database Service on Cloud@Customer provides a backup facility, which you can configure individually on each database.

Copyright © 2008, 2026, Oracle and/or its affiliates.

See: *Managing Databases on Exadata Cloud@Customer* and *Managing Database Backup and Recovery on Exadata Cloud@Customer*.

If you want to store backups on a Recovery Appliance, or on a network file storage (NFS) location that you manage, then you must first create a backup destination. Each backup destination defines the properties that are required to connect to the Recovery Appliance or NFS location, and each backup destination must be accessible in your data center from the VM cluster nodes.

The Oracle Exadata Database Service on Cloud@Customer backup facility can also store backups on Oracle Cloud Infrastructure object storage, or on local Exadata storage on your Oracle Exadata Database Service on Cloud@Customer system. However, you do not need to create a backup destination for any of these other locations. Instead, applicable options for backup to cloud object storage or local Exadata storage are available directly when you create a database.

> **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.

Related Topics

- Zero Data Loss Recovery Appliance
- Manage Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems
  Learn to manage Oracle Database homes on Oracle Exadata Database Service on Cloud@Customer.
- Using the Console to Create a Backup Destination
  To create a backup destination, be prepared to provide values for the backup destination configuration.
- Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer
  Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Cloud@Customer.

Prerequisites for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

To configure backup destinations on a Zero Data Loss Recovery Appliance location, or an NFS backup location, review the prerequisites.

> **Note**
>
> Starting with dbaascli release 25.3.1.0.0, the Data Guard broker is mandatory when configuring Zero Data Loss Recovery Appliance as a backup destination in Data Guard-enabled databases.

- For a Zero Data Loss Recovery Appliance backup destination:

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - The appliance must be configured with a virtual private catalog (VPC) user, which is used for taking the backups.
  - The appliance must be configured with the unique database name of the database being backed up, and a mapping to the VPC user.
  - The appliance must be accessible from the Oracle Exadata Database Service on Cloud@Customer system using the Oracle Net Services connection string, which is provided by the Zero Data Loss Recovery Appliance administrator.
- For an NFS backup destination:
  - Oracle Exadata Database Service on Cloud@Customer non-autonomous databases:
    * You must mount the NFS server location to a local mount point directory on each node in the VM cluster.
    * The local mount point directory and the NFS server must be identical across all nodes in the cluster.
    * You must ensure that the NFS mount is maintained continuously on all of the VM cluster nodes.
    * The NFS-mounted file system must be readable and writable by the `oracle` operating system user on all of the VM cluster nodes.
  - Oracle Autonomous AI Database on Exadata Cloud@Customer:

    > **Note**
    >
    > Oracle Autonomous AI Database on Exadata Cloud@Customer supports NFS versions 3 and 4. Mounts with both versions enabled, however, default to version 4.

    * To ensure that the Autonomous VM cluster can access the NFS server over the (preferred) Backup Network, enter valid Backup Network IP addresses while configuring the VM Cluster Network.
      If the backup destination is not accessible over the backup network but accessible over the client network, the backup traffic will be routed over the client network.
    * The NFS-mounted file system must be readable and writable by the `oracle` operating system user on all of the VM cluster nodes.
    * If permissions are being controlled at the user level, then the `uid:gid` of the `oracle` user for the Autonomous VM cluster is 1001:1001.
    * For NFSv4 servers that restrict port range, you must use the insecure option to enable clients other than an Administrator user to connect to the NFS server.

Using the Console for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

Learn how to use the console to create, edit, move, and terminate a backup destination for your infrastructure for Oracle Exadata Database Service on Cloud@Customer.

- Using the Console to Create a Backup Destination
  To create a backup destination, be prepared to provide values for the backup destination configuration.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Edit a Backup Destination
  To edit a backup destination, be prepared to provide values for the backup destination configuration.
- Using the Console to Move a Backup Destination to Another Compartment
  To move a backup destination, be prepared to provide values for the backup destination configuration.
- Using the Console to Delete a Backup Destination
  To delete a backup destination, be prepared to provide values for the backup destination configuration.

Using the Console to Create a Backup Destination

To create a backup destination, be prepared to provide values for the backup destination configuration.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Backup Destinations**.
4. Click **Create Backup Destination**.
5. Provide the requested information in the **Create Backup Destination** page:
   a. Choose a compartment.

      From the list of available compartments, choose the compartment that you want to contain the backup destination.

   b. Name your backup destination.

      Specify a user-friendly name that you can use to identify the backup destination. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the backup destination.

   c. Choose either a Zero Data Loss Recovery Appliance or a network file system (NFS) backup destination.

      > **Note**
      >
      > You can also set OCI Object Store as a backup destination. However, you cannot set it from this screen. You can configure OCI Object Store as a backup destination when creating a database. For more information, see *Backup Destination Type* in *Using the Console to Create a Database*.

      Select **Recovery Appliance** or **Network Storage (NFS)**.

      - If you select **Recovery Appliance**, then you must also specify the following for Zero Data Loss Recovery Appliance:
        – **Provide the Recovery Appliance connection string:** Specify the Oracle Net Services connection string that connects to the appliance. This information is typically provided by the Zero Data Loss Recovery Appliance administrator.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> ADB customers must ensure to provide the connection string that connects to the recovery appliance in an Oracle "easy connect" string format, as shown below:
>
> *host*:*port*/*service name*
>
> Where *host* is the SCAN hostname of the Zero Data Loss Recovery Appliance.

  - **Provide the Virtual Private Catalog (VPC) Users:** Provide a VPC user name for connecting to the Zero Data Loss Recovery Appliance. You can specify multiple VPC user names in case you want to use the appliance as a backup destination for multiple databases. This information is typically provided by the Zero Data Loss Recovery Appliance administrator.

- If you select **Network Storage (NFS)**, then you must also specify the following:
  - **Self-mount for non-autonomous databases:**

    Provide the local NFS mount point path: Specify the local directory path on each VM cluster node where the NFS server location is mounted. The local directory path and the NFS server location must each be the same across all of the VM cluster nodes.
  - **Auto-mount for Autonomous Databases:**

    Use this destination for Autonomous Databases:
    - **NFS server:** Specify the IP address of the NFS server. Optionally, you can specify up to four IP addresses. These are just extra paths to the same NFS server. Note that Oracle supports only one NFS device for backups.
    - **NFS export share:** Specify the directory path where the exported file system is mounted.

d. Configure Advanced Options.

- **Tags:** (Optional) You can choose to apply 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, refer to information about 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.

6. Click **Create Backup Destination**.

   The Backup Destination Details page displays the newly created backup destination.

Related Topics

- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.
- Resource Tags

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Edit a Backup Destination

To edit a backup destination, be prepared to provide values for the backup destination configuration.

You can only edit a backup destination if it is not currently associated with database.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to edit.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to edit.

   The Backup Destination Details page displays information about the selected backup destination.
5. Click **Edit**.
6. Use the **Edit Backup Destination** dialog to edit the backup destination attributes:

   > **Note**
   >
   > You cannot edit a Backup Destination if there is already a database attached to it.

   - If you are editing a Zero Data Loss Recovery Appliance backup destination:
     - **Provide the Recovery Appliance connection string:** Specify the Oracle Net Services connection string that connects to the Recovery Appliance. This information is typically provided by the Recovery Appliance administrator.
     - **Provide the Virtual Private Catalog (VPC) Users:** Provide a VPC user name for connecting to the Recovery Appliance. You can specify multiple VPC user names in case you want to use the Recovery Appliance as a backup destination for multiple databases. This information is typically provided by the Recovery Appliance administrator.
   - If you are editing an NFS backup destination:
     - **Self-mount for non-autonomous databases:**

       Provide the local NFS mount point path: Specify the local directory path on each VM cluster node where the NFS server location is mounted. The local directory path and the NFS server location must each be the same across all of the VM cluster nodes.
     - **Auto-mount for Autonomous Databases:**

       Use this destination for Autonomous Databases:

       * **NFS server:** Specify the IP address of the NFS server. Optionally, you can specify up to four IP addresses.
       * **NFS export share:** Specify the directory path where the exported file system is mounted.
7. Click **Save Changes**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Move a Backup Destination to Another Compartment

To move a backup destination, be prepared to provide values for the backup destination configuration.

You can change the compartment that contains your backup destination by moving it.

When you move a backup destination, the compartment change does not affect other associated resources. These other resources, such as the associated databases, remain in their current compartment.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to move.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to move.

   The Backup Destination Details page displays information about the selected backup destination.
5. Click **Move Resource**.
6. In the resulting dialog, choose the new compartment for the backup destination and click **Move Resource**.

Using the Console to Delete a Backup Destination

To delete a backup destination, be prepared to provide values for the backup destination configuration.

Before you can delete a backup destination, you must ensure that it is not associated with any databases.

Deleting a backup destination:

- Does not remove any residual backups that are left in the backup destination
- Removes all references to the deleted backup destination from the Cloud Control Plane

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to delete.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to delete.

   The Backup Destination Details page displays information about the selected backup destination.
5. Click **Delete**.
6. In the resulting dialog, enter the backup destination name and click **Delete Backup Destination** to confirm the action.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Creating First Database Home on an Oracle Exadata Database Service on Cloud@Customer System

After provisioning a VM Cluster, create your first Oracle Database Home on Exadata Cloud@Customer.

- About Creating Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer 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.
- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- Using the API to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home, review the list of API calls.

About Creating Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer 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 with Oracle Exadata Database Service on Cloud@Customer
  Learn to use the `dbaascli` utility on Oracle Exadata Database Service on Cloud@Customer.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer

To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

> **Note**
>
> Minimum requirements for provisioning a Database 26ai home:
>
> - Grid Infrastructure 26ai
> - Exadata Guest VM running Exadata System Software 23.1.8

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

   VM Clusters is selected by default.
2. Choose your Compartment.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the VM cluster on which you want to create the Database Home.
4. Under **Resources**, click **Database Homes**.
5. Click **Create Database Home**.
6. In the Create Database Home dialog, enter the following:
   - **Database Home display name**: The display name for the Database Home.
   - **Unified Auditing:** Select this check box to enable Unified Auditing framework.

     > **Note**
     >
     > You cannot disable Unified Auditing after provisioning the Database Home.

     - **For Oracle Database versions lower than 12.1:** You cannot use the Unified Auditing framework and should instead use the Traditional Audit, the legacy Oracle Database audit framework.
     - **For Oracle Database versions 12.1 or higher:** You can enable Unified Auditing from the OCI Console. For Oracle Database versions 12.1 or higher but lower than version 26ai, the **Unified Auditing** check box is not selected by default. However, it is selected by default for Oracle AI Database 26ai.
       Unified Auditing field in the **General Information** section on the **Database Home Details** page displays if Unified Auditing is **Enabled** or **Disabled**.
   - **Database image**: Determines what Oracle Database version is used for the database. You can mix database versions on the Exadata VM Cluster, but not editions. By default, the latest Oracle-published database software image is selected.

Copyright © 2008, 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 choose one of the options below on the **Select a Database Software Image** window.

Oracle Provided Database Software Images: These images contain generally available versions of Oracle Database software.

Custom Database Software Images: These images are created by your organization and contain customized configurations of software updates and patches.

Use the **Select a compartment** and **Select a Database version** selectors to limit the list of custom Database Software Images to a specific compartment or Oracle Database software major release version.

Use the **Region** filter to access the software images created in a different region.

After choosing a software image, click **Select** to return to the Create Database Home dialog.

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

A Database Software Image will not be available for Database Home creation if:

- The database version of Database Software Image is out of support. For example, Database Software Images created using 11.2.0.4 will not be available for Database Home provisioning after 31-Dec-2022.
- The Exadata model should support the PSU/RU version of the Database Software Image. For example, for the 19c release, the X8M-2 model supports RU version 19.4 and greater.
- Only Database Software Images created specifically in the context of Exadata Cloud@Customer service can be used while provisioning and patching Database Homes within the Exadata Cloud@Customer service.
- The Database Software Image is not in `Available` state, that is, `Deleted` or is being `Updated`.

- **Show Advanced Options**

  You have the option to configure advanced options.

  - **Tags:** (Optional) You can choose to apply 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 if you should apply tags, then skip this option (you can apply tags later) or ask your administrator.

Note that after Home install, patch to the latest if the latest patch is available.

7. Click **Create**.

   When the Database Home creation is complete, the status changes from `Provisioning` to `Available`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Resource Tags

Using the API to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer

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 Cloud@Customer, 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

Creating First Database on an Oracle Exadata Database Service on Cloud@Customer System

After provisioning an Oracle Database Home and any needed backup destinations, you are ready to create your first database on Oracle Exadata Database Service on Cloud@Customer.

- Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer
  Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer.
- Oracle Database Releases Supported by Oracle Exadata Database Service on Cloud@Customer
  Learn about the versions of Oracle Database that Oracle Exadata Database Service on Cloud@Customer supports.
- About Provisioning and Configuring Oracle Databases on Oracle Exadata Database Service on Cloud@Customer
  Learn about provisioning and configuring Oracle Database on Oracle Exadata Database Service on Cloud@Customer
- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer

Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer.

Before you can create and use an Oracle Database on Exadata Database Service on Cloud@Customer, you must:

- Provision Exadata Database Service on Cloud@Customer infrastructure
- Configure a VM cluster
- Create any required backup destinations

You can create one or more databases on each Oracle Exadata Database Service on Cloud@Customer system. Other than the storage and processing limits of your Oracle Exadata system, there is no maximum for the number of databases that you can create. By default, databases on Exadata Database Service on Cloud@Customer use Oracle Database Enterprise Edition - Extreme Performance. This edition provides all the features of Oracle Database Enterprise Edition, plus all of the database enterprise management packs, and all of the Enterprise Edition options, such as Oracle Database In-Memory, and Oracle Real Application Clusters (Oracle RAC). If you use your own Oracle Database licenses, then your ability to use various features is limited by your license holdings. TDE Encryption is required for all cloud databases. All new tablespaces will automatically be enabled for encryption.

> **Note**
>
> Custom locale objects (such as language, territory, character set, and collation) require deploying custom locale data files to both the database and storage servers, which is not supported in Exadata Database Service on Cloud@Customer.

Oracle Database Releases Supported by Oracle Exadata Database Service on Cloud@Customer

Learn about the versions of Oracle Database that Oracle Exadata Database Service on Cloud@Customer supports.

Exadata Database Service on Cloud@Customer supports the following Oracle Database software releases:

- Oracle AI Database 26ai
- Oracle Database 19c (19.x)
- Oracle Database 12c Release 2 (12.2.0.1) (requires a valid Upgrade Support contract)
- Oracle Database 12c Release 1 (12.1.0.2) (requires a valid Upgrade Support contract)
- Oracle Database 11g Release 2 (11.2.0.4) (requires a valid Upgrade Support contract)

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

Related Topics

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

About Provisioning and Configuring Oracle Databases on Oracle Exadata Database Service on Cloud@Customer

Learn about provisioning and configuring Oracle Database on Oracle Exadata Database Service on Cloud@Customer

Each Oracle Database is configured as follows:

- When you provision a database, you can associate it with a backup destination, and enable automatic backups.
- When a database is provisioned an archivelog maintenance job is added to the `crontab` for the database.
  - If the database is not enabled for backups, then the archivelog job will maintain FRA space by deleting Archive Redo Logs older than 24 hours.
  - If the database is enabled for backups, then the archivelog job will backup archivelogs that have not been backed up. Once an archived log is backed up, it will be purged when older than 24 hours.
- Each database is configured with Oracle Real Application Clusters (Oracle RAC) database instances running on every node in the virtual machine (VM) cluster.
- Each database is created in an Oracle home, which uses a separate set of Oracle binaries in a separate Oracle home location.
- Each database is configured with default instance parameter settings. While the defaults are reasonable for many cases, you should review the instance parameter settings to ensure that they meet your specific application needs.
  In particular, review the Oracle Database system global area (SGA) and program global area (PGA) instance parameter settings, especially if your VM cluster supports multiple databases. Also, ensure that the sum of all Oracle Database memory allocations never exceeds the available physical memory on each virtual machine.
  - 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.
  - The database initialization parameter `INMEMORY_FORCE` is set to `CELLMEMORY_LEVEL` for all 19.8 and later databases created via the cloud automation. This setting will enable the Exadata Columnar Cache feature, which dramatically speeds up analytic queries. It is available for 19.8 and later databases and no In Memory license is required when running in Exadata Cloud. For more information, see INMEMORY_FORCE
- Exadata Database Service will only create databases with 8K block size. This parameter cannot be changed.
- Each database using Oracle Database 12c Release 1 or a later release is configured as a container database (CDB). One pluggable database (PDB) is created inside the CDB. By default:
  - The first PDB is configured with a local PDB administration user account, named `PDBADMIN`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- The PDBADMIN user account is initially configured with the same administration password as the CDB SYS and SYSTEM users.
- The PDBADMIN user account is initially configured with basic privileges assigned through two roles; CONNECT and PDB_DBA. However, for most practical administrative purposes you must assign extra privileges to the PDBADMIN user account, or to the PDB_DBA role.

You can use native Oracle Database facilities to create extra PDBs, and to manage all of your PDBs. The dbaascli utility also provides a range of convenient PDB management functions, including the ability to create nonCDB databases (release 19.x and earlier).

> **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.

Using the Console to Create a Database

To create an Oracle Database with the console, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
   VM Clusters is selected by default.
2. Choose your **Compartment**.
   A list of VM Clusters is displayed for the chosen Compartment.
3. Click the name of a VM cluster where you want to create the database.
   In the VM Cluster Details page, under **Resources**, **Databases** is selected by default.
4. Click **Create Database**.
   (or)
   a. Click **Database Homes**.
   b. Click the name of the Database Home where you want to create the database.
   c. Click **Create Database**.
5. Provide the requested information in the Create Database page:

> **Note**
>
> You cannot modify the db_name, db_unique_name, and SID prefix after creating the database.

- **Provide the database name**: Specify a user-friendly name that you can use to identify the database. The database name must contain only the permitted characters. Review the following guidelines when selecting a database name.
  - maximum of 8 characters
  - contain only alphanumeric characters
  - begin with an alphabetic character
  - cannot be part of first 8 characters of a db_unique_name on the VM cluster

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - unique within a VM cluster
  - **DO NOT** use `grid` because `grid` is a reserved name
  - **DO NOT** use `ASM` because `ASM` is a reserved name
- **Provide a unique name for the database**: Optionally, specify a unique name for the database. This attribute defines the value of the `db_unique_name` database parameter. The value is case insensitive.

  The `db_unique_name` must contain only the permitted characters. Review the following guidelines when selecting a database name.

  - maximum of 30 characters
  - can contain alphanumeric and underscore (_) characters
  - begin with an alphabetic character
  - unique across the fleet/tenancy

  If a unique name is not provided, then the `db_unique_name` defaults to the following format `<db_name>_<3 char unique string>_<region-name>`.

  If you plan to configure the database for backup to a Recovery Appliance backup destination, then the unique database name must match the name that is configured in the Recovery Appliance.
- **Select a database version**: From the list, choose the Oracle Database software release that you want to deploy.
- **Provide the name of the first PDB**: (Optional) Specify the name for the first PDB. A PDB is created with the database.
  To avoid potential service name collisions when using Oracle Net Services to connect to the PDB, ensure that the PDB name is unique across the entire VM cluster. If you do not provide the name of the first PDB, then a system-generated name is used.
- **Database Home**: Select an existing Database Home or create one as applicable. Note that this field is not available when you create a Database from the Database Home details page.
  - **Select an existing Database Home**: If one or more Database Homes already exist for the database version you have selected, then this option is selected by default. And, you will be presented with a list of Database Homes. Select a Database Home from the list.
  - **Create a new Database Home**: If no Database Homes exist for the database version you have selected, then this option is selected by default.
    a. Enter **Database Home display name**.
    b. Click **Change Database Image** to select your software version.
    Select a **Database Software Image** window is displayed.
    c. Select an **Image Type**, **Oracle Provided Database Software Images**, or **Custom Database Software Images**.
    If you choose **Oracle Provided Database Software Images**, then you can use the **Display all available version** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

- **Provide the administration password**: Provide and confirm the Oracle Database administration password. This password is used for administration accounts and functions in the database, including:
  - The password for the Oracle Database `SYS` and `SYSTEM` users.
  - The Transparent Data Encryption (TDE) Keystore password.

  For Oracle Database 12c Release 1 or later releases, the password for the PDB administration user in the first PDB (`PDBADMIN`) must be nine to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. In addition, the password must not contain the name of the tenancy or any reserved words, such as `Oracle` or `Table`, regardless of casing.
  - **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.
- **Backup Destination Type:** Select a backup destination for the database. From the list, choose an option:
  - **Object Storage**: Select to store backups in an Oracle-managed object storage container on Oracle Cloud Infrastructure.
    To use this option, your Oracle Exadata Cloud@Customer system must have egress connectivity to Oracle Cloud Infrastructure Object Storage.
  - **NFS**: Select to store backups in one of your previously defined backup destinations that use Network File System (NFS) storage. For more information, refer to the information about backup destinations in this publication.
    If you select this option, then you must also choose from the list of NFS **Backup Destinations**.
  - **Recovery Appliance**: Select to store backups in one of your previously defined backup destinations that use Oracle Zero Data Loss Recovery Appliance. Refer to the information about backup destination options in this document.
    If you select Oracle Zero Data Loss Recovery Appliance as your backup option, then you must also:
    * Choose from the list of appliance **Backup Destinations**.
    * Choose from the **VPC User** list, which contains the list of virtual private catalog (VPC) user names that are defined in the Oracle Zero Data Loss Recovery Appliance backup destination.
    * Provide the **Password** for the VPC user.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you select a backup destination, then you cannot change a backup location after the database is created.

- **Enable automatic backups:** Specify the settings for backing up the database to Object Storage, Network File System, or Zero Data Loss Recovery Appliance (ZDLRA).

  > **Note**
  >
  > You can change this setting after database creation.

  * **Object Storage**
    * **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.
      * **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.
    * **Backup Retention Period:** You can choose one of the following preset retention periods.
      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.
    * **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 as per backup retention period:** Select this option if you want to retain database backups for the entire period defined in the Backup retention period 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.
  * **Network File System (NFS)**

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Compartment:** Choose the compartment where the NFS backup destination exists.
- **Backup Destinations:** Choose from the list of NFS backup destinations.
- **NFS (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.
  - **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.
- **NFS (L1):** Incremental backup scheduling time (UTC): Specify the time window when the incremental backups start.
- **Backup Retention Period:** You can choose one of the following preset retention periods.
  NFS 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.
- **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 as per backup retention period:** Select this option if you want to retain database backups for the entire period defined in the Backup retention period 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.

- **Recovery Appliance**
  - **Compartment:** Choose the compartment where the Recovery Appliance backup destination exists.
  - **Backup Destinations:** Choose from the list of Recovery Appliance backup destinations.
  - **VPC User:** Provide a VPC user name and password for connecting to the Zero Data Loss Recovery Appliance.
  - **Recovery Appliance (L0):**
    - **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

Copyright © 2008, 2026, Oracle and/or its affiliates.

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.

* **Recovery Apliance (L1):** Incremental backup scheduling time (UTC): Specify the time window when the incremental backups start.
* **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 Backup retention period 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.

- **Key Management**: Choose an encryption option, **Oracle Wallet**, **Oracle Key Vault**, or **External HSM**. 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.
  - **Oracle Key Vault**: Select the compartment where you have created the Key Store, and then select the Key Store. As part of the CDB creation, a new wallet is created for the CDB in Oracle Key Vault (OKV). Also, a TDE Master Key is generated for the CDB and added to the wallet in OKV.

    Note

    * Container Databases (CDB) and pluggable databases (PDB) only support 256-bit Hardware Security Module (HSM) Vault keys.
    * Validate OKV Key encryption post restart: OKV TDE Maser Key is validated every time you start or restart your CBD.
    * Start or restart fails if the key is not validated. Work requests and life cycle states indicate the reason for failure.
    * View OKV keys post database restore: When you restore a CDB, the master key associated with that backup is restored as well.
    * Enable CDB backups to capture wallet name: CDB backups information about the wallet associated with the backup.
    * OKV Wallet or TDE Master Key on CDB deletion: If you delete a CDB, then the wallet and TDE Master Key remain in OKV and will not be deleted.

  - **External Keystore:**
    * **TDE wallet password:** Enter the TDE wallet password.

Copyright © 2008, 2026, Oracle and/or its affiliates.

    - **External Keystore credential:** Enter user credentials in the format your external keystore provider supports.

- (Optional) Select **Show Advanced Options**. From this window, you can select the following options:
  - **Provide the Oracle SID prefix**:

    > **Note**
    >
    > Entering a SID prefix is only available for 12.1 databases and above.

    Optionally, specify the Oracle SID prefix for the database. The instance number is automatically appended to the SID prefix to become the `instance_name` database parameter. If not provided, then the SID prefix defaults to the `db_name`.

    Review the following guidelines when selecting a database name:

    - maximum of 12 characters
    - contain only alphanumeric characters
    - begin with an alphabetic character
    - unique in the VM cluster
  - Real-time data protection: Real-time protection is the continuous transfer of redo changes from a protected database to ZDLRA. This reduces data loss and provides a recovery point objective (RPO) near 0.
    Once Real-time data protection is enabled, you can see the status in the Backup section of the Database details page.
  - **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`.
  - **Tags**: (Optional) You can choose to apply tags. If you have permissions to create a resource, 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, refer to information about 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.

6. Click **Create**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> You can now:
>
> - Create or delete a CDB while a Data Guard setup is running on another database within the same Oracle home, and vice versa.
> - Create or delete a CDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.
> - Create or delete a CDB while concurrently creating or deleting a PDB within the same Oracle home, and vice versa.
> - Create or delete a CDB concurrently on different databases within the same Oracle home.
> - Create or delete a CDB while simultaneously updating VM Cluster tags.

Related Topics

- Resource Tags
- Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer
  Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Cloud@Customer.
- Manage Keys Using an External Keystore
  Review the use cases and implementation details of an external keystore.

Connecting to an Oracle Exadata Database Service on Cloud@Customer System

After deploying your Oracle Exadata Database Service on Cloud@Customer system with a VM cluster, Oracle Database Home, and Oracle Database, learn how to connect to your VM Cluster virtual machine using SSH and to connect to an Oracle Exadata Database Service on Cloud@Customer database using Oracle Net Services (SQL*Net).

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- Connecting to a Database with Oracle Net Services
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system using Oracle Net Services.

Connecting to a Virtual Machine with SSH

You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer 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".

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Prerequisites for Connecting to an Oracle Exadata Database Service on Cloud@Customer System
  To access a virtual machine in an Oracle Exadata Database Service on Cloud@Customer system using SSH, be prepared to provide the host name or IP address of the virtual machine.
- Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY
  Learn to access a virtual machine from a Microsoft Windows system using PuTTY.
- Accessing a Database After You Connect to the Virtual Machine
  After you connect to a virtual machine, you can use the following series of commands to identify a database and connect to it.
- Connecting from a Unix-Style System
  To access a virtual machine on an Oracle Oracle Exadata Database Service on Cloud@Customer system from a Unix-style system using SSH, use this procedure.

Related Topics

- http://www.putty.org/

Prerequisites for Connecting to an Oracle Exadata Database Service on Cloud@Customer System

To access a virtual machine in an Oracle Exadata Database Service on Cloud@Customer system using SSH, be prepared to provide the host name or IP address of the virtual machine.

- An SSH private key file that corresponds to a public key that is registered in the system. When you create a VM cluster on your Oracle Exadata Database Service on Cloud@Customer system, you must specify the public key portion of one or more SSH key pairs. You can also register extra keys separately after you create the VM cluster.

  > **Note**
  >
  > The public keys are stored in the `authorized_keys` file at `~/.ssh/authorized_keys`. Separate `authorized_keys` files are located under the home directories of the operating system users. By default, only the `opc` user account has an authorized_keys entry and is able to log in remotely. Do not remove or alter the automatically generated entry in `authorized_keys` for the `opc` user.

- The host name or IP address for the virtual machine that you want to access.
  See, *Using the Console to Check the Status of a VM Cluster Virtual Machine*.

Related Topics

- Using the Console to Check the Status of a VM Cluster Virtual Machine
  Review the health status of a VM cluster virtual machine.

Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY

Learn to access a virtual machine from a Microsoft Windows system using PuTTY.

Before you begin

Before you use the PuTTY program to connect to a virtual machine, you need the following:

- The IP address of the virtual machine

Copyright © 2008, 2026, Oracle and/or its affiliates.

- The SSH private key file that matches the public key associated with the deployment. This private key file must be in the PuTTY `.ppk` format. If the private key file was originally created on the Linux platform, you can use the PuTTYgen program to convert it to the `.ppk` format.

To connect to a virtual machine using the PuTTY program on Windows:

1. Download and install PuTTY.

   To download PuTTY, go to http://www.putty.org/ and click the **You can download PuTTY here** link.

2. Run the PuTTY program (`putty.exe`).

   The PuTTY Configuration window is displayed, showing the **Session** panel.

3. In the **Host Name (or IP address)** field, enter the host name or IP address of the virtual machine that you want to access.

4. Confirm that the **Connection type** option is set to **SSH**.

5. In the **Category** tree, expand **Connection** if necessary and then click **Data**.

   The **Data** panel is displayed.

6. In the **Auto-login username** field, enter the operating system user you want to connect as:

   - Connect as the user `opc` to perform operations that require `root` or `oracle` access to the virtual machine, such as backing up or patching; this user can use the `sudo` command to gain `root` or `oracle` access to the VM.

7. Confirm that the **When username is not specified** option is set to **Prompt**.

8. In the **Category** tree, expand **SSH** and then click **Auth**.

   The **Auth** panel is displayed.

9. Click the **Browse** button next to the **Private key file for authentication** field. Then, in the **Select private key file** window, navigate to and open the private key file that matches the public key that is associated with the deployment.

10. In the **Category** tree, click **Session**.

    The **Session** panel is displayed.

11. In the **Saved Sessions** field, enter a name for the connection configuration. Then, click **Save**.

12. Click **Open** to open the connection.

    The PuTTY Configuration window closes and the PuTTY terminal window displays.

    If this is the first time you are connecting to the VM, the PuTTY Security Alert window is displayed, prompting you to confirm the public key. Click **Yes** to continue connecting.

Accessing a Database After You Connect to the Virtual Machine

After you connect to a virtual machine, you can use the following series of commands to identify a database and connect to it.

1. SSH in as the `opc` user.

2. `sudo su oracle`

Copyright © 2008, 2026, Oracle and/or its affiliates.

3. Use the `srvctl` utility located under the Oracle Grid Infrastructure home directory to list the databases on the system. For example:

```
/u01/app/12.2.0.1/grid/bin/srvctl config database -v
nc122   /u02/app/oracle/product/12.2.0/dbhome_6 12.2.0.1.0
s12c    /u02/app/oracle/product/12.2.0/dbhome_2 12.2.0.1.0
```

4. Identify the database instances for the database that you want to access. For example:

```
/u01/app/12.2.0.1/grid/bin/srvctl status database -d s12c
Instance s12c1 is running on node node01
Instance s12c2 is running on node node02
```

5. Configure the environment settings for the database that you want to access. For example:

```
. oraenv
ORACLE_SID = [oracle] ? s12c
The Oracle base has been set to /u02/app/oracle


export ORACLE_SID=s12c1
```

6. You can use the `svrctl` command to display more detailed information about the database. For example:

```
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
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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
```

Connecting from a Unix-Style System

To access a virtual machine on an Oracle Oracle Exadata Database Service on Cloud@Customer 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 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 Database with Oracle Net Services

You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system using Oracle Net Services.

- Using Oracle Net Services to Connect to a Database
  Oracle Database Oracle Exadata Database Service on Cloud@Customer 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 Oracle Exadata Database Service on Cloud@Customer 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 Cloud@Customer with a connect descriptor that bypasses the SCAN listeners, use this procedure to route your connection directly to a node listener.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using Oracle Net Services to Connect to a Database

Oracle Database Oracle Exadata Database Service on Cloud@Customer supports remote database access by using Oracle Net Services.

Because Oracle Exadata Database Service on Cloud@Customer 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 Cloud@Customer databases by using Oracle Net Services.

Prerequisites for Connecting to a Database with Oracle Net Services

Review the prerequisites to connect to an Oracle Database instance on Oracle Oracle Exadata Database Service on Cloud@Customer using Oracle Net Services.

To connect to an Oracle Database on Oracle Exadata Database Service on Cloud@Customer with Oracle Net Services, you need the following:

- The IP addresses for your SCAN VIPs, or the hostname or IP address for a virtual machine that hosts the database that you want to access.
- The database identifier: Either the database system identifier (SID), or a service name.

Connecting to a Database Using SCAN

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

- 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 Cloud@Customer System using multiple SCAN listeners.

Copyright © 2008, 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 Cloud@Customer 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 Cloud@Customer 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 Cloud@Customer 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 © 2008, 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 Cloud@Customer 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 © 2008, 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
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

5

How-to Guides

A collection of tasks and procedures for managing Oracle Exadata Database Service on Cloud@Customer.

- Manage Database Security with Oracle Data Safe
- Connect to the Oracle Exadata Database Service on Cloud@Customer Service
  Learn how to connect to an Oracle Exadata Database Service on Cloud@Customer system using SSH, and how to connect to an Exadata Cloud@Customer database using Oracle Net Services (SQL*Net).
- Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure
  Use the provided tools to manage the Infrastructure.
- Configure Oracle-Managed Infrastructure Maintenance
  Oracle performs the updates to all of the Oracle-managed infrastructure components on Oracle Exadata Database Service on Cloud@Customer.
- Manage VM Cluster Networks
  Learn how to use the console to create, edit, download a configuration file, validate, and terminate your infrastructure network.
- Manage VM Clusters
  Learn how to manage your VM clusters on Oracle Exadata Database Service on Cloud@Customer.
- Manage Software Images
- Create Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to create Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer.
- Manage Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems
  Learn to manage Oracle Database homes on Oracle Exadata Database Service on Cloud@Customer.
- Manage Databases on Oracle Exadata Database Service on Cloud@Customer
- Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer
  Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Cloud@Customer.
- Patch and Update an Oracle Exadata Database Service on Cloud@Customer System
  Learn to update and patch the Oracle Exadata Database Service on Cloud@Customer System
- Interim Software Updates
  For authorized environments, learn how to download interim software updates.
- Use Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer
  Learn to configure and manage Data Guard Groups in your VM cluster.
- Manage Keys Using an External Keystore
  Review the use cases and implementation details of an external keystore.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Managing Encryption Keys on External Devices
  Learn how to store and manage database encryption keys.
- Migrate to Oracle Exadata Database Service on Cloud@Customer
  For general guidance on methods and tools to migrate databases to Oracle Cloud Infrastructure database services, including Oracle Exadata Database Service on Cloud@Customer see "Migrating Databases to the Cloud".
- Overview of Exadata Cloud@Customer Gen1 to Out-of-Place Cloud Upgrade to Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure
  Gen1 is the first generation of Exadata Database Service on Cloud@Customer, which is deployed in conjunction with Gen1 Oracle Cloud At Customer (OCC) as Control Plane deployed in the customer data center. Oracle Exadata Database Service on Cloud@Customer Gen2 is managed from Oracle Cloud Infrastructure (OCI) Control Plane, which runs in OCI public cloud.

Manage Database Security with Oracle Data Safe

- About Oracle Data Safe
- Get Started
- Using Oracle Data Safe

About Oracle Data Safe

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

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

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

Data Safe provides the following features:

- **Security Assessment**: Configuration errors and configuration drift are significant contributors to data breaches. Use security assessment to evaluate your database's configuration and compare it to Oracle and industry best practices. Security assessment reports on areas of risk and notifies you when configurations change.
- **User Assessment**: Many breaches start with a compromised user account. User Assessment helps you spot the riskiest database accounts - those accounts which, if compromised, could cause the most damage - and take proactive steps to secure them. User Assessment Baselines make it easy to know when new accounts are added, or an account's privileges are modified. Use OCI events to receive proactive notifications when a database deviates from its baseline.
- **Activity Auditing**: Understanding and reporting on user activity, data access, and changes to database structures supports regulatory compliance requirements and can aid in post-

Copyright © 2008, 2026, Oracle and/or its affiliates.

incident investigations. Activity auditing collects audit records from databases and helps you manage audit policies. Audit insights make it easy to identify inefficient audit policies, while alerts based on audit data proactively notify you of risky activity.

- **Sensitive Data Discovery**: Knowing what sensitive data is managed in your applications is critical for security and privacy. Data discovery scans your database for over 150 different types of sensitive data, helping you understand what types and how much sensitive data you are storing. Use these reports to formulate audit policies, develop data masking templates, and create effective access control policies.
- **Data Masking**: Minimizing the amount of sensitive data your organization maintains helps you meet compliance requirements and satisfy data privacy regulations. Data masking helps you remove risk from your non-production databases by replacing sensitive information with masked data. With reusable masking templates, over 50 included masking formats, and the ability to easily create custom formats for your organization's unique requirements, data masking can streamline your application development and testing operations.
- **SQL Firewall Management**: Protect against risks such as SQL injection attacks or compromised accounts. Oracle SQL Firewall is a new security capability built into the Oracle AI Database 26ai kernel and offers best-in-class protection against these risks. The SQL Firewall feature in Oracle Data Safe lets you centrally manage and monitor the SQL Firewall policies for your target databases. Data Safe lets you collect authorized SQL activities of a database user, generate and enable the policy with allowlists of approved SQL statements and database connection paths and provides a comprehensive view of any SQL Firewall violations across the fleet of your target databases.

**Includes 1 million audit records per database per month if using the audit collection for Activity Auditing*

Get Started

To get started you just need to register your database with Oracle Data Safe:

- Pre-requisite: Obtain the necessary Identity and Access Management (IAM) permissions to register your target database in Data Safe: Permissions to register an Oracle Cloud@Customer Database
- Select an option for connecting your database to Data Safe
  - Connect over VPN or FastConnect using a **Data Safe private endpoint**
    If you have FastConnect or VPN Connect set up between your Cloud@Customer environment and a virtual cloud network (VCN) in the Oracle Cloud Infrastructure (OCI), you can register your database with Oracle Data Safe by using an Oracle Data Safe private endpoint. You can create the private endpoint during the registration or before. You can find more details on how to create the private endpoint under Create an Oracle Data Safe Private Endpoint.
  - Connect using a **Data Safe on-premises connectors**
    If you don't have FastConnect or VPN set up between your Cloud@Customer environment and OCI or you don't want to use it for Data Safe, you can register your database with Data Safe by using an Oracle Data Safe on-premises connector. You can create and install the on-premises connector during the registration or before. You can find more details on how to create the connector under Create an Oracle Data Safe On-Premises Connector.
- Register your Cloud@Customer database in Data Safe

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using Oracle Data Safe

Once your database is registered in Data Safe, you can leverage all features.

Security Assessment

Security Assessments are automatically scheduled once a week in Data Safe and provide an overall picture of your database security posture. It analyzes your database configurations, users and user entitlements, as well as security policies to uncover security risks and improve the security posture of Oracle Databases within your organization. A security assessment provides findings with recommendations for remediation activities that follow best practices to reduce or mitigate risk.

Start by reviewing the security assessment report for your database: View the latest assessment for a target database

You can find more details on Security Assessment under Security Assessment Overview.

User Assessment

User Assessments are automatically scheduled once a week in Data Safe and help you to identify highly privileged user accounts that could pose a threat if misused or compromised. User Assessment reviews information about your users in the data dictionaries on your target databases and then calculates a potential risk for each user, based on system privileges and role grants.

Start by reviewing the user assessment report for your database: View the latest user assessment for a target database

You can find more details on User Assessment under User Assessment Overview.

Data Discovery

Data Discovery searches for sensitive columns in your database. It comes with over 150 pre-defined sensitive types and you can also create your own sensitive types. You tell Data Discovery if you want to scan your entire database or just certain schemas and what type of sensitive information to look for, and it finds the sensitive columns that meet your criteria and stores them in a sensitive data model (SDM).

Start by discovering sensitive data in your database: Create Sensitive Data Models

You can find more details on Data Discovery under Data Discovery Overview.

Data Masking

Data masking, also known as static data masking helps you to replace sensitive or confidential information in your non-production databases with realistic and fully functional data with similar characteristics as the original data. Data Safe comes with pre-defined masking formats for each of the pre-defined sensitive types that can also be leveraged for your own sensitive types.

Once you know where sensitive data is stored in your database (for instance after running Data Discovery in Data Safe), you can start by creating a masking policy: Create Masking Policies

After you created a masking policy and copied your production database, you can mask your non-production copy: Mask Sensitive Data on a Target Database

You can find more details on Data Masking under Data Masking Overview.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Activity Auditing

Activity Auditing in Oracle Data Safe helps to ensure accountability and improve regulatory compliance. With Activity Auditing, you can collect and retain audit records per industry and regulatory compliance requirements and monitor user activities on Oracle databases with pre-defined reports and alerts. For example, you can audit access to sensitive data, security-relevant events, administrator and user activities, activities recommended by compliance regulations like the Center for Internet Security (CIS), and activities defined by your own organization.

If you are using the audit collection in Data Safe, up to 1 million audit records per target database per month are included for your Cloud@Customer database.

To use activity auditing, start the audit trail for your target database in Data Safe: Start an Audit Trail

Once the audit trail is started, you can monitor and analyze your audit data with pre-defined audit reports: View a Predefined or Custom Audit Report

You can find more details on Activity Auditing under Activity Auditing Overview.

SQL Firewall*

SQL Firewall in Oracle Data Safe lets you centrally manage the SQL Firewalls and provides a comprehensive view of SQL Firewall violations across the fleet of your target databases. Data Safe lets you collect authorized SQL activities of a database user you wish to protect, monitor the progress of the collection, generate and enable the policy with allowlists of approved SQL statements and database connection paths.

Start by enabling the SQL Firewall in your 26ai target database: Enable SQL Firewall On Your Target Database.

Next, you need to generate and enable a SQL Firewall policy with allowlists for the database user you wish to protect: Generate and Enforce SQL Firewall Policies.

Once you start enforcing the SQL Firewall policy, you can monitor and analyze the violations in the pre-defined violation reports: View and Manage Violations Reports.

You can find more details on SQL Firewall under SQL Firewall Overview.

*SQL Firewall is only available for Oracle AI Database 26ai.

Connect to the Oracle Exadata Database Service on Cloud@Customer Service

Learn how to connect to an Oracle Exadata Database Service on Cloud@Customer system using SSH, and how to connect to an Exadata Cloud@Customer database using Oracle Net Services (SQL*Net).

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- Connecting to a Database with Oracle Net Services
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system using Oracle Net Services.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Connecting to a Virtual Machine with SSH

You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer 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".

- Prerequisites for Connecting to an Oracle Exadata Database Service on Cloud@Customer System
  To access a virtual machine in an Oracle Exadata Database Service on Cloud@Customer system using SSH, be prepared to provide the host name or IP address of the virtual machine.
- Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY
  Learn to access a virtual machine from a Microsoft Windows system using PuTTY.
- Connecting from a Unix-Style System
  To access a virtual machine on an Oracle Oracle Exadata Database Service on Cloud@Customer system from a Unix-style system using SSH, use this procedure.
- 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/

Prerequisites for Connecting to an Oracle Exadata Database Service on Cloud@Customer System

To access a virtual machine in an Oracle Exadata Database Service on Cloud@Customer system using SSH, be prepared to provide the host name or IP address of the virtual machine.

- An SSH private key file that corresponds to a public key that is registered in the system. When you create a VM cluster on your Oracle Exadata Database Service on Cloud@Customer system, you must specify the public key portion of one or more SSH key pairs. You can also register extra keys separately after you create the VM cluster.

  > **Note**
  >
  > The public keys are stored in the `authorized_keys` file at `~/.ssh/authorized_keys`. Separate `authorized_keys` files are located under the home directories of the operating system users. By default, only the `opc` user account has an authorized_keys entry and is able to log in remotely. Do not remove or alter the automatically generated entry in `authorized_keys` for the `opc` user.

- The host name or IP address for the virtual machine that you want to access. See, *Using the Console to Check the Status of a VM Cluster Virtual Machine*.

Related Topics

- Using the Console to Check the Status of a VM Cluster Virtual Machine
  Review the health status of a VM cluster virtual machine.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY

Learn to access a virtual machine from a Microsoft Windows system using PuTTY.

Before you begin

Before you use the PuTTY program to connect to a virtual machine, you need the following:

- The IP address of the virtual machine
- The SSH private key file that matches the public key associated with the deployment. This private key file must be in the PuTTY `.ppk` format. If the private key file was originally created on the Linux platform, you can use the PuTTYgen program to convert it to the `.ppk` format.

To connect to a virtual machine using the PuTTY program on Windows:

1. Download and install PuTTY.

   To download PuTTY, go to http://www.putty.org/ and click the **You can download PuTTY here** link.

2. Run the PuTTY program (`putty.exe`).

   The PuTTY Configuration window is displayed, showing the **Session** panel.

3. In the **Host Name (or IP address)** field, enter the host name or IP address of the virtual machine that you want to access.

4. Confirm that the **Connection type** option is set to **SSH**.

5. In the **Category** tree, expand **Connection** if necessary and then click **Data**.

   The **Data** panel is displayed.

6. In the **Auto-login username** field, enter the operating system user you want to connect as:

   - Connect as the user `opc` to perform operations that require `root` or `oracle` access to the virtual machine, such as backing up or patching; this user can use the `sudo` command to gain `root` or `oracle` access to the VM.

7. Confirm that the **When username is not specified** option is set to **Prompt**.

8. In the **Category** tree, expand **SSH** and then click **Auth**.

   The **Auth** panel is displayed.

9. Click the **Browse** button next to the **Private key file for authentication** field. Then, in the **Select private key file** window, navigate to and open the private key file that matches the public key that is associated with the deployment.

10. In the **Category** tree, click **Session**.

    The **Session** panel is displayed.

11. In the **Saved Sessions** field, enter a name for the connection configuration. Then, click **Save**.

12. Click **Open** to open the connection.

    The PuTTY Configuration window closes and the PuTTY terminal window displays.

    If this is the first time you are connecting to the VM, the PuTTY Security Alert window is displayed, prompting you to confirm the public key. Click **Yes** to continue connecting.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Connecting from a Unix-Style System

To access a virtual machine on an Oracle Oracle Exadata Database Service on Cloud@Customer 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 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.

Accessing a Database After You Connect to the Virtual Machine

After you connect to a virtual machine, you can use the following series of commands to identify a database and connect to it.

1. SSH in as the `opc` user.
2. `sudo su oracle`
3. Use the `srvctl` utility located under the Oracle Grid Infrastructure home directory to list the databases on the system. For example:

   ```
   /u01/app/12.2.0.1/grid/bin/srvctl config database -v
   nc122   /u02/app/oracle/product/12.2.0/dbhome_6 12.2.0.1.0
   s12c    /u02/app/oracle/product/12.2.0/dbhome_2 12.2.0.1.0
   ```

4. Identify the database instances for the database that you want to access. For example:

   ```
   /u01/app/12.2.0.1/grid/bin/srvctl status database -d s12c
   Instance s12c1 is running on node node01
   Instance s12c2 is running on node node02
   ```

5. Configure the environment settings for the database that you want to access. For example:

   ```
   . oraenv
   ORACLE_SID = [oracle] ? s12c
   The Oracle base has been set to /u02/app/oracle


   export ORACLE_SID=s12c1
   ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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
```

Connecting to a Database with Oracle Net Services

You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system using Oracle Net Services.

- Using Oracle Net Services to Connect to a Database
  Oracle Database Oracle Exadata Database Service on Cloud@Customer 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 Oracle Exadata Database Service on Cloud@Customer using Oracle Net Services.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- 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 Cloud@Customer 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 Cloud@Customer supports remote database access by using Oracle Net Services.

Because Oracle Exadata Database Service on Cloud@Customer 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 Cloud@Customer databases by using Oracle Net Services.

Prerequisites for Connecting to a Database with Oracle Net Services

Review the prerequisites to connect to an Oracle Database instance on Oracle Oracle Exadata Database Service on Cloud@Customer using Oracle Net Services.

To connect to an Oracle Database on Oracle Exadata Database Service on Cloud@Customer with Oracle Net Services, you need the following:

- The IP addresses for your SCAN VIPs, or the hostname or IP address for a virtual machine that hosts the database that you want to access.
- The database identifier: Either the database system identifier (SID), or a service name.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 Cloud@Customer 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 Cloud@Customer 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 Cloud@Customer 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.

Copyright © 2008, 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 Cloud@Customer System using a custom SCAN name.

Using this approach, you define a custom single client access name (SCAN) name in your domain name server (DNS), which resolves to the three SCAN virtual IP addresses (VIPs).

- Use the following template to define a Net Services alias that references the custom SCAN name:

  ```
  alias-name = (DESCRIPTION=
    (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=scan-name)(PORT=1521)))
    (CONNECT_DATA=(sid-or-service-entry)))
  ```

  Where:

  alias-name is the name you use to identify the alias.

  scan-name is the custom SCAN name.

  sid-or-service-entry identifies the database SID or service name using one of the following formats:

  - `SID=`*sid-name*. For example: `SID=S12C1`.
  - `SERVICE_NAME=`*service-name*. For example: `SERVICE_NAME=PDB1.example.yourcloud.com`.

  Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

  ```
  scan-name:1521/sid-or-service-entry
  ```

  For example:

  ```
  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 Cloud@Customer 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 5-1 Defining a Net Service Alias That Directly References the Node

```
alias-name = (DESCRIPTION=
  (CONNECT_TIMEOUT=timeout)
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=node)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

timeout specifies a timeout period (in seconds), which enables you to terminate a connection attempt without having to wait for a TCP timeout. The (CONNECT_TIMEOUT=*timeout*) parameter is optional.

node is the hostname or IP address for the virtual machine that you want to use.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- SID=*sid-name*. For example, SID=S12C1.
- SERVICE_NAME=*service-name*. For example, SERVICE_NAME=PDB1.example.oraclecloudatcust.com.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

node:1521/*sid-or-service-entry*

For example:

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

Or

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

Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure

Use the provided tools to manage the Infrastructure.

- About Provisioning Oracle Exadata Database Service on Cloud@Customer Systems
  To provision an Oracle Exadata Database Service on Cloud@Customer system, you must work with Oracle to set up and configure the system.
- Defining Maintenance Contacts for Exadata Cloud@Customer
  Oracle Cloud Operations will use these contacts to communicate with customers about infrastructure-related events, including hardware failures and replacements, quarterly maintenance patches, and other operational updates or incidents
- About Changing the Home Region of Oracle Exadata Database Service on Cloud@Customer Infrastructure

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Overview of Elastic Storage Expansion
  With elastic storage expansion, you can dynamically increase your storage capacity to meet your growing workload requirements.
- Using the Console to Provision Oracle Exadata Database Service on Cloud@Customer
  Learn how to provision an Oracle Exadata Database Service on Cloud@Customer system.
- Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure
  Oracle Exadata Database Service on Cloud@Customer uses the same API as Oracle Cloud Infrastructure.

About Provisioning Oracle Exadata Database Service on Cloud@Customer Systems

To provision an Oracle Exadata Database Service on Cloud@Customer system, you must work with Oracle to set up and configure the system.

Provisioning an Oracle Exadata Database Service on Cloud@Customer system is a collaborative process. The process is performed in the following sequence:

1. You create the Oracle Exadata Database Service on Cloud@Customer infrastructure.
2. You generate a file containing the infrastructure configuration details, and provide it to Oracle.
3. The Oracle Exadata Database Service on Cloud@Customer system is physically installed in your data center.
4. Oracle uses the infrastructure configuration file to perform initial system configuration. At the end of this task, Oracle supplies you with an activation file.
5. You activate the Exadata Database Service on Cloud@Customer infrastructure by using the supplied activation file.

When the provisioning process is complete, the Oracle Exadata Database Service on Cloud@Customer system is ready for you to use. You can then create a virtual machine (VM) cluster, and later create some databases.

> ⚠ **Caution**
>
> Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, the APIs, or the command-line interface.

Defining Maintenance Contacts for Exadata Cloud@Customer

Oracle Cloud Operations will use these contacts to communicate with customers about infrastructure-related events, including hardware failures and replacements, quarterly maintenance patches, and other operational updates or incidents

- Add Your Primary Maintenance Contact
  The primary contact is used to coordinate onsite visits for that infrastructure. If Oracle detects an issue that requires an on-site engineer, we will create a Service Request (SR) on the customer's behalf using the primary contact.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Add Secondary Contacts for Cloud Notifications (CNs)
  The secondary contact will not be used to coordinate onsite visits, however secondary contacts will receive Cloud Notifications (CNs) related to your infrastructure. You can add up to nine secondary contacts, distribution list is allowed.
- Edit Maintenance Contacts
  Edit maintenance contacts to update details.
- Remove a Secondary Contact
  You can remove a secondary contact anytime you want to.

Add Your Primary Maintenance Contact

The primary contact is used to coordinate onsite visits for that infrastructure. If Oracle detects an issue that requires an on-site engineer, we will create a Service Request (SR) on the customer's behalf using the primary contact.

Primary Contact Requirements

The primary contact listed on this page must meet the following requirements:

- My Oracle Cloud Support (MOS) Access
  - Must have access to My Oracle Support. OC1 tenancy customers should use "Sign in with your commercial account" and their OCI credentials.
  - Non-OC1 customers should use "Sign in with oracle account."

> **Note**
>
> It is essential to provide accurate and up-to-date contact information. If you are unavailable, another user from MOS user group will be chosen. Please ensure that only individuals approved to coordinate visits are added. Refer to the MOS doc KB182930 on how to setup the user group and add contacts to it.

Prerequisite

Set up your MOS access and determine the ID of the user group you will use to manage the contacts Oracle should use for coordinating onsite visits.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and compartment that contains the Oracle Exadata infrastructure for which you want to view contact details.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure for which you want to view contact details.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Select **Action** then **Manage Contacts**.
6. Add a primary contact:
   a. ***<first name of primary contact> <last name of primary contact> <UGID>***: UGID is the ID of the user group you will use to manage the contacts Oracle should use for coordinating data center visits. You can either use an existing user group or create a

Copyright © 2008, 2026, Oracle and/or its affiliates.

new one. It is essential to keep the primary contact information current. If the primary contact is not available, we'll choose any other user from the user group. Refer to the MOS doc KB182930 on how to setup the user group and add contacts to it.

b. **Email:** Your email address associated with My Oracle Cloud Support Portal.

c. **Phone:** Must include a valid phone number to allow coordination for any necessary on-site visits related to hardware maintenance.

7. Click **Update**.

Add Secondary Contacts for Cloud Notifications (CNs)

The secondary contact will not be used to coordinate onsite visits, however secondary contacts will receive Cloud Notifications (CNs) related to your infrastructure. You can add up to nine secondary contacts, distribution list is allowed.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and compartment that contains the Oracle Exadata infrastructure for which you want to add secondary contacts.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure for which you want to add secondary contacts.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Manage Contacts**.
6. In the Manage Exadata Infrastructure Contacts window, click **Add Contact**.
7. In the Exadata Infrastructure Contacts window, enable the **Secondary maintenance contacts** toggle, click **Another Contact**, and then enter the contact details.
8. Click **Update**.

Edit Maintenance Contacts

Edit maintenance contacts to update details.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and compartment that contains the Oracle Exadata infrastructure for which you want to edit maintenance contact details.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure for which you want to edit maintenance contact details.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Manage Contacts**.
6. Edit the contact details, and then click **Update**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Remove a Secondary Contact

You can remove a secondary contact anytime you want to.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and compartment that contains the Oracle Exadata infrastructure for which you want to remove a secondary contact.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure for which you want to remove a secondary contact.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Manage Contacts**.
6. Click the cross icon next to the secondary contact, and then click **Update**.

About Changing the Home Region of Oracle Exadata Database Service on Cloud@Customer Infrastructure

Oracle Exadata Database Service on Cloud@Customer home region is the region where the OCI Control Plane resides for management and monitoring. You can change the association of provisioned infrastructure from the home region (from where it was provisioned) to a different target region. Most customers will not need to perform this change since it is expected to be a rare operation. You can choose to do that in case the home OCI region is not available for an extended period or another OCI region is preferred due to lower latency or some other reason. There is no service downtime while the change of home region is in progress. This change is a field engineer-assisted operation and a service request needs to be opened to perform this.

> **Note**
>
> Changing the home region of your ExaDB-C@C Infrastructure will not affect billing.

Prerequisites

- The customer must provide the new region for changing the home region of their ExaDB-C@C Infrastructure.
- Subscribe the tenancy to the new region.
- Keep the permissions and network connectivity configuration details handy.

Deployment process

1. Customer opens an SR for changing the home region for their ExaDB-C@C Infrastructure.
2. Oracle processes the SR and assigns a field engineer (FE).
3. FE schedules time to visit the customer data center collects configuration details, and runs precheck.
4. FE works with the customer to obtain credentials (API key/IAM access) for the new region.
5. Customer opens the required endpoints/ports for the new region.

Copyright © 2008, 2026, Oracle and/or its affiliates.

6. FE executes the scripts and validates connectivity.
7. FE validates OCI resources created in the OCI Control Plane for the new region.
8. Customer recreates the needed configuration.

After the home region change

- Failed operations in the previous home region must be resolved in the source region (applicable to the planned move to new region use cases).
- Failed or in-flux operations in the source region will have to be performed again from the new home region (applicable to disaster recovery and unplanned move to new region use cases).

Configuration changes after changing the home

After changing the home region, the following resources are created in the OCI control plane for the new region with new OCIDs

- Exadata Infrastructure
- VM Cluster Networks
- VM Clusters
- Databases (container and pluggable databases)
- Backups
- Backup Destinations for NFS and ZDLRA

By default, these will be created in the same compartment as before the region change. Users will have the option to choose a different compartment during the process of changing the region.

> **Note**
>
> Customers will need to set the following attributes in the new region again as these are not currently carried over after changing the home region.

- Infrastructure maintenance preferences
- OCI Events
- OCI Tags
- Customer Contacts
- Database Software Images
- Key Stores

Operator Access Control

- Create an Operator Control with a default name <*cps_name*>_opctl
  - **Resource Type** will be set to `exadatainfrastructure`
  - **Platform** will be set to `cloud at customer`
  - No actions will be preapproved
  - **Group allowed to approve access requests** will be set to `Use IAM Policy`
  - Message to the operator will be the same as the previous Operator Control

Copyright © 2008, 2026, Oracle and/or its affiliates.

- The compartment used will be the same compartment as the Exadata Infrastructure.
- A new operator assignment will be created using the Exadata OCID present in the Control Plane Server configuration files and associated with the Operator Control created in step 1.
  - **Assignment** will be set to `Always`.
  - **syslog forwarding** will be configured if it's configured by the customer.
  - Auto-approve access requests during the maintenance window will be set to false.
- For each rack, there will be one Operator Control and a corresponding assignment.
- None of the history such as access request history, audit log history, or more info history will be migrated from the old region.

Data Guard

After region change, verify the Data Guard Associations in the OCI control plane. There will be no impact to the Data Guard between the databases and manual switchover/failover might be needed.

Overview of Elastic Storage Expansion

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

Expand the storage capacity on-demand by scaling up the infrastructure with additional storage servers. You can allocate additional storage capacity available from the newly added storage servers to the already deployed VM Cluster without disrupting the current running workloads. Additional storage capacity from newly added storage servers is also available for provisioning new VM Clusters on the infrastructure.

With the elastic storage expansion capability, you can:

- Provision new Exadata Infrastructure with custom storage capacity.
- Start with a smaller storage footprint for the Exadata Infrastructure at install time.
- Expand the storage capacity on existing deployed Exadata Infrastructure on-demand in an automated, elastic fashion.
- Allocate additional storage capacity available from newly added storage servers to already deployed VM clusters and/or use them for provisioning new VM clusters on the infrastructure.

Table 5-1 Exadata Storage Server X7-2 key additional resources

Specification	Exadata Base System Storage Server X7-2	Exadata Storage Server X7-2
Additional Raw Flash Storage Capacity	6.4 TB	25.6 TB
Additional Raw Disk Storage Capacity	48 TB	120 TB
Additional Usable Storage Capacity	14 TB	35.3 TB

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-2 Exadata Storage Server X8-2 key additional resources

Specification	Exadata Base System Storage Server X8-2	Exadata Storage Server X8-2
Additional Raw Flash Storage Capacity	12.8 TB	25.6 TB
Additional Raw Disk Storage Capacity	84 TB	168 TB
Additional Usable Storage Capacity	24.6 TB	49.6 TB

Table 5-3 Exadata Storage Server X8M-2 key additional resources

Specification	Exadata Base System Storage Server X8M-2	Exadata Storage Server X8M-2
Additional Raw Flash Storage Capacity	12.8 TB	25.6 TB
Additional Raw Disk Storage Capacity	84 TB	168 TB
Additional Usable Storage Capacity	24.6 TB	49.6 TB
Additional Persistent Memory	-	1.5 TB

Table 5-4 Exadata Storage Server X9M-2 key additional resources

Specification	Exadata Base System Storage Server X9M-2	Exadata Storage Server X9M-2
Additional Raw Flash Storage Capacity	12.8 TB	25.6 TB
Additional Raw Disk Storage Capacity	84 TB	216 TB
Additional Usable Storage Capacity	24.6 TB	63.6 TB
Additional Persistent Memory	-	1.5 TB

Table 5-5 Exadata Storage Server X10M key additional resources

Specification	Exadata Storage Server X10M
Additional Raw Flash Storage Capacity	27.2 TB
Additional Raw Disk Storage Capacity	264 TB
Additional Usable Storage Capacity	80 TB
Additional Exadata RDMA Memory	1.25 TB

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-6 Exadata Storage Server X11M key additional resources

Specification	Exadata Storage Server X11M-BASE	Exadata Storage Server X11M-HC
Additional Raw Flash Storage Capacity	13.6 TB	27.2 TB
Additional Raw Disk Storage Capacity	132 TB	264 TB
Additional Usable Storage Capacity	35.6 TB	80 TB
Additional Exadata RDMA Memory	-	1.25 TB

Elastic scaling of Exadata Storage Servers is subject to the following conditions:

- The Exadata Cloud@Customer system configuration must be based on Oracle Exadata X7 hardware, Oracle Exadata X8 hardware, Oracle Exadata X8M hardware, or Oracle Exadata X9M hardware.
- Each Exadata Cloud@Customer system configuration can have an absolute maximum of:
  - X7, X8: 12 Exadata Storage Servers
  - X8M, X9M, X10M, X11M: 64 Storage Servers
- Exadata Infrastructure deployed with base configuration shape can only be expanded using base expansion SKU storage servers.
- Exadata Infrastructures deployed with X7 generation at install time can be scaled with X8 generation storage servers. X8 storage servers used to scale X7 infrastructure will only present the same total usable capacity as all other X7 storage servers that are already part of the infrastructure.
- Exadata Infrastructures deployed with X8M generation at install time can only be scaled with X8M or higher generation storage servers. X9M and X10M storage servers used to scale X8M infrastructure will only present the same total usable capacity as other X8M storage servers that are already part of the infrastructure.
- Exadata Infrastructures deployed with X9M generation at install time can only be scaled with X9M or higher generation storage servers. X10M storage servers used to scale X9M infrastructure will only present the same total usable capacity as other X9M storage servers that are already part of the infrastructure.
- Exadata Infrastructures deployed with X10M generation at install time can only be scaled with X10M or higher generation storage servers. X11M storage servers used to scale X10M infrastructure will only present the same total usable capacity as other X10M storage servers that are already part of the infrastructure.

Exadata Infrastructures deployed with additional storage servers will be configured as an Elastic shape with the total number of storage servers and usable capacity clearly called out for the given infrastructure.

Before you can scale the number of Exadata storage servers, review the site and network requirements, and the checklists to prepare and deploy Exadata Cloud@Customer. Ensure that you have worked with sales and followed the procurement process. The following figure provides you with an overview of the order and deployment process.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Figure 5-1 Overview of Order and Deployment Process

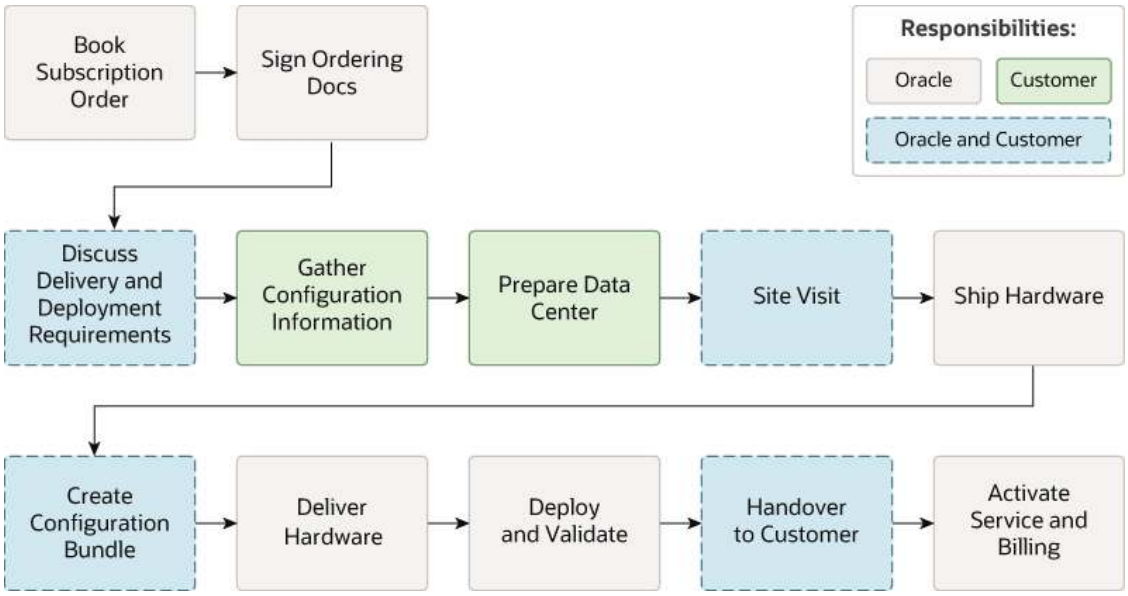

Multi-Rack Expansion

- **Rack Minimum Configuration**
  - **X11M/X10M**
    - Maximum 16 total servers in the initial rack
    - Maximum 18 total servers in the expansion rack
  - **X8M/X9M**
    - Base system expandable to a single rack
    - Maximum 8 database and 12 storage servers in the first rack
    - Maximum 11 database and 18 storage servers in the expansion racks
- **Minimum Expansion Rack Configuration**
  - Start with one storage server and add at least 3 other servers (storage and compute).
- **Maximum Overall Configuration**
  - **Maximum total database servers:** 32
  - **Maximum total storage servers:** 64
  - **Maximum total racks (including initial):** 6

Related Topics

- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.
- Using the Console to Download Scale Infrastructure Compute and Storage Configuration File
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Activate New Database and Storage Servers
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.
- Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption
  To make storage capacity from the new servers for VM clusters consumption, complete this procedure.
- Using the Console to View Details of Oracle Exadata Database Service on Cloud@Customer Infrastructure with Scaled Compute and Storage Capacity
  To view the storage capacity from the new storage server, or scaled compute capacity, use this procedure.

Using the Console to Provision Oracle Exadata Database Service on Cloud@Customer

Learn how to provision an Oracle Exadata Database Service on Cloud@Customer system.

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Edit Network Bonding
  To edit the client and backup network bonding mode, follow these steps.
- Using the Console to View Exadata Infrastructure Network Configuration Details
  To view network configuration details, follow these steps. Save this information for later use to troubleshoot if you face network issues.
- Using the Console to Edit Oracle Exadata Database Service on Cloud@Customer Infrastructure Networking Configuration
  To edit your Oracle Exadata Database Service on Cloud@Customer infrastructure networking configuration, be prepared to provide values for the infrastructure configuration.
- Using the Console to Download a File Containing Configuration Data
  To download an Oracle Exadata Database Service on Cloud@Customer configuration file, complete this procedure.
- Using the Console to Activate Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To activate Oracle Exadata Database Service on Cloud@Customer infrastructure, ensure that you meet the prerequisites, and complete this procedure.
- Using the Console to Check the Status of Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To find the status of your Oracle Exadata Database Service on Cloud@Customer infrastructure, use this procedure to check the Infrastructure Details page.
- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.
- Using the Console to Download Scale Infrastructure Compute and Storage Configuration File
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.
- Using the Console to Activate New Database and Storage Servers
  To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption
  To make storage capacity from the new servers for VM clusters consumption, complete this procedure.
- Using the Console to View Details of Oracle Exadata Database Service on Cloud@Customer Infrastructure with Scaled Compute and Storage Capacity
  To view the storage capacity from the new storage server, or scaled compute capacity, use this procedure.
- Using the Console to Move Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To relocate Oracle Exadata Database Service on Cloud@Customer infrastructure to another compartment, use this procedure.
- Using the Console to Terminate Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To terminate Oracle Exadata Database Service on Cloud@Customer infrastructure, complete the prerequisites, and then complete this procedure.
- Using the Console to Manage Tags for Your Oracle Exadata Database Service on Cloud@Customer Resources

Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure

To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Region**, select the region that you want to associate with the Oracle Exadata infrastructure.
   The region that is associated with your Oracle Exadata infrastructure cannot be changed after the Oracle Exadata infrastructure is created. Therefore, ensure that you select the most appropriate region for your infrastructure. Consider the following factors:
   - Consider any business policies or regulations that preclude the use of a particular region. For example, you can be required to maintain all operations within national boundaries.
   - Consider the physical proximity of the region to your data center. Needless extra physical separation adds unnecessary latency to network communications between Oracle Cloud Infrastructure and your corporate data center.
3. Click **Exadata Infrastructure**.
4. Click **Create Exadata Infrastructure**.
5. In the Create Exadata Infrastructure page, provide the requested information:
   - **Provide the display name:** The display name is a user-friendly name that you can use to identify the Exadata infrastructure. The name doesn't need to be unique, because an Oracle Cloud Identifier (OCID) uniquely identifies the Oracle Exadata infrastructure.
   - **Choose a compartment:** From the list of available compartments, choose the compartment that you want to contain the Oracle Exadata infrastructure.
     For more information, see *Understanding Compartments*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Oracle Cloud Infrastructure region:** The region that is associated with your Oracle Exadata infrastructure cannot be changed after the Oracle Exadata infrastructure is created. Therefore, check the displayed region to ensure that you are using the most appropriate region for your infrastructure.
  See step 2 (earlier in this procedure) for further considerations. To switch regions now, use the **Region** menu at the top of the console.
- **Select the Exadata system model:** From the list, choose an Oracle Exadata hardware.
  The Oracle Exadata system model and system shape combine to define the amount of CPU, memory, and storage resources that are available in the Exadata infrastructure. For more information, see *System Configuration*.
- **Compute and storage configuration:**
  - **Database server type (displayed if Exadata Cloud@Customer X11M is selected):**
    - **Choices:** X11M-BASE, X11M, X11M-L, and X11M-XL
  - **Storage server type (displayed if Exadata Cloud@Customer X11M is selected):**
    - **Choices:** X11M-BASE and X11M-HC
  - **Database servers:**
    - **X11M/X10M:** You can expand the database and storage servers in the initial rack up to 16 total servers.
    - **X8M/X9M:** You can add a minimum of 2 and extend up to a maximum of 8 database servers.
  - **Storage servers:**
    - **X11M/X10M:** You can expand the database and storage servers in the initial rack up to 16 total servers.
    - **X8M/X9M:** You can add a minimum of 3 and extend up to a maximum of 12 storage servers. For each storage server you add, the storage capacity that will be added is displayed on the right.
  - **Multi-rack deployment:** If you scale beyond the maximum limits, Database servers: 8 and Storage servers: 12, then the multi-rack deployment toggle is automatically enabled. You can also enable that option if the database servers count and storage server count are within the minimum and maximum bounds; however, you want the infrastructure to span multiple racks.
    A new field, **Upload multi-rack configuration file**, appears when you enable this toggle button. You can drag and drop the JSON file or select one from the local filesystem. For multi-rack deployments, the Control Plane cannot generate the configuration bundle correctly without the JSON file. The JSON file will be used to generate host names and IP addresses for the additional components. Contact your Oracle Field Engineer for more information.

    If you fail to identify an infrastructure as multi-rack and subsequently determine it is multi-rack, then you will have to delete that infrastructure and recreate it.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Resource allocation will display ECPUs for X11M database servers and OCPUs for X10M and older generation Exadata infrastructure.
> **ECPU:** An ECPU is an abstracted measure of compute resources. ECPUs are based on the number of cores elastically allocated from a pool of compute and storage servers.

- **Site network details**
  - **Time Zone:** The default time zone for the Exadata Infrastructure is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

    > **Note**
    >
    > If you want to set a time zone other than UTC or the browser-detected time zone, then select the **Select another time zone** option, select a **Region** or **country**, and then select the corresponding **Time zone**.
    >
    > If you do not see the region or country you want, then select **Miscellaneous**, and then select an appropriate **Time zone**.

- **Configure the control plane server network**
  Each Oracle Exadata Database Service on Cloud@Customer system contains two control plane servers, which enable connectivity to Oracle Cloud Infrastructure. The control plane servers are connected to the control plane network, which is a subnet on your corporate network. The following settings define the network parameters:
  - **Control Plane Server 1 IP Address:** Provide the IP address for the first control plane server. This IP address is for the network interface that connects the first control plane server to your corporate network using the control plane network.
  - **Control Plane Server 2 IP Address:** Provide the IP address for the second control plane server. This IP address is for the network interface that connects the second control plane server to your corporate network using the control plane network.
  - **Netmask:** Specify the IP netmask for the control plane network.
  - **Gateway:** Specify the IP address of the control plane network gateway.
  - **HTTP Proxy:** (Optional) You can choose to use this field to specify your corporate HTTP proxy. The expected format is as follows, where *`server`* is the server name, *`domain`* is the domain name, and *`port`* is the assigned port:

    `http://server.domain:port`

    For example:

    `http://proxy.example.com:80`

    For enhanced security, when possible, Oracle recommends that you use an HTTP proxy.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Enable Control Plane Server Offline Report:** Enabling the Control Plane Server (CPS) offline report helps in diagnosing connectivity issues between the CPS and OCI endpoints, should they arise.
  To view the report, do the following:

  a. Find the CPS IP addresses.
     For more information, see Using the Console to View Exadata Infrastructure Network Configuration Details.

  b. From your local network, access the report over HTTP.
     To view the report in HTML format, use `http://<CPSPublicIP>:18080/report`

     To view the report in JSON format, use `http://<CPSPublicIP>:18080/report/json`

  For more information, see *ExaCC gen2: Troubleshooting VPN/WSS connection from Customer Side*.

- **Configure the Oracle Exadata system networks**
  Each Oracle Exadata Database Service on Cloud@Customer system contains two system networks, which are not connected to your corporate network. The following settings define IP address allocations for these networks:

  - **Administration Network CIDR Block:** Specifies the IP address range for the administration network using CIDR notation. The administration network provides connectivity that enables Oracle to administer the Exadata system components, such as the Exadata compute servers, storage servers, network switches, and power distribution units. You can accept the suggested default, or specify a custom value.
    The minimum CIDR block is /23. The maximum CIDR block is /21. Because this CIDR block cannot be changed in the future, it is recommended to use a /21 CIDR block to maximize flexibility for future possible expansion. However, the below table can be used for further guidance.

  - **Secure Fabric Network:** Specifies the IP address range for the Exadata secure fabric network using CIDR notation. The Exadata secure fabric network provides the high-speed low-latency interconnect used by Exadata software for internal communications between various system components. You can accept the suggested default, or specify a custom value.
    The minimum CIDR block is /22. The maximum CIDR block is /19. Because this CIDR block cannot be changed in the future, it is recommended to use a /19 CIDR block to maximize flexibility for future possible expansion. However, the below table can be used for further guidance.

Table 5-7 X11M and X10M (Quarter Rack X10M, Quarter Rack X10M-L, and Quarter Rack X10M-XL) CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
6	64	/23	/22
7	45		
8	26		
9	9		
7	64	/22	/22
8	55		
9	30		

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-7 (Cont.) X11M and X10M (Quarter Rack X10M, Quarter Rack X10M-L, and Quarter Rack X10M-XL) CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
10	5		
15	64	/22	/21
16	60		
17	45		
18	30		
19	15		
17	64	/21	/21
18	61		
19	36		
20	11		
32	64	/21	/20

Table 5-8 X9M CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
6	64	/23	/22
7	48		
8	30		
9	9		
7	64	/22	/22
8	55		
9	30		
10	5		
15	64	/22	/21
16	62		
17	45		
18	30		
19	15		
17	64	/21	/21
18	61		
19	36		
20	11		
32	64	/21	/20

Table 5-9 X8M CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
8	64	/23	/22

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-9 (Cont.) X8M CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
9	61		
10	47		
11	37		
12	28		
13	13		
11	64	/22	/22
12	51		
13	34		
14	17		
23	64	/22	/21
24	56		
25	45		
26	38		
27	21		
26	64	/21	/21
27	52		
28	35		
29	18		
32	64	/21	/20

> **Note**
>
> You cannot change this value in the future. Consider all future expansions when making this choice. Otherwise, your future expandability may be limited.

- **Configure DNS and NTP services**
  Each Exadata Database Service on Cloud@Customer system requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. The following settings specify the servers that provide these services to the Exadata infrastructure:
  - **DNS Servers:** Provide the IP address of a DNS server that is accessible using the control plane network. You may specify up to three DNS servers.
  - **NTP Servers:** Provide the IP address of an NTP server that is accessible using the control plane network. You may specify up to three NTP servers.
- **Network bonding mode**

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> The use of LACP requires both the server and the switch(es) have compatible settings to allow LACP to work properly. To use LACP for ExaDB-C@C, you must configure your network switches such that they are compatible with the following parameter from the Linux `ifcfg-bondethx` configuration file:
>
> ```
> BONDING_OPTS="mode=802.3ad miimon=100 downdelay=200 updelay=200
> lacp_rate=1 xmit_hash_policy=layer3+4"
> ```
>
> Note that the `BONDING_OPTS` line in the host operating system cannot be changed, so customer switch settings must be compatible with the parameters stated above without alteration.

The client and backup networks are configured to use active-backup bonding (mode 1) by default. Since the setting gets applied to each network interface individually, you have the flexibility to configure client and backup network interfaces independently of each other. For example, you can configure the backup network to LACP and the client network to active-backup, if desired.

- **Enable LACP for Client Network**: Select to configure the client network with LACP mode 4 (mode=802.3ad) active/active dynamic link.
- **Enable LACP for Backup Network**: Select to configure the backup network with LACP mode 4 (mode=802.3ad) active/active dynamic link.

You can check the bonding mode of the client and backup networks under the Network section on the Exadata Infrastructure Details page.

- **Active/Backup** indicates that the network interfaces are configured to use the default mode 1.
- **LACP** indicates that the network interfaces are mode 4 enabled.

- **Provide maintenance details**
  - Click **Configure maintenance**.
    On the resulting Configure maintenance page, configure the following:
    - **Maintenance scheduling preference: Oracle managed schedule**
      - **Choose a maintenance method:**
        - **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
        - **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
      - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The

Copyright © 2008, 2026, Oracle and/or its affiliates.

maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.

- **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

> **Note**
>
> Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 30 minutes

Maximum: 120 minutes

- **Maintenance scheduling preference: Customer managed schedule**
  - **Maintenance schedule:** Define maintenance preferences for this infrastructure
    - **Configure maintenance preference:** Define maintenance time preferences for each quarter. If more than one preference is defined for a quarter, Oracle automation will select one of them to perform maintenance on all components in your infrastructure.
      Select at least one month every two quarters.
    - **Specify a schedule:** Choose your preferred week, weekday, start time, and lead time for infrastructure maintenance.
      - Optional. Under **Week of the month**, specify which week of the month, maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days of the month, and have a duration of 7 days. Weeks start and end based on calendar dates, not days of the week. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days. If you do not specify a week of the month, Oracle will run the maintenance update in a week to minimize disruption.
      - Optional. Under **Day of the week**, specify the day of the week on which the maintenance will occur. If you do not specify a day of the week, Oracle will run the maintenance update on a weekend day to minimize disruption.
      - Optional. Under **Hour of the day**, specify the hour during which the maintenance run will begin. If you do not specify a start hour, Oracle will pick the least disruptive time to run the maintenance update.
      - Under **Notification Lead Time**, specify the minimum number of weeks ahead of the maintenance event you would like to receive a notification message. Your lead time ensures that a newly

Copyright © 2008, 2026, Oracle and/or its affiliates.

released maintenance update is scheduled to account for your required minimum period of advanced notification.

* **Choose a maintenance method:**
  * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
  * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
* **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
  * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

    > **Note**
    >
    > Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

    Default: 30 minutes

    Maximum: 120 minutes
* **Show advanced options:**
  * Enable monthly security infrastructure maintenance: Select this check box to perform monthly security infrastructure maintenance.

* **Maintenance schedule:** Use maintenance window preferences from a scheduling policy. During infrastructure provisioning, after the scheduling policy is selected, Oracle generates a recommended maintenance scheduling plan to apply updates to all the components in your infrastructure. The recommended plan schedules all DB Servers, followed by Storage Servers and Network Switches, into the maintenance windows from your policy based on duration. After provisioning the infrastructure,

Copyright © 2008, 2026, Oracle and/or its affiliates.

you can update the scheduling plan by editing the 'Maintenance Scheduling Plan' resource and customize the update to specific components to align with different windows in your scheduling policy.

* Click **Select policy**.
* In the resulting Select maintenance scheduling policy window, choose a compartment and a policy.
  You can also create a maintenance scheduling policy and use it. For more information, see Create a Maintenance Scheduling Policy. Note that you can add additional maintenance windows to the policy after creating it. For more information, see Add Additional Maintenance Windows to a Maintenance Scheduling Policy.
* Click **Save**.

Click **Save**.

If you switch from rolling to non-rolling maintenance method, then **Confirm Non-rolling Maintenance Method** dialog is displayed.

Enter the name of the infrastructure in the field provided to confirm the changes.

Click **Save Changes**.

> **Note**
>
> After creating the infrastructure, you can find the maintenance method, maintenance schedule, DB Server version, and Storage Server version details under the **Maintenance** and **Version** sections on the **Infrastructure Details** page.

- **Show Advanced Options**

  You have the option to configure advanced options.

  – **Tags:** (Optional) You can choose to apply tags. If you have permission to create a resource, then you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, then skip this option (you can apply tags later) or ask your administrator.

6. Click **Create**.
   If all of your inputs are valid, then the Infrastructure Details page appears. The page outlines the next steps in the provisioning process. Initially, after creation, the state of the Oracle Exadata infrastructure is **Requires-Activation**.

Related Topics

- Understanding Compartments
- System Configuration Options for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.
- ExaCC gen2: Troubleshooting VPN/WSS connection from Customer Side
- Resource Tags

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Edit Network Bonding

To edit the client and backup network bonding mode, follow these steps.

> **Note**
>
> The use of LACP requires both the server and the switch(es) have compatible settings to allow LACP to work properly. To use LACP for ExaDB-C@C, you must configure your network switches such that they are compatible with the following parameter from the Linux `ifcfg-bondethx` configuration file:
>
> ```
> BONDING_OPTS="mode=802.3ad miimon=100 downdelay=200 updelay=200
> lacp_rate=1 xmit_hash_policy=layer3+4"
> ```
>
> Note that the `BONDING_OPTS` line in the host operating system cannot be changed, so customer switch settings must be compatible with the parameters stated above without alteration.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Click **Exadata Infrastructure**.
3. Click the name of the Exadata infrastructure that you want to edit.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
4. Click **Edit network bonding**.
   Edit Network Bonding page is displayed.
5. Select or clear the **Enable LACP for Client Network** and **Enable LACP for Backup Network** checkboxes.
   Apply Network Bonding Change dialog is displayed.

   > **Note**
   >
   > This action initiates a non-rolling update process and bonding modes are changed simultaneously on all Database Servers. You are required to manage your own switch settings on your network accordingly. You may experience a network outage and should plan for application downtime until your switch settings and ExaDB-C@C settings are made to match.

6. Click **Save** to apply the changes or **Cancel** to retain the existing settings.
   The state of Exadata Infrastructure changes to **Updating** state until the back-end change on the bonding mode setting is completed.

   You must change your switch settings to correspond to the `BONDING_OPTS`, as stated above. The state changes to **Waiting for Connectivity** until you make the correct switch changes. While in this state, the bonding mode validation is continuously re-tried until it passes successfully.

   When the network validation succeeds, then the state of the Exadata Infrastructure changes to **Active**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you need to revert your settings while in the "**Waiting for Connectivity**" state, you can edit the network bonding mode again using the same procedure described above. Note that Oracle has no control over your switch settings, so you will need to manage your switch settings to return them to the original settings if changed.

Using the Console to View Exadata Infrastructure Network Configuration Details

To view network configuration details, follow these steps. Save this information for later use to troubleshoot if you face network issues.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**.
2. Under **Region**, select the region that you want to associate with the Oracle Exadata infrastructure.
3. Click **Exadata Infrastructure**.
4. From the list of infrastructures, click the name of the infrastructure that you're interested in. Note that the infrastructure must be in **Active** state.
5. On the **Infrastructure Details** page, find the network configuration details under the **Network** section.

Using the Console to Edit Oracle Exadata Database Service on Cloud@Customer Infrastructure Networking Configuration

To edit your Oracle Exadata Database Service on Cloud@Customer infrastructure networking configuration, be prepared to provide values for the infrastructure configuration.

You can only edit Oracle Exadata Database Service on Cloud@Customer infrastructure networking configuration only if the current state of the Oracle Exadata infrastructure is **Requires Activation**. Also, ensure that you do not edit the Exadata infrastructure after you download the configuration file and provide it to Oracle.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Edit Infrastructure Networking**.
6. Use the Edit Infrastructure Networking dialog to edit the Oracle Exadata infrastructure networking:
   a. **Time zone:** The default time zone for the Exadata Infrastructure is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, then select the **Select another time zone** option, select a **Region** or **country**, and then select the corresponding **Time zone**.
>
> If you do not see the region or country you want, then select **Miscellaneous**, and then select an appropriate **Time zone**.

b. **Configure the cloud control plane network**

Each Oracle Exadata Database Service on Cloud@Customer system contains two Control Plane Servers, which enable connectivity to Oracle Cloud Infrastructure. The Control Plane Servers are connected to the control plane network, which is a subnet on your corporate network. The following settings define the network parameters:

- **Control Plane Server 1 IP Address:** Provide the IP address for the first control plane server. This IP address is for the network interface that connects the first Control Plane Server to your corporate network using the control plane network.
- **Control Plane Server 2 IP Address:** Provide the IP address for the second control plane server. This IP address is for the network interface that connects the second Control Plane Server to your corporate network using the control plane network.
- **Netmask:** Specify the IP netmask for the control plane network.
- **Gateway:** Specify the IP address of the control plane network gateway.
- **HTTP Proxy:** Optionally, you can use this field to specify your corporate HTTP proxy to use for the HTTPS connection from the Control Plane Server to Oracle Cloud Infrastructure. The expected format is:

  `http://server.domain:port`

  For example:

  `http://proxy.example.com:80`

  For enhanced security, when possible, Oracle recommends that you use an HTTP proxy.

c. **Configure the Exadata system networks**

Each Oracle Exadata Database Service on Cloud@Customer system contains two system networks, which are not connected to your corporate network. The following settings define IP address allocations for these networks:

- **Administration Network CIDR Block:** Specifies the IP address range for the administration network using CIDR notation. The administration network provides connectivity that enables Oracle to administer the Exadata system components, such as the Exadata compute servers, storage servers, network switches, and power distribution units.

  The maximum CIDR block prefix length is `/23`, which defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Oracle Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth. The minimum CIDR block prefix length is `/16`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Ensure that the IP address range does not conflict with other hosts your corporate network, and does not overlap with the InfiniBand network CIDR block.

- **Secure Fabric Network CIDR Block:** Specifies the IP address range for the Exadata InfiniBand network using CIDR notation. The Exadata InfiniBand network provides the high-speed low-latency interconnect used by Exadata software for internal communications between various system components.

  The maximum CIDR block prefix length is `/22`, which defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Oracle Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth. The minimum CIDR block prefix length is `/19`.

  Ensure that the IP address range does not conflict with other hosts your corporate network, and does not overlap with the administration network CIDR block.

d. **Configure DNS and NTP services**
Each Oracle Exadata Database Service on Cloud@Customer system requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. The following settings specify the servers that provide these services to the Exadata infrastructure:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the control plane network. You can specify up to three DNS servers.
- **NTP Servers:** Provide the IP address of an NTP server that is accessible using the control plane network. You may specify up to three NTP servers.

7. Click **Save Changes**.

Using the Console to Download a File Containing Configuration Data

To download an Oracle Exadata Database Service on Cloud@Customer configuration file, complete this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Oracle Exadata infrastructure for which you want to download a file containing the infrastructure configuration details.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure for which you want to download a file containing the infrastructure configuration details.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Show guide** on the banner displayed.
6. Click **Download Configuration**.

   Your browser downloads a file containing the infrastructure configuration details.

The generated configuration file includes all the relevant configuration details for the additional storage servers included as part of the create infrastructure flow.

When you provide the generated infrastructure configuration file to Oracle, ensure that it has not been altered in any way. Also, ensure that you do not edit the Oracle Exadata infrastructure after you download the configuration file and provide it to Oracle.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Activate Oracle Exadata Database Service on Cloud@Customer Infrastructure

To activate Oracle Exadata Database Service on Cloud@Customer infrastructure, ensure that you meet the prerequisites, and complete this procedure.

- Ensure that you have added a primary contact. You cannot activate your infrastructure without adding a primary maintenance contact.
- Locate the activation file. This file is supplied to you by Oracle after installation and initial configuration of your Oracle Exadata Database Service on Cloud@Customer system.
- Ensure that the current state of your infrastructure is **Requires Activation**. You can only activate Oracle Exadata if its state is **Requires Activation**.

1. Download the activation file.
2. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
3. Choose **Region** and **Compartment**, and select the region and compartment that contains the Oracle Exadata infrastructure that you want to activate.
4. Click **Exadata Infrastructure**.
5. Click the name of the Oracle Exadata infrastructure that you want to activate.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
6. Click **Show activation guide** on the banner displayed.

   The **Show activation guide** button is only available if the Oracle Exadata infrastructure requires activation. You cannot activate Oracle Exadata infrastructure multiple times.
7. Use the Activate Exadata Infrastructure dialog to upload the activation file, and then click **Activate**.

   The activation file includes all the relevant details for the additional storage servers included as part of the create infrastructure flow.

   After activation, the state of the Oracle Exadata infrastructure changes to **Active**.

Using the Console to Check the Status of Oracle Exadata Database Service on Cloud@Customer Infrastructure

To find the status of your Oracle Exadata Database Service on Cloud@Customer infrastructure, use this procedure to check the Infrastructure Details page.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Oracle Exadata infrastructure that you are interested in.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure that you are interested in.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. Check the icon on the Infrastructure Details page. The color of the icon and the text below it indicates the status of the Oracle Exadata infrastructure.
   - **Creating**: Yellow icon. The Oracle Exadata infrastructure definition is being created in the control plane.
   - **Requires Activation**: Yellow icon. The Oracle Exadata infrastructure is defined in the control plane, but it must be provisioned and activated before it can be used.
   - **Active**: Green icon. The Oracle Exadata infrastructure is successfully provisioned and activated.
   - **Deleting**: Gray icon. The Oracle Exadata infrastructure is being deleted by using the Console or API.
   - **Deleted**: Gray icon. The Oracle Exadata infrastructure is deleted, and is no longer available. This state is transitory. It is displayed for a short time, after which the Oracle Exadata infrastructure is no longer displayed.
   - **Activation Failed**: Red icon. An error condition currently prevents the activation of the Oracle Exadata infrastructure. Typically, this state is auto-correcting, and does not require user intervention.

Using the Console to Scale Infrastructure Compute and Storage

To scale infrastructure compute and storage, complete this procedure.

You can scale infrastructure compute and storage when the current state of the Oracle Exadata infrastructure is **Active** or **Requires Activation**.

> **Note**
>
> You can add database servers and storage servers together or add them individually.

You cannot scale infrastructure beyond the point you previously reserved network resources for when creating your infrastructure originally. The below table identifies the required network ranges. Note that the network addresses and ranges originally identified for these networks at initial infrastructure deployment cannot be changed. If you do not have sufficient network resources reserved already, you will not be able to scale the infrastructure at this point.

Table 5-10 X11M and X10M (Quarter Rack X10M, Quarter Rack X10M-L, and Quarter Rack X10M-XL) CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
6	64	/23	/22
7	45		
8	26		
9	9		
7	64	/22	/22
8	55		
9	30		
10	5		
15	64	/22	/21

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-10 (Cont.) X11M and X10M (Quarter Rack X10M, Quarter Rack X10M-L, and Quarter Rack X10M-XL) CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
16	60		
17	45		
18	30		
19	15		
17	64	/21	/21
18	61		
19	36		
20	11		
32	64	/21	/20

Table 5-11 X9M CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
6	64	/23	/22
7	48		
8	30		
9	9		
7	64	/22	/22
8	55		
9	30		
10	5		
15	64	/22	/21
16	62		
17	45		
18	30		
19	15		
17	64	/21	/21
18	61		
19	36		
20	11		
32	64	/21	/20

Table 5-12 X8M CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
8	64	/23	/22
9	61		

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-12 (Cont.) X8M CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
10	47		
11	37		
12	28		
13	13		
11	64	/22	/22
12	51		
13	34		
14	17		
23	64	/22	/21
24	56		
25	45		
26	38		
27	21		
26	64	/21	/21
27	52		
28	35		
29	18		
32	64	/21	/20

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Oracle Exadata infrastructure that you are interested in.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure for which you want to download a file containing the infrastructure configuration details.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Scale Infrastructure**.
6. Select the number of database servers from the **Additional database servers** field.
7. Select the number of storage servers from the **Additional storage servers** field.
8. Select the **Multi-rack deployment** checkbox, if you want the infrastructure to span multiple racks.
9. Click **Scale**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Download Scale Infrastructure Compute and Storage Configuration File

To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Oracle Exadata infrastructure that you are interested in.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure for which you want to download a file containing the infrastructure configuration details.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Download Configuration**.
   Your browser downloads a file containing the infrastructure configuration details.

> **Note**
>
> When you provide the generated infrastructure configuration file to Oracle, ensure that it has not been altered in any way. Also, ensure that you do not edit the Oracle Exadata infrastructure after you download the configuration file and provide it to Oracle.

Using the Console to Activate New Database and Storage Servers

To download an Oracle Exadata Cloud@Customer scale configuration file, complete this procedure.

Upload the activation file once the field engineer finishes deploying the compute and storage servers and shares the activation file with you.

> **Note**
>
> Once the activation file is uploaded and the activate process is initiated, you cannot change the Scale Infrastructure request or cancel this whole operation. If there is an activation failure, then contact the field engineer to resolve the issue.

1. Download the activation file.
2. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
3. Choose the **Region** and **Compartment** that contains the Oracle Exadata infrastructure that you are interested in.
4. Click **Exadata Infrastructure**.
5. Click the name of the Oracle Exadata infrastructure that you want to activate.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.

6. Click **Activate New Storage Server(s)**.
   The Activate button is only available if the Oracle Exadata infrastructure requires activation. You cannot activate Oracle Exadata infrastructure multiple times.
7. Use the **Activate New Server** dialog to upload the activation file, and then click **Activate Now**.
   After activation, the state of the Oracle Exadata infrastructure changes to **Active**.

Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption

To make storage capacity from the new servers for VM clusters consumption, complete this procedure.

1. Download the activation file.
2. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
3. Choose the **Region** and **Compartment** that contains the Oracle Exadata infrastructure that you are interested in.
4. Click **Exadata Infrastructure**.
5. Click the name of the Oracle Exadata infrastructure that you want to activate.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
6. Click **Activate New Storage Server(s)**.
7. Review the advisory on the **Activate New Server** dialog, and then click **Activate**.

> **Note**
>
> While **Activate New Server** operation is in progress,
>
> - The system rebalances the storage to ensure that the capacity from the newly added storage servers is available for VM Cluster consumption. This rebalance will run with a ASM power limit of 4 independent of what you have set within your ASM configuration. If you would like the rebalance to complete faster, you can update the power limit after the rebalance process has begun. Monitor the ASM rebalance process to ensure it successfully completes (which will allow the add storage capacity workflow to complete).
> - You cannot create or delete VM Clusters.
> - Existing VM Clusters provisioned are in the Available life cycle state. However, they do not support scale up or down of resources allocated to the VM cluster, except for OCPU allocation. OCPU allocation changes are allowed even while **Activate New Server** operation is in progress.
> - During the add storage operation, if ASM detects insufficient space in the VM cluster to perform the disk group shrinking operation, it automatically allocates additional space to the VM to complete the process. This space is then deducted from the total storage being added to the infrastructure. The VM cluster reflects the additional space allocated to the VM.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to View Details of Oracle Exadata Database Service on Cloud@Customer Infrastructure with Scaled Compute and Storage Capacity

To view the storage capacity from the new storage server, or scaled compute capacity, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Oracle Exadata infrastructure that you are interested in.
3. Click **Exadata Infrastructure**.
   After changing the storage capacity, **Shape** will change to either **Elastic** or **Elastic Base**.

> **Note**
>
> When you perform scale VM Cluster Exadata storage operation, newly added Exadata storage capacity is also available for consumption. Similarly, when you create a VM cluster, the Console displays the newly added Exadata storage capacity (shared Exadata storage) as available storage.

Related Topics

- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Using the Console to Move Oracle Exadata Database Service on Cloud@Customer Infrastructure

To relocate Oracle Exadata Database Service on Cloud@Customer infrastructure to another compartment, use this procedure.

You can change the compartment that contains your Exadata Database Service on Cloud@Customer infrastructure by moving it.

When you move Exadata infrastructure, the compartment change is also applied to the associated VM cluster networks. However, the compartment change does not affect any other associated resources, such as the VM clusters, which remain in their current compartment.

To move Oracle Exadata Database Service on Cloud@Customer infrastructure:

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and compartment that contains the Oracle Exadata infrastructure that you want to move.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to move.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.

5. Click **More actions**, and then select **Move Resource**.
6. In the resulting dialog, choose the new compartment for the Oracle Exadata infrastructure, and click **Move Resource**.

Using the Console to Terminate Oracle Exadata Database Service on Cloud@Customer Infrastructure

To terminate Oracle Exadata Database Service on Cloud@Customer infrastructure, complete the prerequisites, and then complete this procedure.

Terminating Exadata Database Service on Cloud@Customer infrastructure removes it from the Cloud Control Plane.

If you are terminating Oracle Exadata infrastructure before activation, then if required, you can create replacement Oracle Exadata infrastructure without any input from Oracle.

If you are terminating active Oracle Exadata infrastructure, then to create replacement Oracle Exadata infrastructure, you must repeat the full provisioning process, including the tasks that Oracle performs.

Before you can terminate active Exadata infrastructure, you must:

- Terminate all of the resources that it contains, including the databases, VM cluster, and VM cluster network.
- Lodge a service request (SR) with Oracle indicating your intention to terminate the Oracle Exadata infrastructure. In response to the SR, Oracle flags the Oracle Exadata infrastructure as ready for termination.

After Oracle has flagged the Oracle Exadata infrastructure, terminate the Exadata infrastructure by using the following process:

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and compartment that contains the Oracle Exadata infrastructure that you want to terminate.
3. Click **Exadata Infrastructure**.
4. Click the name of the Oracle Exadata infrastructure that you want to terminate.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Actions**, and then **Terminate**.
6. In the resulting dialog, enter the Oracle Exadata infrastructure name and click **Terminate Exadata Infrastructure** to confirm the action.

Using the Console to Manage Tags for Your Oracle Exadata Database Service on Cloud@Customer Resources

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**.
2. Choose your **Compartment**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

3. Find the Exadata Infrastructure, VM Cluster Network, VM Cluster, Backup Destination, Database Home, or Database resource you're interested in, and click the name.
4. Click the **Tags** tab to view or edit the existing tags. Or, click **More Actions** and then **Apply Tags** to add new ones.

Related Topics

- Resource Tags

Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure

Oracle Exadata Database Service on Cloud@Customer uses the same API as Oracle Cloud 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 Exadata Database Service on Cloud@Customer infrastructure:

- `ActivateExadataInfrastructure`
- `CreateExadataInfrastructure`
- `DeleteExadataInfrastructure`
- `DownloadExadataInfrastructureConfigFile`
- `GenerateRecommendedVmClusterNetwork`
- `GetExadataInfrastructure`
- `ListExadataInfrastructure`
- `UpdateExadataInfrastructure`
- `AddStorageCapacityExadataInfrastructure`

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- ActivateExadataInfrastructure
- CreateExadataInfrastructure
- DeleteExadataInfrastructure
- DownloadExadataInfrastructureConfigFile
- GenerateRecommendedVmClusterNetwork
- GetExadataInfrastructure
- ListExadataInfrastructure
- UpdateExadataInfrastructure
- AddStorageCapacityExadataInfrastructure

Copyright © 2008, 2026, Oracle and/or its affiliates.

Configure Oracle-Managed Infrastructure Maintenance

Oracle performs the updates to all of the Oracle-managed infrastructure components on Oracle Exadata Database Service on Cloud@Customer.

Oracle performs periodic maintenance on your Exadata Cloud@Customer infrastructure to ensure it is free of potential issues that could affect database availability, integrity, and security. Updating the software running on the infrastructure with the latest product and security fixes protects your data and the overall compliance of the Oracle Cloud. These updates are performed in an automated way and include all the best practices, relieving you of the need to invest any effort in maintaining your infrastructure. Oracle updates include the physical database server hosts, storage servers, network fabric switches, management switches, power distribution units (PDUs), integrated lights-out management (ILOM) interfaces, and control plane servers.

Oracle performs two types of infrastructure maintenance:

- Quarterly maintenance is applied every three months and can include product fixes, enhancements, and security fixes.
- Monthly maintenance only applies critical security fixes that can be applied online to ensure components are maintained at the highest security standards with any security vulnerabilities fixed as soon as possible.

Quarterly Maintenance

Oracle minimizes the impact of quarterly maintenance on your applications using rolling maintenance operations, preserving database availability throughout the update process. Rolling maintenance reboots each database server, one at a time, with at most one server offline at any time. Applications designed for high availability automatically and transparently migrate their database connections between available database instances without disruption, eliminating the need for scheduling downtime. Storage server updates are also applied in a rolling manner. Rebooting storage servers has no effect on the database service, and thus has no impact on your applications.

Oracle allows you to fully control quarterly maintenance schedules, so you can schedule maintenance during a period which will have the least impact on your business users. You have full control and visibility over when quarterly maintenance will be applied, even allowing you to schedule the maintenance across multiple maintenance windows. Scheduling is simplified using a Maintenance Scheduling Policy, which aims to standardize scheduling across the fleet to ensure consistency and efficiency. By defining the policy once and applying it to multiple resources, it streamlines the scheduling process. You may also reschedule maintenance should unexpected business issues occur.

Monthly Security Maintenance

Monthly security maintenance is performed on the database servers online, with no reboot, and no impact to your applications. Monthly updates are applied to storage servers in a rolling manner, also with no impact to your applications.

Monthly security maintenance can also be scheduled at a specific time during the month, albeit in a single maintenance window. Oracle will publish a schedule for monthly maintenance at least one week prior to start of the maintenance period, and you can reschedule if required.

You may manage contacts who are notified regarding infrastructure maintenance, set a maintenance window to determine the time your quarterly infrastructure maintenance will begin, and also view scheduled maintenance runs and the maintenance history of your Exadata Cloud@Customer in the Oracle Cloud Infrastructure Console. For details regarding

Copyright © 2008, 2026, Oracle and/or its affiliates.

the infrastructure maintenance process and configuring the maintenance controls refer to the following:

- About Oracle Managed Oracle Exadata Database Service on Cloud@Customer Infrastructure Maintenance Updates
  Oracle performs patches and updates to all of the Oracle-managed system components on Oracle Exadata Database Service on Cloud@Customer.
- Infrastructure Maintenance Contacts
  Maintenance contacts are required for service request based communications for hardware replacement and other maintenance events.
- Maintenance Scheduling Policy
  Learn how to use the OCI Console to configure and manage maintenance scheduling policies. If a maintenance scheduling policy is used, all the scheduling preferences for the infrastructure maintenance are derived from the policy.
- Using the Console to Configure Oracle-Managed Infrastructure Updates
  Full Exadata infrastructure software updates are scheduled on a quarterly basis. In addition, important security updates are scheduled monthly. While you cannot opt-out of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal and allows scheduling flexibility to help you plan for them.
- Manage Quarterly/Monthly Maintenance Run created from Scheduling Plan
- Monitor Infrastructure Maintenance Using Lifecycle State Information
  The lifecycle state of your Exadata Infrastructure resource enables you to monitor when the maintenance of your infrastructure resource begins and ends.
- Receive Notifications about Your Infrastructure Maintenance Updates
  There are two ways to receive notifications. One is through email to infrastructure maintenance contacts and the other one is to subscribe to the maintenance events and get notified.
- Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure Maintenance Controls
  Oracle Exadata Database Service on Cloud@Customer uses the same API as Oracle Cloud Infrastructure to manage infrastructure maintenance controls.

About Oracle Managed Oracle Exadata Database Service on Cloud@Customer Infrastructure Maintenance Updates

Oracle performs patches and updates to all of the Oracle-managed system components on Oracle Exadata Database Service on Cloud@Customer.

In all but rare exceptional circumstances, you receive advance communication about these updates to help you plan for them. If there are corresponding recommended updates for your VM cluster virtual machines (VMs), then Oracle provides notifications about them.

Wherever possible, scheduled updates are performed in a manner that preserves service availability throughout the update process. However, there can be some noticeable impact on performance and throughput while individual system components are unavailable during the update process.

For example, database server patching typically requires a reboot. In such cases, wherever possible, the database servers are restarted in a rolling manner, one at a time, to ensure that the service remains available throughout the process. However, each database server is unavailable for a short time while it restarts, and the overall service capacity diminishes accordingly. If your applications cannot tolerate the restarts, then take mitigating action as needed. For example, shut down an application while database server patching occurs.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Overview of the Quarterly Infrastructure Maintenance Process
  By default, infrastructure maintenance updates the Exadata database server hosts in a rolling fashion, followed by updating the storage servers.
- Overview of Monthly Security Maintenance
  Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.
- Understanding Monthly and Quarterly Maintenance in the Same Month

Overview of the Quarterly Infrastructure Maintenance Process

By default, infrastructure maintenance updates the Exadata database server hosts in a rolling fashion, followed by updating the storage servers.

Rolling infrastructure maintenance begins with the Exadata database server hosts. For the rolling maintenance method, database servers are updated one at a time. Each of the database server host's VMs is shut down, the host is updated, restarted, and then the VMs are started, while other database servers remain operational. This rolling maintenance does not impact applications designed for high availability. Older applications not written to handle a rolling instance restart can be impacted. This process continues until all servers are updated.

After database server maintenance is complete, storage server maintenance begins. For the rolling maintenance method, storage servers are updated one at a time and this does not impact database or application availability. However, the rolling storage server maintenance can result in reduced IO performance as storage servers are taken offline one at a time (reducing available IO capacity) and resynced when brought back online (small overhead on database servers). Properly sizing the database and storage infrastructure to accommodate increased work distributed to database and storage servers that are not under maintenance will minimize (or eliminate) any performance impact.

You can also choose non-rolling maintenance to update database and storage servers. The non-rolling maintenance method first updates your storage servers at the same time, then your database servers at the same time. Although non-rolling maintenance minimizes maintenance time, it incurs full system downtime while the storage servers and database servers are being updated.

Note that while databases are expected to be available during the rolling maintenance process, the automated maintenance verifies Oracle Clusterware is running but does not verify that all database services and pluggable databases (PDBs) are available after a server is brought back online. The availability of database services and PDBs after maintenance can depend on the application service definition. For example, a database service, configured with certain preferred and available nodes, may be relocated during the maintenance and wouldn't automatically be relocated back to its original node after the maintenance completes. Oracle recommends reviewing the documentation on *Achieving Continuous Availability for Your Applications* on Exadata Cloud Systems to reduce the potential for impact to your applications. By following the documentation's guidelines, the impact of infrastructure maintenance will be only minor service degradation as database servers are sequentially updated.

Oracle recommends that you follow the *Maximum Availability Architecture (MAA) best practices* and use Data Guard to ensure the highest availability for your critical applications. For databases with Data Guard enabled, Oracle recommends that you separate the maintenance windows for the infrastructure instances running the primary and standby databases. You may also perform a switchover prior to the maintenance operations for the infrastructure instance hosting the primary database. This allows you to avoid any impact on your primary database during infrastructure maintenance.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Prechecks are performed on the Exadata Cloud@Customer infrastructure components prior to the start of the maintenance window. The goal of the prechecks is to identify issues that may prevent the infrastructure maintenance from succeeding. The Exadata infrastructure and all components remain online during the prechecks. An initial precheck is run approximately two weeks prior to the maintenance start and another precheck is run approximately 24 hours prior to maintenance start. If the prechecks identify an issue that requires rescheduling the maintenance notification is sent to the maintenance contacts.

Review MOS Note KB181723 for the Exadata Infrastructure maintenance timeframes.

Related Topics

- Achieving Continuous Availability For Your Applications
- Maximum Availability Architecture (MAA) Best Practices
- https://support.oracle.com/rs?type=doc&id=2333222.1

Overview of Monthly Security Maintenance

Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.

> **Note**
>
> For more information about the CVE release matrix, see Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1).
> To view the CVE release matrix specific to an Exadata Infrastructure version, click the Exadata version, for example, Exadata 23. Version-specific CVE release matrices are listed in the **Notes** column of the table.

Security maintenance, when needed, is scheduled to be applied during a 21-day window that begins between the 18th-21st of each month and will run till the 9th-12th of the next month. Customers will receive notification of the proposed schedule at least 7 days before the start of the monthly maintenance window and can reschedule monthly maintenance to another date in the window if desired. The monthly security maintenance process updates the physical database servers to fix critical security vulnerabilities and critical product issues. Monthly maintenance also updates storage servers to an Exadata Storage Software image that resolves known security vulnerabilities and product issues. No updates are applied to the customer-managed guest VMs. Monthly maintenance also updates storage servers to an Exadata Storage Software image that resolves known security vulnerabilities and product issues.

Updates to database servers are applied online via Ksplice technology, and have no impact to workloads running on the compute (database) servers, as database server security updates are applied online to the host server while your VM and all processes within the VM, including databases, remain up and running. Servers and VMs are not restarted. Updates to storage servers are applied in a rolling fashion. As with quarterly maintenance, the impact of rebooting storage servers should be minimal to applications.

While updating your services infrastructure, some operations including memory, and storage scaling, operating system and Grid Infrastructure patching (including prechecks), and elastic expansion of compute and storage servers may be blocked.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> CPU scaling to zero cannot be performed during an active Monthly Security Maintenance window, which starts 10 minutes before the scheduled start time.
> If a VM is shut down as part of a standard VM stop operation, the Control Plane does not perform any additional actions. This behavior applies only when the VM shutdown is initiated as part of a CPU scale-to-zero operation.

Please plan to defer these operations until after the updates are complete. If you attempt an affected operation, the console will notify you of the ongoing security updates. No software is updated in the guest VMs.

Related Topics

- https://support.oracle.com/rs?type=doc&id=888828.1
- View or Edit a Scheduled Security Maintenance
  Learn how to view and edit the next scheduled security maintenance.
- Exadata Infrastructure Maintenance Event Types
  Review the list of event types that Exadata Infrastructure Maintenance emits.
- https://support.oracle.com/rs?type=doc&id=2333222.1

Understanding Monthly and Quarterly Maintenance in the Same Month

Special considerations are made when both quarterly and monthly security maintenance are scheduled to run in the same month. Quarterly maintenance will reapply any security fixes already applied by security maintenance, and neither quarterly nor monthly maintenance will apply a storage server update if the existing storage server version is the same or newer than the version contained in the update.

- The contents of the updates applied during quarterly maintenance are determined at the start of the maintenance quarter and use the latest Exadata release from the month prior to the start of the maintenance quarter. If any additional security fixes are available at that time, those updates are included in the quarterly maintenance. That image is then used throughout the quarter. For example, the January release is used for quarterly maintenance in Feb, March, and April.
- When quarterly maintenance is applied it is possible there are security updates previously installed on the database servers are not included in the quarterly maintenance release to be applied. In that case, the automation will apply the same security fixes to new release installed by the quarterly maintenance so there will not be any regression in security fixes. If the current image on the storage server is the same or newer than that to be applied by the quarterly or monthly security maintenance, that maintenance will be skipped for the storage servers.

If quarterly maintenance is scheduled within 24 hours of the time the monthly is scheduled, the scheduled monthly maintenance will be skipped, and the monthly update will instead be applied immediately following the quarterly maintenance.

- When scheduled at the same time, the monthly update is executed immediately following the completion of the quarterly maintenance.
- If monthly maintenance is scheduled to begin 0-24 hours ahead of the quarterly maintenance, then the monthly maintenance will not execute as scheduled, but instead, wait and be executed immediately following the quarterly maintenance. If the quarterly

Copyright © 2008, 2026, Oracle and/or its affiliates.

maintenance is subsequently rescheduled, then the monthly security maintenance will begin immediately. Oracle, therefore, recommends scheduling quarterly and monthly maintenance at the same time. As a result, if you reschedule the quarterly at the last moment, the monthly maintenance will run at the scheduled time instead of immediately upon editing the schedule. You can also reschedule the monthly security maintenance when rescheduling the quarterly maintenance as long as you keep the monthly within the current maintenance window. Monthly maintenance can be rescheduled to another time in the maintenance window, but cannot be skipped.

Monthly Security Maintenance before Quarterly Maintenance

- To apply security maintenance before quarterly maintenance, reschedule the monthly security maintenance to occur more than 24 hours prior to the quarterly maintenance. The security maintenance will online apply security patches to the database servers with no impact to applications, and apply an update to the storage servers with minimal to no impact (may be slight performance degradation) on applications. The quarterly maintenance will follow as scheduled, and will perform rolling maintenance on the database servers, which will impact applications not written to handle a rolling reboot. As part of the quarterly maintenance, it will apply the same security updates to the database server that are already installed on the system (no security regression).
- If you are concerned about getting the latest security updates applied, schedule the monthly security maintenance to run after the new monthly maintenance window opens (usually on the 21st of the month).
- The impact of the monthly security maintenance rebooting the storage servers should be minimal, so impact to the applications during this month will only be due to the restart of the database servers during the quarterly maintenance. However, if you must coordinate a maintenance window with your end users for the security maintenance, this will require two maintenance windows.

Quarterly Maintenance before Monthly Security Maintenance

- To run the quarterly maintenance before the monthly security maintenance, reschedule the security maintenance to run no earlier than 24 hours before the quarterly maintenance is scheduled to start. The security maintenance will be deferred until the quarterly maintenance is completed. The quarterly maintenance will perform rolling maintenance on the database servers, which will impact applications not written to handle a rolling reboot. The quarterly maintenance may or may not skip the storage server patching. That depends on if it is newer or older than the release currently installed. In most cases, the version installed should be newer than the version associated with the quarterly maintenance. Exceptions to this rule may occur if it is the first month of a maintenance quarter, or you skipped the security maintenance in one or more prior months. The security maintenance will run either immediately after the quarterly maintenance is completed, or when scheduled, whichever is later. It will apply online updates to the database servers (no application impact) and will likely update the storage servers in a rolling manner. In some corner cases. the quarterly maintenance may contain the same storage server release as the security maintenance and the security maintenance storage server updates will be skipped.
- The impact to end users of running the quarterly maintenance before the security maintenance should be roughly the same as running the security maintenance first. The quarterly maintenance will be a disruptive event, but the security maintenance rebooting the storage servers should cause minimal disruption, and the security maintenance is applied to the database servers online. However, if you must coordinate a maintenance window with your end users for the security maintenance, this will require two maintenance windows. You can schedule those two maintenance windows to run back-to-back, to appear as single maintenance window to end users. To do this, reschedule the security

Copyright © 2008, 2026, Oracle and/or its affiliates.

maintenance to start at the same time (or up to 24 hours prior) as the quarterly maintenance. The security maintenance will be deferred until the quarterly maintenance is completed. Assuming you have been regularly applying monthly security maintenance, the storage servers will be skipped by the quarterly maintenance and will be updated by the security maintenance immediately upon the completion of the quarterly maintenance.

Minimizing Maintenance Windows

- To minimize the number of maintenance windows (you have to negotiate those with end users), schedule the quarterly maintenance and monthly maintenance at the same time. The security maintenance will be blocked. The quarterly maintenance will update the database servers in a rolling manner and will most likely skip the storage server. The security maintenance will follow up immediately and update the database servers online and the storage servers in a rolling manner. The result is a single database and storage server restart in a single maintenance window.
- There are two exceptions to this. 1. If the quarterly and monthly maintenance contain the same storage server release, the quarterly maintenance will apply the storage server update, and the security maintenance will be skipped. From your perspective, this is still a single rolling reboot in a single maintenance window. 2. The currently installed release on the storage servers is older than that contained in the quarterly maintenance, which in turn is older than that in the security maintenance. That would cause the quarterly maintenance to update the storage, and then the security maintenance to do it as well. This can only happen if you skipped a prior month's security maintenance, because it requires the current image to be at least 2 months out of date. In such a scenario, you may want to schedule the security maintenance first and then the quarterly maintenance. This would result in one storage server reboot, but two distinct maintenance windows — the first for the security maintenance, and then later the quarterly maintenance.
- To minimize the impact to your end users, always apply the monthly security updates, and in months where both are scheduled, schedule them at the same time.

> **Note**
>
> - If the Exadata Infrastructure is provisioned before Oracle schedules the security maintenance, then it will be eligible for security maintenance.
> - Any time before the scheduled monthly Exadata Infrastructure maintenance, you can reschedule it.

Infrastructure Maintenance Contacts

Maintenance contacts are required for service request based communications for hardware replacement and other maintenance events.

Add a primary maintenance contact and optionally add a maximum of nine secondary contacts. Both the primary and secondary contacts receive all notifications about hardware replacement, network issues, and software maintenance runs.

You can promote any secondary contacts as the primary anytime you want. When you promote a secondary contact to primary, the current primary contact will be demoted automatically to secondary.

For more information, see: *Using the Console to Create Infrastructure* and *Managing Infrastructure Maintenance Contacts*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Using the Console to Create Oracle Exadata Database Service on Cloud@Customer Infrastructure
  To create your Oracle Exadata Database Service on Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Defining Maintenance Contacts for Exadata Cloud@Customer
  Oracle Cloud Operations will use these contacts to communicate with customers about infrastructure-related events, including hardware failures and replacements, quarterly maintenance patches, and other operational updates or incidents

Maintenance Scheduling Policy

Learn how to use the OCI Console to configure and manage maintenance scheduling policies. If a maintenance scheduling policy is used, all the scheduling preferences for the infrastructure maintenance are derived from the policy.

The Maintenance Scheduling Policy aims to standardize scheduling across the fleet, ensuring consistency and efficiency. By defining the policy once and applying it to multiple resources, it streamlines the scheduling process. The policy aligns with business best practices, scheduling maintenance activities in accordance with these standards. It also serves as a central repository for documenting and coordinating maintenance commitments with various stakeholders, enhancing compliance and efficiency. The centralized management of resources subscribing to the policy ensures adherence to compliance requirements, and any changes can be efficiently coordinated from a single location. Additionally, the policy improves communication about planned maintenance across different environments in the fleet, facilitating better coordination and awareness.

> **Note**
>
> If you use a scheduling policy, scheduling preferences defined locally on the infrastructure are not used by Oracle automation to apply quarterly updates.

- View the List of Maintenance Scheduling Policy
- Create a Maintenance Scheduling Policy
- View Details of a Maintenance Scheduling Policy
- Edit a Maintenance Scheduling Policy
- Edit a Maintenance Scheduling Policy that's in Needs Attention State
- View the Maintenance Windows of a Maintenance Scheduling Policy
- Edit a Maintenance Window of a Maintenance Scheduling Policy
- Add Additional Maintenance Windows to a Maintenance Scheduling Policy
- View Resources Associated with a Policy
- Delete a Maintenance Window of a Maintenance Scheduling Policy
- Move a Maintenance Scheduling Policy to a Different Compartment
- Add Tags to a Maintenance Scheduling Policy
- Delete a Maintenance Scheduling Policy

Copyright © 2008, 2026, Oracle and/or its affiliates.

View the List of Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.
3. Use the **Compartment** filter to view the list of policies created in a specific compartment.
4. Use the **State** to filter the policies by their states.

   The policy states include:

   - Creating
   - Needs attention
   - Available
   - Updating
   - Failed
   - Deleting
   - Deleted

Create a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.
3. Click **Create maintenance scheduling policy**.
4. In the resulting Create maintenance scheduling policy window, enter the following:
   a. **Name:** Enter a descriptive name for the policy.
   b. **Compartment:** Choose a compartment where you want to create this resource.
   c. **Cadence:** Choose a frequency.

      Available options:

      i. Every quarter

      > **Note**
      >
      > To use the scheduling policy to plan and automate quarterly infrastructure maintenance, you must create a policy with a 'Quarterly' cadence and a scheduled start month value set to 'February.'

      ii. Every month

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> To plan and automate monthly infrastructure maintenance using a scheduling policy, configure the policy with an 'Every month' cadence. The start month is not significant for monthly schedules and can be set to any month.

d. **Schedule start month:** Choose a start month.
Oracle automation requires a quarterly maintenance schedule to start in February if this policy is used for infrastructure maintenance.

e. **Maintenance window:** Oracle automation will perform maintenance to run scheduled actions as per your maintenance windows defined in the schedule.

   i. **Month:** (Applies to Every Quarter cadence) Choose a month. The month selection option depends on the selected policy cadence. For a quarterly cadence, you can choose between the First, Second, or Third months of every quarter.

   ii. **Week:** A week for maintenance purposes starts on the 1st, 8th, 15th, and 22nd days of the month.

   iii. **Day:** You can choose between the week's days: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.

   iv. **Start time:** Specify the start time when you want to begin maintenance.

   v. **Duration:** Plan your maintenance window based on average time estimates to complete Database Cloud service actions. Refer to Maintenance timeframes for updating the database server and storage servers during Exadata Infrastructure maintenance (KB181723) for time estimates.

   vi. **Enforce window duration:** With this option enabled, any scheduled action that goes over the configured window duration will be paused and re-scheduled to resume in a future maintenance window.

> **Note**
>
> If an update or action is already underway and cannot be paused and resumed without causing disruption, we will continue and complete the action. Oracle automation will reschedule any action planned to start after the configured window duration to a future maintenance window.

f. **Show Advanced Options:**

   i. **Tags:** (Optional) You can choose to apply tags. If you have permission to create a resource, then you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, then skip this option (you can apply tags later) or ask your administrator.

> **Note**
>
> You can add additional maintenance windows to the scheduling policy after creation.

For detailed guidance on performing quarterly infrastructure maintenance using the scheduling policy,see Using the Console to Configure Oracle-Managed Infrastructure Updates.

Copyright © 2008, 2026, Oracle and/or its affiliates.

View Details of a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Maintenance scheduling policy page displays the list of policies.
3. Click the name of the policy that you want to view details.

Edit a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to edit.
4. On the resulting Maintenance scheduling policy details page, click **Edit scheduling policy**.
5. On the resulting Edit scheduling policy window, you can only edit the name of the policy.
6. Edit the name and then click **Save changes**.

Edit a Maintenance Scheduling Policy that's in Needs Attention State

> **Note**
>
> You cannot use a scheduling policy without maintenance windows to plan and automate maintenance activity; hence, it has a 'Needs Attention' life cycle state. While the policy has no windows defined, you can change all its properties, including cadence and schedule start month.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that's in **Needs attention** state.
4. On the resulting Maintenance scheduling policy details page, click **Edit scheduling policy**.
5. On the resulting Edit scheduling policy window, update the policy settings.
6. Click **Save changes**.

View the Maintenance Windows of a Maintenance Scheduling Policy

Copyright © 2008, 2026, Oracle and/or its affiliates.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to view Maintenance window details.
4. Click the **Maintenance windows** tab to view the maintenance windows associated with the policy.

Edit a Maintenance Window of a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to edit Maintenance window details.
4. Click the **Maintenance windows** tab to view the maintenance windows associated with the policy.
5. Click the Actions menu (three dots) of the Maintenance window you want to edit and then select **Edit maintenance window**.
6. On the resulting Edit maintenance window page, make changes and then click **Save**.

Add Additional Maintenance Windows to a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to add Maintenance windows.
4. Click the **Maintenance windows** tab to view the maintenance windows associated with the policy.
5. Click **Add maintenance window**.

   > **Note**
   >
   > All maintenance windows in a scheduling policy should be no more than 3 weeks apart.

6. In the resulting Add maintenance window page, add details and then click **Add**.

View Resources Associated with a Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The resulting Scheduling Policy page displays the list of policies.

3. Click the name of the policy that you want to view the associated resources.

   The **Associated resources** section in the resulting Maintenance scheduling policy details page lists the resources associated with the policy.

4. Click the **Associated resources** tab to view the resources associated with the policy.

Delete a Maintenance Window of a Maintenance Scheduling Policy

> **Note**
>
> Only maintenance windows not used by any resources to plan and automate maintenance activity can be deleted from the policy. Any window already used by services to automate maintenance cannot be deleted. Deleting the window will also delete all its associated empty actions.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.

3. Click the name of the policy that you want to delete Maintenance windows.
4. Click the **Maintenance windows** tab to view the maintenance windows associated with the policy.
5. Click the Actions menu of the Maintenance window you want to delete and then select **Delete**.
6. On the resulting Delete maintenance window dialog, click **Delete** to confirm the action.

Move a Maintenance Scheduling Policy to a Different Compartment

> **Note**
>
> You cannot move a policy across compartments if any resource uses it to plan and automate maintenance activity.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.

3. Click the name of the policy that you want to move to a different compartment.
4. On the resulting Maintenance scheduling policy details page, click the Actions menu and then select **Move resource**.
5. On the resulting Move resource to a different compartment dialog, choose a compartment, and then click **Move resource**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Add Tags to a Maintenance Scheduling Policy

> **Note**
>
> You cannot move a policy across compartments if any resource uses it to plan and automate maintenance activity.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to add tags to.
4. In the resulting Maintenance scheduling policy details page, click the Actions menu and then select **Add tags**.
5. In the resulting Add tags dialog, add tags, and then click **Add**.

Delete a Maintenance Scheduling Policy

> **Note**
>
> You cannot delete a policy if any resource is using it to plan and automate maintenance activity. Deleting the policy will also delete all its associated empty windows.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

   The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to delete.
4. On the resulting Maintenance scheduling policy details page, click the Actions menu and then select **Delete**.
5. On the resulting Delete maintenance policy dialog, enter the name of the policy, and then click **Delete**.

> **Note**
>
> You cannot delete a policy if it's used in a maintenance scheduling plan.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Configure Oracle-Managed Infrastructure Updates

Full Exadata infrastructure software updates are scheduled on a quarterly basis. In addition, important security updates are scheduled monthly. While you cannot opt-out of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal and allows scheduling flexibility to help you plan for them.

For quarterly infrastructure maintenance, you can set a maintenance window to determine when the maintenance will begin. You can also edit the maintenance method, enable custom action, and view the scheduled maintenance runs and the maintenance history of your Exadata Cloud@Customer in the Oracle Cloud Infrastructure Console. For security maintenance, you may edit the scheduled start time within the 21-day window.

- View or Edit Quarterly Maintenance Preferences
  To edit your Oracle Exadata Database Service on Cloud@Customer infrastructure quarterly maintenance preferences, be prepared to provide values for the infrastructure configuration. The changes you make will only apply to future maintenance runs, not those already scheduled.
- View or Edit Monthly Maintenance Preferences
  This task describes how to set monthly maintenance preferences. The changes you make will only apply to future maintenance runs, not those already scheduled.
- Manage Quarterly/Monthly Maintenance Plan using Scheduling Policy
  After the scheduling policy is selected, Oracle generates a recommended maintenance scheduling plan to apply updates to all the components in your infrastructure.
- View or Edit a Scheduled Quarterly Maintenance
  Learn how to view and edit the time of the next scheduled maintenance.
- View or Edit a Scheduled Security Maintenance
  Learn how to view and edit the next scheduled security maintenance.
- View the Maintenance History
  Learn how to view the maintenance history for an Oracle Exadata Database Service on Cloud@Customer Infrastructure.

View or Edit Quarterly Maintenance Preferences

To edit your Oracle Exadata Database Service on Cloud@Customer infrastructure quarterly maintenance preferences, be prepared to provide values for the infrastructure configuration. The changes you make will only apply to future maintenance runs, not those already scheduled.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click the **Quarterly maintenance scheduling plan** tab.
6. Click **Edit Maintenance Preferences**.
   Edit Maintenance Preferences page is displayed.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Changes made to maintenance preferences apply only to future maintenance, not the maintenance that has already been scheduled.

7. On the Edit Maintenance Preferences page, configure the following:
   - **Maintenance scheduling preference: Oracle managed schedule**
     An Oracle-managed schedule is a maintenance scheduling preference in which Oracle automatically determines and manages the timing of maintenance activities. Oracle selects the maintenance windows and performs required updates within those windows, without requiring customer-defined schedules.
     - **Choose a maintenance method:**
       - **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
       - (Applies to quarterly maintenance) **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
     - (Applies to quarterly maintenance) **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
       - **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

       > **Note**
       >
       > Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

       Default: 15 minutes

       Maximum: 120 minutes
     - Click **Save Changes**.

       > **Note**
       >
       > From the next maintenance run onwards, executions will occur according to Oracle's schedules.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Maintenance scheduling preference: Customer managed schedule**
  - **Maintenance schedule:** Define maintenance preferences for this infrastructure.

    > **Note**
    >
    > Changes will take effect from the next maintenance run.

    * **Configure maintenance preference:** Define maintenance time preferences for each quarter. If more than one preference is defined for a quarter, Oracle automation will select one of them to perform maintenance on all components in your infrastructure.
      Select at least one month every two quarters.
    * **Specify a schedule:** Choose your preferred week, weekday, start time, and lead time for infrastructure maintenance.
      * Optional. Under **Week of the month**, specify which week of the month, maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days of the month, and have a duration of 7 days. Weeks start and end based on calendar dates, not days of the week. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days. If you do not specify a day of the week, then Oracle will run the maintenance update on a weekend day to minimize disruption.
      * Optional. Under **Day of the week**, specify the day of the week on which the maintenance will occur. If you do not specify a day of the week, Oracle will run the maintenance update on a weekend day to minimize disruption.
      * Optional. Under **Hour of the day**, specify the hour during which the maintenance run will begin. If you do not specify a start hour, Oracle will pick the least disruptive time to run the maintenance update.
      * Under **Notification Lead Time**, specify the minimum number of weeks ahead of the maintenance event you would like to receive a notification message. Your lead time ensures that a newly released maintenance update is scheduled to account for your required minimum period of advanced notification.
    * **Choose a maintenance method:**
      * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
      * (Applies to quarterly maintenance) **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
    * (Applies to quarterly maintenance) **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.

Copyright © 2008, 2026, Oracle and/or its affiliates.

* **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

> **Note**
>
> Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 15 minutes

Maximum: 120 minutes

* **Show advanced options:**
  * Enable monthly security infrastructure maintenance: Select this check box to perform monthly security infrastructure maintenance.

– **Maintenance schedule:** Use maintenance window preferences from a scheduling policy
During infrastructure provisioning, after the scheduling policy is selected, Oracle generates a recommended maintenance scheduling plan to apply updates to all the components in your infrastructure. The recommended plan schedules all DB Servers, followed by Storage Servers and Network Switches, into the maintenance windows from your policy based on duration. After provisioning the infrastructure, you can update the scheduling plan by editing the 'Maintenance Scheduling Plan' resource and customize the update to specific components to align with different windows in your scheduling policy.

* Click **Select policy**.
* In the resulting Select maintenance scheduling policy window, choose a compartment and a policy.
  You can also create a maintenance scheduling policy and use it. For more information, see Create a Maintenance Scheduling Policy. Note that you can add additional maintenance windows to the policy after creating it. For more information, see Add Additional Maintenance Windows to a Maintenance Scheduling Policy.
* Click **Save**.

> **Note**
>
> Changes will take effect from the next maintenance run.

You must confirm your choice by entering the currently used policy name in a confirmation dialog before making any changes that delete the associated maintenance plan created with the attached policy.

* Changing from one scheduling policy to another scheduling policy after the recommended maintenance plan is created and saved for the infrastructure

Copyright © 2008, 2026, Oracle and/or its affiliates.

* Changing from using a scheduling policy to not using a policy and defining your maintenance preference in line with your infrastructure
* Change from using the scheduling policy to not using the policy and apply updates as per the Oracle-managed schedule.

All of the above changes delete the scheduling plan for your infrastructure created with the current policy, and you will lose any customizations made to the Oracle recommended plan if you attach the same policy later.

8. Click **Save**.
   If you switch from rolling to non-rolling maintenance method, then Confirm Non-rolling Maintenance Method dialog is displayed.
   a. Enter the name of the infrastructure in the field provided to confirm the changes.
   b. Click **Save**.

View or Edit Monthly Maintenance Preferences

This task describes how to set monthly maintenance preferences. The changes you make will only apply to future maintenance runs, not those already scheduled.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.
   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click the **Monthly maintenance scheduling plan** tab.
6. Click **Edit Maintenance Preferences**.
   Edit Maintenance Preferences page is displayed.

   NoteChanges made to maintenance preferences apply only to future maintenance, not the maintenance that has already been scheduled.
7. On the Edit Maintenance Preferences page, configure the following:
   - **Maintenance scheduling preference: Oracle managed schedule** An Oracle-managed schedule is a maintenance scheduling preference in which Oracle automatically determines and manages the timing of maintenance activities. Oracle selects the maintenance windows and performs required updates within those windows, without requiring customer-defined schedules.
   - **Maintenance scheduling preference: Customer managed schedule**
     – **Maintenance schedule:** Use maintenance window preferences from a scheduling policy
     During infrastructure provisioning, after the scheduling policy is selected, Oracle generates a recommended maintenance scheduling plan to apply updates to all the components in your infrastructure. The recommended plan schedules all DB Servers, followed by Storage Servers and Network Switches, into the maintenance windows from your policy based on duration. After provisioning the infrastructure, you can update the scheduling plan by editing the 'Maintenance Scheduling Plan' resource and customize the update to specific components to align with different windows in your scheduling policy.

Copyright © 2008, 2026, Oracle and/or its affiliates.

* Click **Select policy**.
* In the resulting Select maintenance scheduling policy window, choose a compartment and a policy.
  You can also create a maintenance scheduling policy and use it. For more information, see Create a Maintenance Scheduling Policy. Note that you can add additional maintenance windows to the policy after creating it. For more information, see Add Additional Maintenance Windows to a Maintenance Scheduling Policy.
* Click **Save**.

> **Note**
>
> Changes will take effect from the next maintenance run.

You must confirm your choice by entering the currently used policy name in a confirmation dialog before making any changes that delete the associated maintenance plan created with the attached policy.

* Changing from one scheduling policy to another scheduling policy after the recommended maintenance plan is created and saved for the infrastructure
* Change from using the scheduling policy to not using the policy and apply updates as per the Oracle-managed schedule.

All of the above changes delete the scheduling plan for your infrastructure created with the current policy, and you will lose any customizations made to the Oracle recommended plan if you attach the same policy later.

8. Click **Save**.

Manage Quarterly/Monthly Maintenance Plan using Scheduling Policy

After the scheduling policy is selected, Oracle generates a recommended maintenance scheduling plan to apply updates to all the components in your infrastructure.

The recommended plan schedules all DB Servers, followed by Storage Servers and Network Switches, into the maintenance windows from your policy based on duration. You can update the maintenance scheduling plan and customize the update to specific components to align with different windows in your scheduling policy.

- View Quarterly/Monthly Maintenance Scheduling Policy
  To view the quarterly maintenance scheduling policy for your infrastructure, use this procedure.
- Change the Quarterly/Monthly Maintenance Scheduling Policy
  To change the quarterly maintenance scheduling policy for your infrastructure, use this procedure.
- View Quarterly/Monthly Maintenance Scheduling Plan
  To view the maintenance scheduling plan for your infrastructure, use this procedure.
- Edit Quarterly/Monthly Maintenance Plan Scheduled Actions
  To edit scheduled actions of an Exadata Cloud@Customer Infrastructure maintenance scheduling plan, use this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

View Quarterly/Monthly Maintenance Scheduling Policy

To view the quarterly maintenance scheduling policy for your infrastructure, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view the quarterly maintenance scheduling policy.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure. In the Maintenance section, find the Quarterly Maintenance Schedule policy or Monthly Maintenance Schedule policy.
5. Click the link to view details of the policy.

Change the Quarterly/Monthly Maintenance Scheduling Policy

To change the quarterly maintenance scheduling policy for your infrastructure, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view the quarterly maintenance scheduling policy or Monthly Maintenance Schedule policy.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure. In the Maintenance section, find the Quarterly Maintenance Schedule policy.
5. Click the Actions menu and then select **Edit quarterly maintenance preferences** or **Edit monthly maintenance preferences**.

   The Customer Managed Schedule and Use maintenance window preferences from a scheduling policy are selected as part your policy selection.
6. Click **Select Policy**.
7. Select an existing policy or create a policy and select it.
8. Click **Save**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Changes will take effect from the next maintenance run.
>
> - To change from one scheduling policy to another scheduling policy after the recommended maintenance plan is created and saved for the infrastructure, you must confirm your choice by entering the currently used policy name in a confirmation dialog before.
> - Policy change deletes the scheduling plan created with the current policy for your infrastructure, and you will lose any customizations made to the Oracle recommended plan if you attach the same policy later.

View Quarterly/Monthly Maintenance Scheduling Plan

To view the maintenance scheduling plan for your infrastructure, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view maintenance scheduling plan.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click the **Quarterly maintenance scheduling plan** or **Monthly maintenance scheduling plan** tab.

   The Maintenance scheduling Plan section displays the maintenance scheduling plan associated with the infrastructure. The details include maintenance windows associated with the plan, lifecycle state, duration in hours, scheduled actions, and the estimated time.

> **Note**
>
> - If all infrastructure components are not scheduled for an update in the maintenance scheduling plan, you will see a banner in the console indicating the components missing from the plan, and the plan will reflect a 'Needs Attention' lifecycle state.
> - If the missing components are not included in the plan before Oracle automation schedules the next quarterly maintenance run, they will be automatically added to an 'Unplanned' maintenance window for that quarter's maintenance run.

Edit Quarterly/Monthly Maintenance Plan Scheduled Actions

To edit scheduled actions of an Exadata Cloud@Customer Infrastructure maintenance scheduling plan, use this procedure.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view maintenance scheduling plan.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click the **Quarterly maintenance scheduling plan** or **Monthly maintenance scheduling plan** tab.
6. In the Maintenance scheduling plan section, click the Actions menu (three dots) of the Maintenance window you want to edit action, and select **Edit**.
7. In the resulting Edit scheduled actions window, do the following:
   - **Add a scheduled action:**
     - **Create new action:** When you are creating a new maintenance action, you can choose to add components already scheduled to update in different maintenance windows to a new window.
       - **Select action type:**
         (Applies to quarterly maintenance) **DB Server Exadata full software update** or (Applies to monthly security maintenance) **DB Server Exadata online software update**
           - **DB Server Exadata full software update**
             - **Configure maintenance method:**
               - **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
               - (Applies to quarterly maintenance) **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
               - (Applies to quarterly maintenance) **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
                 - **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 15 minutes

Maximum: 120 minutes

- **Add DB Servers:**
  - **Select DB Servers:** Updates to selected DB Servers will be moved from their currently scheduled window to the window you are adding the maintenance action.
- **Storage Server Exadata full software update**
  - **Configure maintenance method:**
    - **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
    - (Applies to quarterly maintenance) **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.

> **Note**
>
> All storage servers in the infrastructure must be scheduled to update in a single maintenance action to apply non-rolling storage updates. While these updates are applied, your database workloads will incur complete downtime.

  - **Select storage server from:** Select a window from where you want to add storage servers from.
  - **Select maintenance action to add from:** Select the action from where you want to add storage servers from.
  - **Select number of storage servers to add:** Select the number of storage servers to add to this action.
- **Network switch software update:** A banner with the message "Network switch update is already scheduled for the selected maintenance window." if the network switch software update is already scheduled in the chosen maintenance window.

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - **Move action from another window:** When you are moving an action, you can choose to move all components scheduled to update in a specific window to a new window.
    * **Select the window to move action from:** Choose a window from the maintenance run.
    * **Select the action to move:** Choose a specific action from the maintenance window to move.

- **Remove a scheduled action:** Click the Actions menu (three dots) of the maintenance action, and then select **Remove**.

  > **Note**
  >
  > – Any action can be removed as long as there are no components scheduled for update in that action.
  >
  > – Any window can be removed as long as there are no actions scheduled for update in that window.

- **Edit a scheduled action:**
  - Click the Actions menu (three dots) of the maintenance action type **Db server full software update**, **Storage server full software update**, and **Network switch software update**, and then select **Edit scheduled action**.

View or Edit a Scheduled Quarterly Maintenance

Learn how to view and edit the time of the next scheduled maintenance.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click **Exadata Infrastructure**.
4. In the list of Exadata Infrastructures, find the infrastructure you want to set the next scheduled maintenance window for and click its highlighted name.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.

   > **Note**
   >
   > An information block is displayed 6 hours before the start of a maintenance run, regardless of whether you've chosen rolling or non-rolling maintenance method. When the maintenance begins, it is automatically removed.

5. On the Infrastructure Details page, click **Maintenance**, click the **View** link in the **Next Quarterly Maintenance** field.

   The Exadata Infrastructure Maintenance page is displayed.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you've opted to use a Scheduling Policy, link to the policy will appear. Click the link to view the Maintenance scheduling policy details. To edit a maintenance window, refer to the topics listed under *Manage Quarterly Maintenance Run created from Scheduling Plan*.

6. On the **Exadata Infrastructure Maintenance** page, scheduled maintenance details are listed.

   Target DB Server Version and **Target Storage Server Version:** These fields display the Exadata software version to be applied by the scheduled maintenance. The version applied will be the most recent certified update for Exadata infrastructures in the cloud. If the next quarterly update is not yet certified when the maintenance is scheduled, then the versions may show "LATEST" until the new quarterly update becomes available. Once the update becomes available the new version will be displayed.

   To find information on the Database Server Exadata software version or the Storage Server Exadata software version, see My Oracle Support note *Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)*.

7. To change the next scheduled maintenance settings, click **Edit Maintenance Run**.

   On the **Edit Maintenance** page, do the following:

   - Select a maintenance method, **Rolling** or **Non-rolling**.

     > **Note**
     >
     > If you select the **Non-rolling** option, an information block appears stating that components will be updated simultaneously, resulting in full system downtime.

   - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
     - **Custom action timeout (in minutes):** Maximum timeout available to perform custom action before starting maintenance on the DB Servers.
       Default: 30 minutes

       Maximum: 120 minutes
   - To reschedule the next quarterly maintenance run, enter a date and time in the **Scheduled Start time** field.
     The following restrictions apply:
     - Oracle expects to be able to perform infrastructure maintenance at least once per quarter. You should not defer maintenance beyond the end of a maintenance quarter unless unexpected issues prevent your accommodating it before the next maintenance quarter.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- In the event unexpected issues prevent your accommodating the scheduled infrastructure maintenance run, you can reschedule the infrastructure maintenance to another date no more than 180 days from the prior infrastructure maintenance. Since normal maintenance should be performed quarterly, this provides approximately 90 additional days for you to reschedule the infrastructure maintenance. Oracle strongly recommends you not schedule maintenance at or close to the 180 day limit, as you will have no flexibility to reschedule further if additional unexpected issues arise.
- If a new maintenance release is announced prior to your rescheduled maintenance run, the newer release will be applied on your specified date.
- You can reschedule your maintenance to take place earlier than it is currently scheduled.
- Oracle reserves certain dates each quarter for internal maintenance operations, and you cannot schedule your maintenance on these dates.

- Click **Save Changes**.

8. To view estimated maintenance time details for various components, click the **View** link is displayed in the **Total Estimated Maintenance Time** field.

   The **View** link is displayed in the **Total Estimated Maintenance Time** field only if the Maintenance Method is **Rolling**.

   The **Estimated Maintenance Time Details** page is displayed with details that include:

   - Total Estimated Maintenance Time
   - Database Servers Estimated Maintenance Time
   - Storage Servers Estimated Maintenance Time
   - Network Switches Estimated Maintenance Time
   - Order in which components are updated. In rolling maintenance, components are updated in the sequence displayed

   a. To view the number of VMs that will be restarted as part of Database Server maintenance, click the **Show details** link.
   The **VM Location** dialog is displayed.

   b. In the **VM Cluster Name** field, you can find out what VM cluster a particular VM belongs to.

   c. Click **Close**.

9. Click **Close** to close the **Estimated Maintenance Time Details** page.

- View and Edit Quarterly/Monthly Maintenance While Maintenance is In Progress
  While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for a custom action, you can resume the maintenance prior to the timeout or extend the timeout.
- (Applies to quarterly maintenance) View and Edit Maintenance While Maintenance is Waiting for Custom Action
  While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for custom action, you can resume the maintenance prior to the timeout or extend the timeout.

Related Topics

- https://support.oracle.com/epmos/faces/DocContentDisplay?id=888828.1

Copyright © 2008, 2026, Oracle and/or its affiliates.

View and Edit Quarterly/Monthly Maintenance While Maintenance is In Progress

While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for a custom action, you can resume the maintenance prior to the timeout or extend the timeout.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.

   The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.
5. Click **Maintenance**.
6. Click **Edit Maintenance Run**.

   Edit Maintenance page is displayed.

   > **Note**
   >
   > You can only make edits to the custom action configuration, not the maintenance method or scheduled start time. Enabling or disabling the custom action or modifying the custom action timeout while maintenance is in progress will apply to all database servers that have yet to be updated.

7. (Applies to quarterly maintenance) On the **Edit Maintenance** page, do the following:
   - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
     - **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.
       Default: 30 minutes

       Maximum: 120 minutes
8. Click **Save**.

   If you have configured the rolling maintenance method, then the **View** link is displayed in the **Total Estimated Maintenance Time** field.

   a. Click **View**.
   Estimated Maintenance Time Details page is displayed with details that include:
      - Total Estimated Maintenance Time
      - Database Servers Estimated Maintenance Time

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Storage Servers Estimated Maintenance Time
- Network Switches Estimated Maintenance Time
- Order in which components are updated. In rolling maintenance, components are updated in the sequence displayed.

b. Click **Close**.

(Applies to quarterly maintenance) View and Edit Maintenance While Maintenance is Waiting for Custom Action

While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for custom action, you can resume the maintenance prior to the timeout or extend the timeout.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud at Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.

   The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.
5. Click **Maintenance**.

> **Note**
>
> - Editing a maintenance run is not available while waiting for custom action.
> - While maintenance is waiting for custom action, an information block is displayed. The information block is removed after the maintenance resumes.

6. On the information block, do one of the following:
   a. Click **Resume Maintenance Now** to resume the maintenance, proceeding to the next database server.
   Resume Maintenance dialog is displayed. Click **Resume Maintenance Now**.
   b. Click **Extend Custom Action Timeout**.
   You can extend timeout multiple times within the maximum allowable time of 2 hours. If you try extending beyond the maximum limit, then the system displays the Cannot Extend Custom Action Timeout dialog indicating that the custom action timeout has already been extended to the maximum allowable 2 hours and you cannot extend it further.

View or Edit a Scheduled Security Maintenance

Learn how to view and edit the next scheduled security maintenance.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.

Copyright © 2008, 2026, Oracle and/or its affiliates.

3. Choose your **Compartment**.
4. Click the name of the Exadata infrastructure that you want to view maintenance details.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.

   Quarterly and monthly maintenance details are displayed under the **Maintenance** section. Monthly security maintenance is labeled as **Next Security Maintenance**.
5. On the **Infrastructure Details** page, click **Maintenance**.

   The Exadata Infrastructure Maintenance page is displayed. The Exadata Infrastructure Maintenance page includes details such as **Type: Monthly Security Maintenance**, **Scheduled Start Time**, and so on.
6. To reschedule monthly security maintenance, click **Edit Maintenance Run** and pick a new date within the 21-day cycle.

   > **Note**
   >
   > Certain black-out dates are not available for security maintenance and are grayed out in the rescheduling calendar.

7. Select a date and then click **Save Changes**.
8. To view the maintenance history, click **Maintenance history**.

   The Maintenance History page displays details including the type of maintenance, **Monthly** or **Quarterly**.

   When a monthly security maintenance is in progress, the Infrastructure resource state will be **Available**.

View the Maintenance History

Learn how to view the maintenance history for an Oracle Exadata Database Service on Cloud@Customer Infrastructure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click **Exadata Infrastructure**.
4. In the list of Exadata Infrastructures, find the infrastructure you want to view the maintenance history and click its highlighted name.

   The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. On the Infrastructure Details page, click **Maintenance history** to see a list of past maintenance events including details on their completion state and the target database and storage server versions.

   To find information on the Database Server Exadata software version or the Storage Server Exadata software version, see My Oracle Support note *Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1).*

Related Topics

- https://support.oracle.com/epmos/faces/DocContentDisplay?id=888828.1

Copyright © 2008, 2026, Oracle and/or its affiliates.

Manage Quarterly/Monthly Maintenance Run created from Scheduling Plan

- View Maintenance Windows Associated with a Maintenance Run
- Edit Maintenance Window Associated with a Maintenance Run
- View Maintenance Actions Associated with a Maintenance Run
- Edit Maintenance Actions of a Maintenance Window Associated with a Maintenance Run
- View and Edit Maintenance While Maintenance is In Progress
- (Applies to quarterly maintenance) View and Edit Maintenance While Maintenance is Waiting for Custom Action
- Cancel a Maintenance Run While In Progress
- Start a Maintenance On Demand
  To start a maintenance on demand, follow these steps.
- View Maintenance Activity in a Compartment
- View the Details of a Maintenance Activity
- View Maintenance History in a Compartment
- View Maintenance History Details in a Compartment
- Review and Respond to Unplanned Maintenance Activity

View Maintenance Windows Associated with a Maintenance Run

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.
   The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
   The resulting Activity page lists maintenance updates scheduled to run in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance windows.
   The **Maintenance windows** tab in the resulting Maintenance run page lists the maintenance windows associated with the chosen activity.
   - **Start time (UTC):** Window start time in UTC For example, Sun, Jun 23, 2024, 18:30:58 UTC.
     - The following restrictions apply to monthly security maintenance:
       * Monthly security maintenance is performed on the database servers online, with no reboot. Monthly updates are applied to storage servers in a rolling manner.
       * Monthly security maintenance can also be scheduled at a specific time during the month, albeit in a single maintenance window. Oracle will publish a schedule for monthly maintenance at least one week prior to start of the maintenance period, and you can reschedule if required.
       * CPU scaling and VM startup/shutdown are the only operations supported during monthly infrastructure maintenance. Please plan to defer other

Copyright © 2008, 2026, Oracle and/or its affiliates.

operations until after the updates are complete. If you attempt an affected operation, the console will notify you of the ongoing security updates. No software is updated in the guest VMs.

    - Security maintenance, when required, can be scheduled by customers at any time within the calendar month. Each maintenance window must be defined within a continuous 21-day period; however, customers can select the start and end dates within the month (for example, from the 1st to the 21st, or from the 10th to the 30th or 31st).
      Customers receive notification of the proposed schedule at least 7 days before the start of the selected maintenance window and can reschedule within the permitted range as needed.

  - The following restrictions apply to quarterly maintenance:
    - Oracle expects to be able to perform infrastructure maintenance at least once per quarter. You should not defer maintenance beyond the end of a maintenance quarter unless unexpected issues prevent your accommodating it before the next maintenance quarter.
    - In the event unexpected issues prevent your accommodating the scheduled infrastructure maintenance run, you can reschedule the infrastructure maintenance to another date no more than 180 days from the prior infrastructure maintenance. Since normal maintenance should be performed quarterly, this provides approximately 90 additional days for you to reschedule the infrastructure maintenance. Oracle strongly recommends you not schedule maintenance at or close to the 180 day limit, as you will have no flexibility to reschedule further if additional unexpected issues arise.
    - If a new maintenance release is announced prior to your rescheduled maintenance run, the newer release will be applied on your specified date.
    - You can reschedule your maintenance to take place earlier than it is currently scheduled.
    - Oracle reserves certain dates each quarter for internal maintenance operations, and you cannot schedule your maintenance on these dates.

- **Type:** Planned vs Unplanned. All windows created from the infrastructure maintenance scheduling plan or added by you to this maintenance run are 'Planned' windows. All other windows Oracle automation creates to address failures, duration enforcement, or unforeseen events are defined as 'Unplanned' windows. Always review activities scheduled to run in an 'Unplanned' window.
- **Maintenance actions:** The summary of actions scheduled to update in a given window.
  The server name identifies updates scheduled for DB servers. For example, Apply full update to DB servers dbServer-1 and dbServer-2. The storage server updates are identified as count since all storage servers have identical storage layouts. For example, Apply full update to 2 Storage Servers. The Network switches are updated as a pair and cannot be scheduled to update in different actions or maintenance windows. For example, Apply full update to 2 Network switches.
- **Estimated time:** The estimated time for Oracle automation to complete maintenance actions scheduled to apply updates to all infrastructure components across all windows in the maintenance run.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Edit Maintenance Window Associated with a Maintenance Run

> **Note**
>
> You can update the window configuration, like window schedule start time, duration, and duration enforcement, while the window is still in the 'Scheduled' life cycle state. Once the window is in progress, you cannot make changes to the configuration. You can choose to cancel a running maintenance for a window. Details covered in the *Cancel Maintenance Window Associated with a Maintenance Run* section.'

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.
   The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
   The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance windows.
   The Maintenance windows tab in the resulting Maintenance run page lists the maintenance windows associated with the chosen activity.
6. Click the Actions menu (three dots) of the Maintenance window you want to edit, and then select **Edit maintenance window**.
7. In the resulting Edit maintenance window page, update the **Maintenance window start time**, **Scheduled start time (UTC)**, **Duration in hours**, and **Enforce window duration** fields.
8. Click **Save changes**.

View Maintenance Actions Associated with a Maintenance Run

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.
   The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
   The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance actions.
6. Click the **Maintenance actions** tab.
7. To add actions:
   a. Click **Add action**.
   b. In the resulting Add maintenance actions window, you can **Create new maintenance action** or **Use existing maintenance action**.
8. Click **Add**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

9. To delete an action:
    - Click the Actions menu (three dots) of the maintenance action, and then select **Remove**.

Edit Maintenance Actions of a Maintenance Window Associated with a Maintenance Run

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.
   The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
   The resulting Activity page lists maintenance activities in the chosen compartment.
    a. Click the name of the activity you want to view associated maintenance windows.
       The Maintenance Windows tab in the resulting Maintenance run page lists the maintenance windows associated with the chosen activity.
    b. Click the Actions menu (three dots) of the Maintenance window you want to edit, and then select **Edit maintenance actions.**The resulting Edit maintenance actions page displays the list of actions. You can either add more actions or delete the existing ones.
    c. To add actions:
        i. Click **Add action**.
        ii. Do the following in the resulting Add maintenance action window:
            - **Create new maintenance action**: When you are creating a new maintenance action, you can choose to add components already scheduled to update in different maintenance windows to a new window.
                - **Select action type**:
                    * (Applies to quarterly maintenance) **DB Server Exadata full software update** or (Applies to monthly security maintenance) **DB Server Exadata online software update**
                        * **Maintenance method:**
                            * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
                            * (Applies to quarterly maintenance) **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
                            * (Applies to quarterly maintenance) **Custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB

Copyright © 2008, 2026, Oracle and/or its affiliates.

servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.

- **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

> **Note**
>
> Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.
> **Default:** 15 minutes
>
> **Maximum:** 120 minutes

  - **Add DB Servers:**
    - **Select DB Servers:** Updates to selected DB Servers will be moved from their currently scheduled window to the window you are adding the maintenance action.

- **Storage Server Exadata full software update**
  - **Maintenance method:**
    - **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
    - (Applies to quarterly maintenance) **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.

> **Note**
>
> All storage servers in the infrastructure must be scheduled to update in a single maintenance action to apply non-rolling storage updates. While these updates are applied, your database workloads will incur complete downtime.

  - **Maintenance window:** Select a window from where you want to add storage servers from.
  - **Maintenance actions:** Select the action from where you want to add storage servers from.
  - **Storage servers:** Select the number of storage servers to add to this action.

Copyright © 2008, 2026, Oracle and/or its affiliates.

d. To delete an action:

i. Click the Actions menu (three dots) of the maintenance action, and then select **Remove**.

> **Note**
>
> - Any action can be removed as long as there are no components scheduled for update in that action.
> - Any window can be removed as long as there are no actions scheduled for update in that window.

View and Edit Maintenance While Maintenance is In Progress

While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for a custom action, you can resume the maintenance prior to the timeout or extend the timeout.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.
   The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.

   Note

   (Applies to quarterly maintenance) **Maintenance In Progress** status is displayed in the **Next Quarterly Maintenance** field.

   (Applies to monthly maintenance) **Maintenance In Progress** status is displayed in the **Next Security Maintenance** field.
5. Click the **View** link in the **Next Quarterly Maintenance** or **Next Security Maintenance** field.

   You will be on the maintenance run details page that is in progress.
6. Click the **Maintenance windows** tab on the Maintenance Run details page.
7. Identify the maintenance window that's in progress
8. Click the Actions menu (three dots) and select **Edit maintenance actions**.
9. (Applies to quarterly maintenance) Click the Actions menu (three dots) and select **Edit custom action configuration**.
10. (Applies to quarterly maintenance) In the resulting **Edit custom action configuration** page, enter **Custom action in minutes**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> While maintenance is in progress you can only change the custom action time for DB Server action type. You cannot change the custom action time for support this option for any other action type.

11. Click **Update**.

(Applies to quarterly maintenance) View and Edit Maintenance While Maintenance is Waiting for Custom Action

While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for a custom action, you can resume the maintenance prior to the timeout or extend the timeout.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.
   The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.

> **Note**
>
> **Maintenance In Progress** status is displayed in the **Next Quarterly Maintenance** field.

5. Click the **View** link in the **Next Quarterly Maintenance** field.
   You will be on the maintenance run details page that is in progress.
6. In the resulting **Maintenance Run** details page, click **Maintenance actions**.
   While maintenance is waiting for custom action, an information block is displayed. And, you cannot edit the maintenance while waiting for customer action. The information block is removed after the maintenance resumes.
7. On the information block, do the following:
   a. Click **Resume Maintenance Now** to resume the maintenance, proceeding to the next database server.
      Resume Maintenance dialog is displayed. Click **Resume Maintenance Now**.
   b. Click **Extend Custom Action Timeout**.
      You can extend timeout multiple times within the maximum allowable time of 2 hours. If you try extending beyond the maximum limit, then the system displays the Cannot Extend Custom Action Timeout dialog indicating that the custom action timeout has already been extended to the maximum allowable 2 hours and you cannot extend it further.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Cancel a Maintenance Run While In Progress

To cancel a Maintenance Window Associated with a Maintenance Run, follow these steps:

> **Note**
>
> You can cancel a running maintenance while the scheduled updates for a window are in progress. Canceling the maintenance while the updates are in progress allows you to reschedule all actions that have not yet started to a future maintenance window of your choice. You can choose a new start time and duration to finish all the actions rescheduled from the maintenance window you decided to cancel.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.
   The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
   The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance windows.
   The Maintenance Windows section in the resulting Maintenance run page lists the maintenance windows associated with the chosen activity.
6. Click the Actions menu (three dots) of the Maintenance window you want to cancel, and then select **Cancel maintenance window.**

To cancel a Maintenance Run While In Progress, follow these steps:

1. Open the navigation menu. Under **Oracle Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.
   The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
   The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the '**In Progress**' activity you want to cancel.
6. Click **Maintenance windows** on the Maintenance Run details page.
7. Identify the maintenance window that's in progress.
8. Click the Actions menu (three dots) and select **Cancel maintenance window**.
9. On the resulting Cancel maintenance run window, reconfigure **Maintenance window start time**.
10. Select the **Enforce window duration** check box to pause and re-schedule any scheduled action that goes over the configured window duration to resume in a future maintenance window.

Copyright © 2008, 2026, Oracle and/or its affiliates.

11. Click **Reschedule maintenance run**.

> **Note**
>
> The maintenance run will complete the current operation. All remaining actions scheduled for this window will be rescheduled to a new maintenance window.

Start a Maintenance On Demand

To start a maintenance on demand, follow these steps.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.
   The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
   The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance windows.
   The Maintenance Windows tab in the resulting Maintenance run page lists the maintenance windows associated with the chosen activity.
6. Click the Actions menu (three dots) of the Maintenance window, and then select **Start maintenance now.**

View Maintenance Activity in a Compartment

> **Note**
>
> Maintenance activity lists all the maintenance updates scheduled to run for all infrastructure resources in a given compartment for the selected Exadata cloud service.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.
   The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
   The resulting Activity page lists maintenance activities in the chosen compartment.

View the Details of a Maintenance Activity

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Scheduling policy**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The resulting Scheduling Policy page displays the list of policies.

3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
   The resulting Activity page lists maintenance updates scheduled to run in the chosen compartment.
5. Click the Actions menu (three dots) of the maintenance activity you want to view details.
   The resulting Maintenance run page displays the details of the chosen maintenance activity.

View Maintenance History in a Compartment

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose a compartment from the **Compartment** filter.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view maintenance details.
5. Click the **Maintenance history** tab.
   The resulting page displays the history of maintenance runs in the chosen compartment.

View Maintenance History Details in a Compartment

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose a compartment from the **Compartment** filter.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view maintenance details.
5. Click the **Maintenance history** tab.
   The resulting History page lists maintenance run status (Succeeded or Failed) and other details of the maintenance activities in the chosen compartment.

   In case of a failure, automation will mark the run as 'Failed' and automatically reschedule a new window with the remaining components that needs update to the target version. The same update order is carried over.
6. Click the Actions menu (three dots) of the maintenance activity you want to view details.
   The resulting Maintenance History page displays the details of the chosen maintenance activity.

Review and Respond to Unplanned Maintenance Activity

Unplanned maintenance activity when infrastructure is scaled after maintenance is planned

- After scaling your infrastructure by adding DB or storage servers, you may need to update your maintenance scheduling plan to include these new components. If any infrastructure component is missing from the maintenance plan, a warning will appear in the infrastructure maintenance plan details section.
- When Oracle automation creates the maintenance run for the quarter, any components not included in the maintenance plan will be automatically added to an 'Unplanned'

Copyright © 2008, 2026, Oracle and/or its affiliates.

maintenance window. This ensures that all components have the correct system software applied each quarter, maintaining OCI software compliance.

- You can edit the scheduled start time of the 'Unplanned' window or move the updates for the missing components to an existing planned window as needed.

Unplanned maintenance activity when a scheduled update fails to apply

- If a scheduled update fails, the Oracle operations team will engage, evaluate the failure, and reschedule the failed update along with any unfinished updates to a future maintenance window. Oracle automation will mark this rescheduled window as 'Unplanned' and notify you to review the rescheduled maintenance activity.
- You can edit the 'Unplanned' window's scheduled start time or move the failed and unfinished updates to an existing planned window as needed.

Unplanned maintenance when schedule activity exceeds enforced window duration

- For maintenance windows configured with duration enforcement, Oracle automation will check if the estimated time to execute and apply the scheduled update is sufficient within the remaining window duration. If not, Oracle automation will automatically reschedule all unfinished updates to a future 'Unplanned' window, mark the current window as 'Duration Exceeded,' and notify you to review the rescheduled maintenance activity.
- Any updates already in progress will continue past the enforced window duration to ensure a consistent state of the underlying resources.

Monitor Infrastructure Maintenance Using Lifecycle State Information

The lifecycle state of your Exadata Infrastructure resource enables you to monitor when the maintenance of your infrastructure resource begins and ends.

In the Oracle Cloud Infrastructure Console, you can see lifecycle state details messages on the **Exadata Infrastructure Details** page when a tooltip is displayed beside the **Status** field. You can also access these messages using the `ListExadataInfrastructures` API, and using tools based on the API, including *SDKs* and the *OCI CLI*.

During infrastructure maintenance operations, you can expect the following:

- If you specify a maintenance window, then patching begins at your specified start time. The infrastructure resource's lifecycle state changes from **Available** to **Maintenance in Progress**.

> **Note**
>
> The prechecks are now done prior to the start of the maintenance.

- When Exadata database server maintenance starts, the infrastructure resource's lifecycle state is **Maintenance in Progress**, and the associated lifecycle state message is, **The underlying infrastructure of this system (dbnodes) is being updated.**
- When storage server maintenance starts, the infrastructure resource's lifecycle state is **Maintenance in Progress**, and the associated lifecycle state message is, **The underlying infrastructure of this system (cell storage) is being updated and this will not impact Database availability.**
- After storage server maintenance is complete, the networking switches are updated one at a time, in a rolling fashion.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- When maintenance is complete, the infrastructure resource's lifecycle state is **Available**, and the Console and API-based tools do not provide a lifecycle state message.

Related Topics

- ListExadataInfrastructures
- Software Development Kits and Command Line Interface
- Command Line Interface (CLI)

Receive Notifications about Your Infrastructure Maintenance Updates

There are two ways to receive notifications. One is through email to infrastructure maintenance contacts and the other one is to subscribe to the maintenance events and get notified.

Oracle schedules maintenance run of your infrastructure based on your scheduling preferences and sends email notifications to all your infrastructure maintenance contacts. You can login to the console and view details of the schedule maintenance run. Appropriate maintenance related events will be generated as Oracle prepares for your scheduled maintenance run, for example, precheck, patching started, patching end, and so on. For more information about all maintenance related events, see *Oracle Exadata Cloud@Customer Events*. In case, if there are any failures, then Oracle reschedules your maintenance run, generates related notification, and notifies your infrastructure maintenance contacts.

For more information about Oracle Cloud Infrastructure Events, see *Overview of Events*. To receive additional notifications other than the ones sent to infrastructure maintenance contacts, you can subscribe to infrastructure maintenance events and get notified using the Oracle Notification service, see *Notifications Overview*.

Related Topics

- Oracle Exadata Database Service on Cloud@Customer Events
  Oracle Exadata Database Service on Cloud@Customer resources emit events, which are structured messages that indicate changes in resources.
- Overview of Events
- Notifications Overview
- Defining Maintenance Contacts for Exadata Cloud@Customer
  Oracle Cloud Operations will use these contacts to communicate with customers about infrastructure-related events, including hardware failures and replacements, quarterly maintenance patches, and other operational updates or incidents

Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Infrastructure Maintenance Controls

Oracle Exadata Database Service on Cloud@Customer uses the same API as Oracle Cloud Infrastructure to manage infrastructure maintenance controls.

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 infrastructure maintenance controls:

- CreateExadataInfrastructure
- GetExadataInfrastructure

Copyright © 2008, 2026, Oracle and/or its affiliates.

- ListExadataInfrastructures
- UpdateExadataInfrastructure
- UpdateMaintenanceRun
- GetMaintenanceRun
- ListMaintenanceRuns
- ListSchedulingPolicies
- CreateSchedulingPolicy
- GetSchedulingPolicy
- UpdateSchedulingPolicy
- DeleteSchedulingPolicy
- ChangeSchedulingPolicyCompartment
- ListRecommendedScheduledActions
- ListSchedulingWindows
- CreateSchedulingWindow
- GetSchedulingWindow
- UpdateSchedulingWindow
- DeleteSchedulingWindow
- ListSchedulingPlans
- CreateSchedulingPlan
- GetSchedulingPlan
- DeleteSchedulingPlan
- ChangeSchedulingPlanCompartment
- ReorderScheduledActions
- CascadingDeleteSchedulingPlan
- ListScheduledActions
- CreateScheduledAction
- GetScheduledAction
- UpdateScheduledAction
- DeleteScheduledAction
- ListParamsForActionType
- ReorderScheduledActions
- ListExecutionWindows
- CreateExecutionWindow
- GetExecutionWindow
- UpdateExecutionWindow
- DeleteExecutionWindow
- ReorderExecutionActions
- CancelExecutionWindow

Copyright © 2008, 2026, Oracle and/or its affiliates.

- ListExecutionActions
- CreateExecutionAction
- GetExecutionAction
- UpdateExecutionAction
- DeleteExecutionAction
- MoveExecutionActionMember

Manage VM Cluster Networks

Learn how to use the console to create, edit, download a configuration file, validate, and terminate your infrastructure network.

- About Managing VM Cluster Networks on Oracle Exadata Database Service on Cloud@Customer
  The VM cluster provides a link between your Oracle Exadata Database Service on Cloud@Customer infrastructure and Oracle Databases you deploy.
- Using the Console to Create a VM Cluster Network
  To create your VM cluster network with the Console, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to View VM Cluster Network Details
  To view VM Cluster Network details, use this procedure.
- Using the Console to View SCAN Listener Port Configured
  You can only edit a VM cluster network that is not associated with a VM cluster.
- Using the Console to Edit a VM Cluster Network
  To edit a VM Cluster Network, use this procedure. You can only edit a VM cluster network that is not associated with a VM cluster.
- Using the Console to Download a File Containing the VM Cluster Network Configuration Details
  To provide VM cluster network information to your network administrator, you can download and supply a file containing the network configuration.
- Using the Console to Validate a VM Cluster Network
  You can only validate a VM cluster network if its current state is **Requires Validation**, and if the underlying Exadata infrastructure is activated.
- Using the Console to Download Network Validation Report
  Learn to validate and inspect the network validation failure report without active involvement from Oracle Cloud Ops in troubleshooting networking configuration issues.
- Using the Console to Terminate a VM Cluster Network
  Before you can terminate a VM cluster network, you must first terminate the associated VM cluster, if one exists, and all the databases it contains.
- Using the Console to Add Another DB Server to the VM Cluster Network
  To a DB server to the VM Cluster Network, complete this procedure.
- Using the Console to Remove DB Servers from a VM Cluster Network
  To remove a database server, use this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

About Managing VM Cluster Networks on Oracle Exadata Database Service on Cloud@Customer

The VM cluster provides a link between your Oracle Exadata Database Service on Cloud@Customer infrastructure and Oracle Databases you deploy.

Before you can create any databases on your Exadata Cloud@Customer infrastructure, you must create a VM cluster network, and you must associate it with a VM cluster.

The VM cluster network specifies network resources, such as IP addresses and host names, that reside in your corporate data center and are allocated to Exadata Cloud@Customer. The VM cluster network includes definitions for the Exadata client network and the Exadata backup network. The client network and backup network contain the network interfaces that you use to connect to the VM cluster virtual machines, and ultimately the databases that reside on those virtual machines.

> **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.

Using the Console to Create a VM Cluster Network

To create your VM cluster network with the Console, be prepared to provide values for the fields required for configuring the infrastructure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure for which you want to create a VM cluster network.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure for which you want to create a VM cluster network.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click **Create VM Cluster Network**.
6. Provide the requested information on the Data Center Network Details page:
   a. Provide the display name.

      The display name is a user-friendly name that you can use to identify the VM cluster network. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the VM cluster network.

   b. Assign IPs to DB servers.

      By default, all DB servers are assigned IP addresses to enable easy addition and removal of VMs to the cluster in the future. Note that while DB Servers can be added to or removed from the VM Cluster Network, addresses cannot be changed in the future. It is advisable to plan ahead and reserve enough addresses with your network

Copyright © 2008, 2026, Oracle and/or its affiliates.

team to accommodate any and all future growth of your VM Cluster to avoid a database migration scenario in the future. If needed, then edit the list of DB servers from assigning IP addresses. If needed, then edit the list of DB servers from assigning IP addresses.

i. Click **Edit**.
Edit DB Server Selection dialog is displayed.

ii. Clear the checkboxes to omit one or more DB servers.

iii. Click **Save**.
Assign IPs to DB servers section displays the list of selected DB servers.

c. Provide client network details.

The client network is the primary channel for application connectivity to Exadata Database Service on Cloud@Customer resources. The following settings define the required network parameters:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the client network between `1` and `4094`, inclusive. To specify no VLAN tagging, enter "`1`". (This is equivalent to a "`NULL`" VLAN ID tag value.)

> **Note**
>
> The values "`0`" and "`4095`" are reserved and cannot be entered.

- **Database servers starting IP address:** Using CIDR notation, provide the IP address range for the client network.
- **Netmask:** Specify the IP netmask for the client network.
- **Gateway:** Specify the IP address of the client network gateway.
- **Hostname Prefix:** Specify the prefix that is used to generate the hostnames in the client network.
- **Domain Name:** Specify the domain name for the client network.

d. Provide backup network details.

The backup network is the secondary channel for connectivity to Exadata Database Service on Cloud@Customer resources. It is typically used to segregate application connections on the client network from other network traffic. The following settings define the required network parameters:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the backup network between `1` and `4094`, inclusive. To specify no VLAN tagging, enter "`1`". (This is equivalent to a "`NULL`" VLAN ID tag value.)

> **Note**
>
> The values "`0`" and "`4095`" are reserved, and cannot be entered.

- **Database servers starting IP address:** Using CIDR notation, provide the IP address range for the backup network.
- **Netmask:** Specify the IP netmask for the backup network.
- **Gateway:** Specify the IP address of the backup network gateway.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Hostname Prefix:** Specify the prefix that is used to generate the hostnames in the backup network.
- **Domain Name:** Specify the domain name for the backup network.

e. (**For authorized customers only**) Optionally, provide disaster recovery network details. The disaster recovery network is typically used to segregate Data Guard network traffic. The following settings define the required network parameters:

- Select the **Enable disaster recovery network** checkbox to configure the third network.
- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the disaster recovery between `1` and `4094`, inclusive. To specify no VLAN tagging, enter "`1`". (This is equivalent to a "`NULL`" VLAN ID tag value.)

> **Note**
>
> The values "0" and "4095" are reserved, and cannot be entered.

- **Database servers starting IP address:** Using CIDR notation, provide the IP address range for the backup network.
- **Netmask:** Specify the IP netmask for the backup network.
- **Gateway:** Specify the IP address of the backup network gateway.
- **Hostname Prefix:** Specify the prefix that is used to generate the hostnames in the backup network.
- **Domain Name:** Specify the domain name for the backup network.

f. Provide DNS and NTP server details.

The VM cluster network requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. The following settings specify the servers that provide these services:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the client network. You may specify up to three DNS servers.
- **NTP Servers:** Provide the IP address of an NTP server that is accessible using the client network. You may specify up to three NTP servers.

g. Configure Advanced Options.

Network:

- (Optional) Assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. Default is 1521.
- (Optional) Assign a SCAN listener port (TCPS) in the range between 1024 and 8999. Default is 2484.

Tags: (Optional) You can choose to apply tags. If you have permission to create a resource, then you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, refer to information about 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.

7. Click **Review Configuration**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The Review Configuration page displays detailed information about the VM cluster network, including the hostname and IP address allocations. These allocations are initially system-generated and are based on your inputs.

8. (Optional) You can choose to adjust the system-generated network definitions on the Review Configuration page.
   a. Click **Edit**.
   b. Use the **Edit** dialog to adjust the system-generated network definitions to meet your requirements.
   c. Click **Save Changes**.
9. Click **Create VM Cluster Network**.

   The VM Cluster Network Details page is now displayed. Initially, after creation, the state of the VM cluster network is **Requires Validation**.

Maximum and recommended CIDR block prefix lengths for each Exadata system shape
The following table specifies the maximum and recommended CIDR block prefix lengths for each Exadata system shape. The maximum CIDR block prefix length defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth.

Max # Database Servers	Client and DR Subnet Size	Backup Subnet Size
5	/28 (16 IPs)	/29 (8 IPs)
13	/27 (32 IPs)	/28 (16 IPs)
29	/26 (64 IPs)	/27 (32 IPs)
32	/25 (128 IPs)	/26 (64 IPs)

Related Topics

- Resource Tags

Using the Console to View VM Cluster Network Details

To view VM Cluster Network details, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure for which you want to create a VM cluster network.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network that you want to view details.

   The VM Cluster Network Details page displays information about the selected VM cluster network.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to View SCAN Listener Port Configured

You can only edit a VM cluster network that is not associated with a VM cluster.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to edit.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure for which you want to view VM cluster network details.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Under **VM Cluster Networks**, click the name of the VM Cluster Network for which you want to view details.

   VM Cluster Network Details page displays SCAN Listener Port under VM Cluster Network Information.

Using the Console to Edit a VM Cluster Network

To edit a VM Cluster Network, use this procedure. You can only edit a VM cluster network that is not associated with a VM cluster.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to edit.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network that you want to edit.

   The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Edit VM Cluster Network**.
7. Use the **Edit** dialog to edit the VM cluster network attributes:
   a. **Client Network**

      The client network is the primary channel for application connectivity to Exadata Database Service on Cloud@Customer resources. You can edit the following client network settings:

      - **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the client network between `1` and `4094`, inclusive. To specify no VLAN tagging, enter "`1`". (This is equivalent to a "`NULL`" VLAN ID tag value.)

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> The values "0" and "4095" are reserved and cannot be entered.

- **Netmask:** Specify the IP netmask for the client network.
- **Gateway:** Specify the IP address of the client network gateway.
- **Domain Name:** Specify the domain name of the client network.
- **DB Server:** Edit the host names and IP addresses of the DB servers if applicable.

b. **SCAN Addresses:** Edit the SCAN IP addresses if applicable.

c. **Backup Network**

The backup network is the secondary channel for connectivity to Exadata Database Service on Cloud@Customer resources. It is typically used to segregate application connections on the client network from other network traffic. You can edit the following backup network settings:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the backup network between 1 and 4094, inclusive. To specify no VLAN tagging, enter "1". (This is equivalent to a "NULL" VLAN ID tag value.)

> **Note**
>
> The values "0" and "4095" are reserved and cannot be entered.

- **Netmask:** Specify the IP netmask for the disaster recovery network.
- **Domain Name:** Specify the domain name of the backup network.
- **DB Server:** Edit the host names and IP addresses of the DB servers if applicable.

d. **Disaster Recovery Network**

If so equipped (for authorized customers only), you can edit the following disaster recovery network settings:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the backup network between 1 and 4094, inclusive. To specify no VLAN tagging, enter "1". (This is equivalent to a "NULL" VLAN ID tag value.)

> **Note**
>
> The values "0" and "4095" are reserved and cannot be entered.

- **Netmask:** Specify the IP netmask for the disaster recovery network.
- **Domain Name:** Specify the domain name of the backup network.

e. **Configure DNS and NTP Servers**

The VM cluster network requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. You can edit the following settings:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the client network. You may specify up to three DNS servers.
- **NTP Servers:** Provide the IP address of an NTP server that is accessible using the client network. You may specify up to three NTP servers.

Copyright © 2008, 2026, Oracle and/or its affiliates.

8. Click **Save Changes**.

   After editing, the state of the VM cluster network is **Requires Validation**.

Using the Console to Download a File Containing the VM Cluster Network Configuration Details

To provide VM cluster network information to your network administrator, you can download and supply a file containing the network configuration.

Use this procedure to download a configuration file that you can supply to your network administrator. The file contains the information needed to configure your corporate DNS and other network devices to work along with Exadata Database Service on Cloud@Customer.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you are interested in.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network for which you want to download a file containing the VM cluster network configuration details.

   The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Download Network Configuration**.

   Your browser downloads a file containing the VM cluster network configuration details.

Using the Console to Validate a VM Cluster Network

You can only validate a VM cluster network if its current state is **Requires Validation**, and if the underlying Exadata infrastructure is activated.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to validate.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network that you want to validate.

   The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Validate VM Cluster Network**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Validation performs a series of automated checks on the VM cluster network. The Validate VM Cluster Network button is only available if the VM cluster network requires validation.

7. In the resulting dialog, click **Validate** to confirm the action.

   After successful validation, the state of the VM cluster network changes to **Validated** and the VM cluster network is ready to use. If validation fails for any reason, examine the error message and resolve the issue before repeating validation.

   If network validation fails, then a banner with the **Download Report** button is displayed. Download the report, edit the network resources assignment, and then try validating again.

Using the Console to Download Network Validation Report

Learn to validate and inspect the network validation failure report without active involvement from Oracle Cloud Ops in troubleshooting networking configuration issues.

You can download the network validation report only when the VM Cluster Network State life cycle state is `Validation Failed`.

You cannot download the network validation report for the following VM Cluster Network State life cycle states:

- `Validated` (Successful validation run does not generate a results file)
- `Requires Validation`
- `Updating`
- `Allocated`

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to validate.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network for which you want to download a file containing the VM cluster network configuration details.

   The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Download Network Configuration**.

Example 5-2 Validation Report

```
{
   "Physical_Links":[
      "Issue detected with Physical Links. Please check the cables connected
to ExaCC, or contact Oracle Support."
   ],
   "VLAN":[
      "Gateway  <G1> is not accessible from source <ipaddr > with vlan id
<vlanId>. Please ensure vlan Id, IP address and Gateway are correct",
      "Gateway  <G2> is not accessible from source <ipaddr > with vlan id
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
<vlanId>. Please ensure vlan Id, IP address and Gateway are correct"
   ],
   "Gateway":[
      "Gateway  <G1> is not accessible from source <ipaddr > with vlan id
<vlanId>. Please ensure vlan Id, IP address and Gateway are correct",
      "Gateway  <G2> is not accessible from source <ipaddr > with vlan id
<vlanId>. Please ensure vlan Id, IP address and Gateway are correct"
   ],
   "DNS":[
      "Missing reverse DNS entry(ies) <hostname> for <IP addr> in the DNS
server <dnsIP>. Please update the DNS Server with appropriate entry(ies).",
      "Missing DNS entry(ies) <IP> for <hostname> in the DNS server <dnsIP>.
Please update the DNS Server with appropriate entry(ies).",
      "Wrong reverse DNS entry(ies) <hostname> found for <IP addr> in the DNS
server <dnsIP>. Please update the DNS Server with appropriate entry(ies).",
      "Wrong DNS entry(ies) <IP addr> found for <hostname> in the DNS server
<dnsIP>. Please update the DNS Server with appropriate entry(ies)."
   ],
   "NTP":[
      "NTP <ntpIP> is not accessible from source <domU_ipaddr> with vlan id
<vlanId> using gateway <gatewayIP>. Please ensure NTP, VlanId, IP address,
Gateway are correct."
   ],
   "IP_Availability":[
      "IP <IP1> is already in use. Please ensure this IP is available.",
      "IP <IP2> is already in use. Please ensure this IP is available.",
      "IP <IP3> is already in use. Please ensure this IP is available."
   ]
}
```

Using the Console to Terminate a VM Cluster Network

Before you can terminate a VM cluster network, you must first terminate the associated VM cluster, if one exists, and all the databases it contains.

Terminating a VM cluster network removes it from the Cloud Control Plane.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to terminate.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

   The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network that you want to terminate.

   The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Terminate**.
7. In the resulting dialog, enter the name of the VM cluster network, and click **Terminate VM Cluster Network** to confirm the action.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Add Another DB Server to the VM Cluster Network

To a DB server to the VM Cluster Network, complete this procedure.

> **Note**
>
> Oracle strongly recommends that customers use the latest GI version, or at least GI version 19.11.

1. Click the name of the **VM Cluster Network** that you want to edit.
   Exadata Infrastructure Details page displays the configured DB servers.
2. Click **Add DB Servers to Network** to assign IP addresses.
   Add DB Servers to Network page is displayed.
3. Select one or more DB servers from the list.
4. Click **Add DB Servers to Network**.
5. Enter IP addresses in the **Client network** and **Backup network** sections.
6. Click **Add DB Server to Network IP**.
   Exadata Infrastructure Details page displays a banner with the count of DB servers selected.

   If you want to change the IP addresses, click **Edit VM Cluster Network**. Make changes on the Edit VM Cluster Network page and then click **Save Changes**.
7. Click **Validate VM Cluster Network**.
   Validate VM Cluster Network goes into **validating** state.

   The statuses of the DB servers change to **Validated** if the validation is successful.

Using the Console to Remove DB Servers from a VM Cluster Network

To remove a database server, use this procedure.

> **Note**
>
> You cannot remove network resources that are currently allocated to a virtual machine.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under Infrastructure, click **Exadata Infrastructure**.
3. In the list of Exadata Infrastructures, click the display name of the infrastructure you wish to view details.
4. Click the VM Cluster Network from the list of VM Cluster Networks.

   VM Cluster Network Details page is displayed.
5. Under **Network Configuration**, in the **Client network** or **Backup network** section, identify the database server that you want to remove.
6. Click the **Actions** icon (three dots) for a virtual machine, and then click **Remove**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

7. Click **Remove** on the confirmation dialog displayed.

Manage VM Clusters

Learn how to manage your VM clusters on Oracle Exadata Database Service on Cloud@Customer.

- About Managing VM Clusters on Oracle Exadata Database Service on Cloud@Customer
  The VM cluster provides a link between your Oracle Exadata Database Service on Cloud@Customer infrastructure and Oracle Databases you deploy.
- Overview of VM Cluster Node Subsetting
  VM Cluster Node Subsetting enables you to allocate a subset of database servers to new and existing VM clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.
- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.
- Using the Console to Manage VM Clusters on Oracle Exadata Database Service on Cloud@Customer
  Learn how to use the console to create, edit, and manage your VM Clusters on Oracle Exadata Database Service on Cloud@Customer.
- Using the API to Manage Oracle Exadata Database Service on Cloud@Customer VM Clusters
  Review the list of API calls to manage your Oracle Exadata Database Service on Cloud@Customer VM cluster networks and VM clusters.
- 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

About Managing VM Clusters on Oracle Exadata Database Service on Cloud@Customer

The VM cluster provides a link between your Oracle Exadata Database Service on Cloud@Customer infrastructure and Oracle Databases you deploy.

The VM cluster contains an installation of Oracle Clusterware, which supports databases in the cluster. In the VM cluster definition, you also specify the number of enabled CPU cores, which determines the amount of CPU resources that are available to your databases

Copyright © 2008, 2026, Oracle and/or its affiliates.

Before you can create any databases on your Exadata Cloud@Customer infrastructure, you must create a VM cluster network, and you must associate it with a VM cluster.

> **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.

Overview of VM Cluster Node Subsetting

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

With VM Cluster Node Subsetting, you can:

- Create a smaller VM cluster to host databases that have low resource and scalability requirements or to host a smaller number of databases that require isolation from the rest of the workload.
- Expand or shrink an existing VM cluster by adding and removing nodes to ensure optimal utilization of available resources.

Consider reviewing the points below that will assist you in subsetting VM cluster nodes.

- VM Cluster Node Subsetting capability is available for new and existing VM clusters in Gen2 Exadata Cloud@Customer service.
- All VMs across a VM cluster will have the same resource allocation per VM irrespective of whether the VM was created during cluster provisioning or added later by extending an existing VM cluster.
- VM Clusters only need a minimum of 1 VM with node subsetting. However, Oracle recommends a minimum of 2 VMs per VM Cluster to provide high availability.
- Each VM cluster network is pre-provisioned with IP addresses for every DB Server in the infrastructure. One cluster network can only be used by a single VM cluster and is validated to ensure the IP addresses do not overlap with other cluster networks. Adding or removing VMs to the cluster does not impact the pre-provisioned IP addresses assigned to each DB server in the associated cluster network.

For the Maximum number of VMs per DB server and Maximum number of VM Clusters per System, see the System Shape and Configuration Tables. The Maximum number of VM Clusters per System depends on the resources available per DB server and is subject to the per DB Server maximum VM limit.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> When a cluster contains a node-subsetted database, the attributed usage and cost feature for pluggable databases will not work because the process of creating node-subsetted databases happens on the backend, and the metadata for node-subsetted databases doesn't get synchronized with the Control Plane Server.
>
> However, if the database was originally created without using node-subsetting and later converted to a node-subsetted database, this issue will not arise since the metadata is already available in the Control Plane.

Related Topics

- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Add VMs to a Provisioned Cluster
  To add virtual machines to a provisioned cluster, use this procedure.
- Using the Console to View a List of DB Servers on an Exadata Infrastructure
  To view a list of database server hosts on an Oracle Exadata Cloud@Customer system, use this procedure.
- Using the Console to Remove a VM from a VM Cluster
  To remove a virtual machine from a provisioned cluster, use this procedure.

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).

Copyright © 2008, 2026, Oracle and/or its affiliates.

Disabling diagnostic events and health monitoring will only stop the collection and notification of data/events from the time you uncheck the checkboxes tied to the options. However, historical data will not be purged from Oracle Cloud Operations data repositories.

Related Topics

- Database Service Events
  The Database Service emits events, which are structured messages that indicate changes in resources.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If have previously opted in for incident log and trace file collection and decide to opt out when Oracle Cloud operations run a log collection job, then the job will run its course and will not cancel. Future log collections won't happen until you opt-in again to the incident logs and trace file collection option.
>
> TFA is shipped with scripts that run when a particular component is requested, for example, for CRS component, `crscollect.pl` will run a number of `crsctl` commands and gather the input. By default, TFA does not redact collected logs.

Table 5-13 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 © 2008, 2026, Oracle and/or its affiliates.

Table 5-13 (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 © 2008, 2026, Oracle and/or its affiliates.

Table 5-13 (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 © 2008, 2026, Oracle and/or its affiliates.

Table 5-13 (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`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Health Metrics

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

> **Note**
>
> Oracle may add more metrics in the future, but if you have already chosen to collect metrics, you need not update your opt-in value. It will remain enabled/disabled based on your current preference.

> **Note**
>
> In addition to the metrics listed below, Oracle analyzes additional metrics to provide highest level of service operations and support for ensuring high availability of services.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Guest VM Health Metrics List - Database Metrics

Table 5-14 Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
`CpuUtilization`	CPU Utilization	Percentage	Mean	One minute	Five minutes	The CPU utilization is expressed as a percentage, which is aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs.
`StorageUtilization`	Storage Utilization	Percentage	Mean	One hour	One hout	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.
`BlockChanges`	DB Block Changes	Changes per second	Mean	One minute	Five minutes	The Average number of blocks changed per second.
`ExecuteCount`	Execute Count	Count	Sum	One minute	Five minutes	The number of user and recursive calls that executed SQL statements during the selected interval.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-14 (Cont.) Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
CurrentLogons	Current Logons	Count	Sum	One minute	Five minutes	The number of successful logons during the selected interval.
TransactionCount	Transaction Count	Count	Sum	One minute	Five minutes	The combined number of user commits and user rollbacks during the selected interval.
UserCalls	User Calls	Count	Sum	One minute	Five minutes	The combined number of logons, parses, and execute calls during the selected interval.
ParseCount	Parse Count	Count	Sum	One minute	Five minutes	The number of hard and soft parses during the selected interval.
StorageUsed	Storage Space Used	GB	Max	One hour	One hour	Total amount of storage space used by the database at the collection time.
StorageAllocated	Storage Space Allocated	GB	Max	One hour	One hour	Total amount of storage space allocated to the database at the collection time.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-14 (Cont.) Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
StorageUsedByTablespace	Storage Space Used By Tablespace	GB	Max	One hour	One hour	Total amount of storage space used by tablespace at the collection time. In the case of container databases, this metric provides root container tablespaces.
StorageAllocatedByTablespace	Allocated Storage Space By Tablespace	GB	Max	One hour	One hour	Total amount of storage space allocated to the tablespace at the collection time. In the case of container databases, this metric provides root container tablespaces.
StorageUtilizationByTablespace	Storage Space Utilization By Tablespace	Percentage	Mean	One hour	One hour	This indicates the percentage of storage space utilized by the tablespace at the collection time. In the case of container databases, this metric provides root container tablespaces.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Guest VM Health Metrics List - Non-Database Metrics

Table 5-15 Guest VM Health Metrics List - Non-Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
`ASMDiskgroupUtilization`	ASM Diskgroup Utilization	Percentage	Max	10 minutes	Percentage of usable space used in a Disk Group. Usable space is the space available for growth. DATA disk group stores our Oracle database files. RECO disk group contains database files for recovery such as archives and flashback logs.
`FilesystemUtilization`	Filesystem Utilization	Percentage	Max	One minute	Percent utilization of provisioned filesystem.
`CpuUtilization`	CPU Utilization	Percentage	Mean	One minute	Percent CPU utilization.
`MemoryUtilization`	Memory Utilization	Percentage	Mean	One minute	Percentage of memory available for starting new applications, without swapping. The available memory can be obtained via the following command: `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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-15 (Cont.) Guest VM Health Metrics List - Non-Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
OcpusAllocated	OCPU Allocated	Integer	Max	One minute	The number of OCPUs allocated.

Related Topics

- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Introduction to Scale Up or Scale Down Operations

With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.

- Scaling Up or Scaling Down the VM Cluster Resources
  You can scale up or scale down the memory, local disk size (/u02), ASM Storage, and CPUs (ECPUs for X11M).
- Resizing Memory and Large Pages
- Calculating the ASM Storage
- Estimating How Much Local Storage You Can Provision On Your VMs
- Scaling Local Storage

Scaling Up or Scaling Down the VM Cluster Resources

You can scale up or scale down the memory, local disk size (/u02), ASM Storage, and CPUs (ECPUs for X11M).

Billing Does Not Stop When VMs Are Stopped

Oracle does not stop billing when a VM or VM Cluster is stopped. To stop billing for a VM Cluster, lower the OCPU (ECPUs for X11M) count to zero.

Scaling up or down these resources requires careful auditing of existing usage and capacity management by the customer DB administrator. Review current usage to avoid failures during

Copyright © 2008, 2026, Oracle and/or its affiliates.

or after a scale-down operation. When scaling up, consider how much of these resources are left for the next VM Cluster you plan to create.

Exadata Cloud@Customer Cloud tooling calculates the current usage of memory, local disk, and ASM storage in the VM cluster, adds headroom, and arrives at a minimum value below which you cannot scale down. You must specify values above this minimum to successfully complete scaling operations.

OCPU/ECPU Reservation and Scaling Implications

- When creating or scaling a VM Cluster, setting the number of OCPUs (ECPUs for X11M) to zero will shut down the VM Cluster and eliminate any billing for that VM Cluster, but the hypervisor will still reserve the minimum 2 OCPUs (8 ECPUs for X11M) for each VM. These reserved OCPUs (ECPUs for X11M)cannot be allocated to any other VMs, even though the VM to which they are allocated is shut down. The Control Plane does not account for reserved OCPUs (ECPUs for X11M)when showing maximum available OCPU (ECPU for X11M), so you should account for these reserved OCPUs (ECPUs for X11M)when performing any subsequent scaling operations to ensure the operation can acquire enough OCPUs (ECPUs for X11M) to successfully complete the operation.
- For memory and `/u02` scale up or scale down operations, if the difference between the current value and the new value is less than 2%, then no change will be made to that VM. This is because memory change involves rebooting the VM, and `/u02` change involves bringing down the Oracle Grid Infrastructure stack and un-mounting `/u02`. Production customers will not resize for such a small increase or decrease, and hence, such requests are a no-op.
- You can scale the VM Cluster resources even if any of the DB servers in the VM Cluster are down:
  - If a DB server is down and scaling is performed, the VMs on that server will not be automatically scaled to the new OCPUs when the DB server and the VMs come back online. It's your responsibility to ensure that all the VMs in the cluster have the same OCPU values.
  - Even if the DB server is down, billing does not stop for the VM Cluster that has the VMs on that DB server.

/u02 Scale-Up on KVM Hosts

On KVM hosts (not XEN-based VMs), `/u02` scale-up does not require downtime, as increasing LVM volumes is an online operation.

Resizing ASM Disk Groups and Storage: Guidelines and Considerations

You can resize the ASM disk group allocation using the same options available when provisioning a VM cluster, as outlined in the *Storage Allocation to ASM Disk Groups* section of the *Using the Console to Create an ASM VM Cluster* topic. In addition to those options, the following restrictions and considerations apply.

- Resizing ASM Disk Group
  You cannot reduce a disk group's size below its current used space plus a reserved capacity equal to 15% of the disk group's total size.
- Scaling Exadata Storage
  When scaling Exadata storage:
  - If you do not specify new disk group space allocations, the disk group sizes are adjusted according to the existing percentage allocations.
  - If you specify new disk group space allocations, the disk group sizes are adjusted according to the new percentage allocations.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- ASM Rebalance Considerations
  - Resizing ASM disk groups, whether by changing total capacity or disk group allocations, will trigger an ASM rebalance operation, which may temporarily impact database performance.
  - If both storage scaling and disk group resizing are required, it is recommended to perform both actions together to avoid multiple rebalance operations.
  - In general, resizing ASM disk groups should be avoided unless absolutely necessary, in order to avoid unintended ASM rebalance operations

Resizing Memory and Large Pages

You can scale the database server memory up and down in a VM Cluster. Scaling memory requires a rolling restart of the database servers to take effect. For memory scaling to succeed, the databases must autostart in the Open state.

Changing the memory in a VM Cluster will affect the large pages (HugePages) settings for the VMs in that cluster. When a VM is initially created, each VM's operating system is configured with 50% of the memory allocated to the VM for large pages, and databases are configured to use that memory for their SGA. Oracle recommends you not modify the large pages configuration unless you understand the implication of any changes you make. Improper configurations can prevent all databases from starting, and even prevent the VM from booting.

Although modifying the large pages configuration is not recommended, it is permitted. However, any changes made may be overridden by cloud automation if the VM's memory is resized later. During a memory resize operation, cloud automation will attempt to maintain the large pages memory as a percentage of total memory, with a maximum limit of 60%. If large pages are configured to use more than 60% of the total memory, cloud automation will automatically resize it to this 60% limit.

Once the new large pages allocation is calculated, the automation will perform the following checks:

- Condition 1: The current HugePages usage, multiplied by 1.15 (15% more than currently used), must be less than the new large pages allocation.
- Condition 2: The current HugePages usage, multiplied by 1.15, must also be less than 60% of the new total memory size.

> **Note**
>
> The current HugePages usage is determined by subtracting the free HugePages from the total current HugePages.

If both conditions are met, cloud automation will apply the memory changes to the VMs. If either condition is not met, the process will terminate with an error similar to the following:

```
EXACLOUD: Requested memory is insufficient. The new hugepage count is <<>>, which is
less than the minimum required for the VM. Not proceeding with the change.
```

This process ensures there is enough conventional memory for the VM to boot. Before proceeding with the resize, automation performs a precheck to determine the current large pages usage by running database instances. If the precheck indicates that there will not be enough large pages memory after the resize to support the existing databases, the resize will fail, and the process will not continue.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Calculating the ASM Storage

Oracle chosen default storage sizes

Each ASM VM Cluster will have its own pool of Exadata Storage. That pool will be split between the following disk groups, with each disk group getting the percent specified in the table. The ratios vary depending on whether the VM Cluster is configured to support local backups, or for Exadata snapshots. The disk groups sizes are fixed and cannot be changed post deployment.

Table 5-16 Disk Group sizes

Disk Group	Local backups = No Snapshots =No	Local backups = Yes Snapshots = No	Local backups = No Snapshots= Yes	Local backups = Yes Snapshots = Yes
+DATA	80%	40%	60%	35%
+RECO	20%	60%	20%	50%
+SPARSE	0%	0%	20%	15%

Use the following formula to calculate the minimum required ASM storage:

- For each disk group, for example, `DATA`, `RECO`, note the total size and free size by running the `asmcmd lsdg` command on any Guest VM of the VM cluster.
- Calculate the used size as (Total size - Free size) / 3 for each disk group. The /3 is used because the disk groups are triple mirrored.
- DATA:RECO ratio is:

  80:20 if **Local Backups** option was NOT selected in the user interface.

  40:60 if **Local Backups** option was selected in the user interface.
- Ensure that the new total size as given in the user interface passes the following conditions:
  Used size for DATA * 1.15 <= (New Total size * DATA %)

  Used size for RECO * 1.15 <= (New Total size * RECO %)

Example 5-3 Calculating the ASM Storage

1. Run the `asmcmd lsdg` command in the Guest VM:
   - Without SPARSE:

```
/u01/app/19.0.0.0/grid/bin/asmcmd lsdg
ASMCMD>
State   Type Rebal Sector Logical_Sector Block AU     Total_MB
Free_MB    Req_mir_free_MB   Usable_file_MB   Offline_disks
Voting_files   Name
MOUNTED HIGH N        512     512        4096 4194304 12591936
10426224   1399104           3009040           0
Y      DATAC5/
MOUNTED HIGH N        512     512        4096 4194304 3135456
3036336    348384            895984            0
N      RECOC5/
ASMCMD>
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

- With SPARSE:

```
/u01/app/19.0.0.0/grid/bin/asmcmd lsdg
ASMCMD>
State   Type Rebal Sector Logical_Sector Block AU       Total_MB
Free_MB   Req_mir_free_MB   Usable_file_MB   Offline_disks
Voting_files   Name
MOUNTED HIGH N        512     512        4096 4194304   12591936
10426224  1399104           3009040            0
Y     DATAC5/
MOUNTED HIGH N        512     512        4096 4194304   3135456
3036336   348384            895984             0
N     RECOC5/
MOUNTED HIGH N        512     512        4096 4194304   31354560
31354500  3483840           8959840            0
N     SPRC5/
ASMCMD>
```

> **Note**
>
> The listed values of all attributes for SPARSE diskgroup (SPRC5) present the virtual size. In Exadata DB Systems and Exadata Cloud@Customer, we use the ratio of 1:10 for `physicalSize:virtualSize`. Hence, for all purposes of our calculation we must use 1/10th of the values displayed above in case of SPARSE for those attributes.

2. Used size for a disk group = (Total_MB - Free_MB) /3
   - Without SPARSE:
     Used size for DATAC5 = (12591936 - 10426224) / 3 = 704.98 GB

     Used size for RECO5 = (3135456 - 3036336) / 3 = 32.26 GB
   - With SPARSE:
     Used size for DATAC5 = (12591936 - 10426224) / 3 ~= 704.98 GB

     Used size for RECO5 = (3135456 - 3036336) /3 ~= 32.26 GB

     Used size for SPC5 = (1/10 * (31354560 - 31354500)) / 3 ~= 0 GB
3. Storage distribution among diskgroups
   - Without SPARSE:
     DATA:RECO ratio is 80:20 in this example.
   - With SPARSE:
     DATA RECO: SPARSE ratio is 60:20:20 in this example.
4. New requested size should pass the following conditions:
   - Without SPARSE: (For example, 5 TB in user interface.)
     5 TB = 5120 GB ; 5120 *.8 = 4096 GB; 5120 *.2 = 1024 GB

     For DATA: (704.98 * 1.15) <= 4096 GB

     For RECO: (32.36 * 1.15) <= 1024 GB
   - With SPARSE: (For example, 8 TB in the user interface.)
     8 TB = 8192 GB; 8192 *.6 = 4915 GB; 8192 *.2 = 1638 GB; 8192 *.2 = 1638 GB

     For DATA: (704.98 * 1.15) <= 4915 GB

Copyright © 2008, 2026, Oracle and/or its affiliates.

For RECO: (32.36 * 1.15) <= 1638 GB

For SPARSE: (0 * 1.15) <= 1638 GB

Above resize will go through. If above conditions are not met by the new size, then resize will fail the precheck.

Custom storage sizes

The following minimum disk group size requirements apply:

- **DATA disk group:** Minimum of 10% of total storage.
- **RECO disk group:** Minimum of 10% of total storage.
- **SPARSE disk group:** Minimum of 0%. If you specify a non-zero value, the minimum is 10%. The maximum is subject to Exadata system limits.

> **Note**
>
> You cannot reduce a disk group's size below its current used space plus a reserved capacity equal to 15% of the disk group's total size. This restriction applies in addition to the requirements described above.

Estimating How Much Local Storage You Can Provision On Your VMs

VM Images include the files necessary to boot and run the VM and its operating system, as well as space for Oracle Homes stored in `/u02`. To estimate how much additional local storage space beyond the minimum can be allocated to any file system associated with a VM, subtract the size of the VM images for all VMs on a server from the total available space. If you have not modified the default VM Image size by expanding any file systems, use the VM Image size (default and minimum) below. If you have or plan to modify your VM Image size, you must use the OCI console and "Scale VM Cluster" action to check the allocated and available for an existing VM Cluster as expanding some non-`/u02` file systems will consume more incremental storage than was added to the file system. This information is also available in the "Configure VM Cluster" action while creating a new VM Cluster.

X8-2 and X7-2 Systems

- Total space available for VM images (X7 All Systems): 1237 GB
- Total space available for VM images (X8 All Systems): 1037 GB
- VM Image size (default and minimum) including `/u02`: 244 GB
- Default (minimum) `/u02`: 60 GB

X8M-2 Systems

- Total space available for VM images (X8M Base System): 1237 GB
- Total space available for VM images (X8M Elastic): 2500 GB
- VM Image size (default and minimum) including `/u02`: 244 GB
- Default (minimum) `/u02`: 60 GB

X11M, X10M, and X9M-2 Systems

- Total Available for VM Images (Base System X9M): 1077 GB

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Total Available for VM Images (Elastic): 2243 GB
- VM Image size (default and minimum) including `/u02`: 244 GB
- Default (minimum) `/u02`: 60 GB

Scaling Local Storage

Scale Local Space Operation Guidelines

You can scale local storage by modifying the size of many of the individual file systems in a VM. By default, the file systems are created at their minimum size. You can increase the size of the file systems as required. However, note that you can only shrink `/u02`. Other file systems can only be increased in size. The maximum supported size of any file system is 900 GB.

The storage consumed by all file systems is greater than the sum of the file system sizes. Refer to the calculations displayed in the OCI console to see the effects on free local storage when resizing a file system.

Using the OCI Console or API, you can increase or decrease the size of the following local file systems:

- `/u02`

Using the OCI Console or API, you can increase the size of following local filesystems:

- `/`
- `/u01`
- `/tmp`
- `/var`
- `/var/log`
- `/var/log/audit`
- `/home`

However, you cannot resize the following local file systems:

- `/crashfiles`
- `/boot`
- `/acfs01`
- `/u01/app/19.0.0.0/grid`

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> - With the exception of `/u02`, you can only expand the file systems and cannot reduce their size once they have been expanded.
> - For X8M and later models, scaling up or scaling down the `/u02` file system, as well as expanding other Guest VM file systems, are online operations that do not require a rolling restart of the Guest VM. However, for X8 systems that use the Xen-based hypervisor, file system resize operations still require a rolling restart of the Guest VM.
> - Each file system can only be expanded to a maximum of 900 GB
> - Ability to increase the size of additional local file systems is only supported on X8M and later systems.

For more information about resizing these file systems, see *Estimating How Much Local Storage You Can Provision to Your VMs*.

Resource Limit Based On Current Utilization

- Any scale-down operation must leave 15% buffer on top of highest local space utilization across all nodes in the cluster.
- The lowest local space per node allowed is higher of the above two limits.
- Run the `df -kh` command on each node to find out the node with the highest local storage.
- You can also use the utility like `cssh` to issue the same command from all hosts in a cluster by typing it just once.
- Lowest value of local storage each node can be scaled down to would be = 1.15x (highest value of local space used among all nodes).

ACFS File Systems

If requested by support, you can also resize the `/acfs01` file system. This file system is used by the system to stage software. It uses Exadata storage and is not subject to the limits described above for `/u02`. It is a shared file system visible from all nodes in the cluster, and can be online resized from the command line of any VM.

- **Default size:** The default size of `/acfs01` is 100 GB.
- **Scaling /acfs01:** You can scale `acfs01` as user grid from any VM via the `/sbin/acfsutil` command. No reboot is required. The resize operation will not affect the availability of the database service running in the VM Cluster. The following command issued by the `grid` user will increase the size of `/acfs01` by 100 GB: `/sbin/acfsutil size +100 GB /acfs01`.
- You can create additional ACFS file systems if required. These will also consume storage from the Exadata Storage diskgroups and may be shared across all VMs in the cluster. Refer to the ACFS documentation for more information.

Using the Console to Manage VM Clusters on Oracle Exadata Database Service on Cloud@Customer

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create an Exascale VM Cluster
  To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- VM Cloud Automation Software Update Management
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Using the Console to Add VMs to a Provisioned Cluster
  To add virtual machines to a provisioned cluster, use this procedure.
- Using the Console to View a List of DB Servers on an Exadata Infrastructure
  To view a list of database server hosts on an Oracle Exadata Cloud@Customer system, use this procedure.
- Using the Console to Remove a VM from a VM Cluster
  To remove a virtual machine from a provisioned cluster, use this procedure.
- 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.
- Using the Console to Add SSH Keys After Creating a VM Cluster
- Using the Console to Scale the Resources on a VM Cluster
  Starting in Oracle Exadata Database Service on Cloud@Customer, you can scale up or down multiple resources at the same time. You can also scale up or down resources one at a time.
- Using the Console to Stop, Start, or Reboot a VM Cluster Virtual Machine
  Use the console to stop, start, or reboot a virtual machine.
- Using the Console to Check the Status of a VM Cluster Virtual Machine
  Review the health status of a VM cluster virtual machine.
- Using the Console to Move a VM Cluster to Another Compartment
  To change the compartment that contains your VM cluster on Oracle Exadata Database Service on Cloud@Customer, 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.

Using the Console to Create an ASM VM Cluster

To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Exadata VM Clusters**.
4. Click **Create Exadata VM Cluster**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. Provide the requested information on the Create Exadata VM Cluster page:
   a. **Provide the display name:** The display name is a user-friendly name that you can use to identify the VM cluster. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the VM cluster.
   b. **Select a compartment:** From the list of available compartments, choose the compartment that you want to contain the VM cluster.
   c. **Select an Exadata Infrastructure:** From the list, choose the Exadata infrastructure to host the VM cluster. You are not able to create a VM cluster without available and active Exadata infrastructure.
   d. **Select a VM Cluster Network:** From the list, choose a VM cluster network definition to use for the VM cluster. You must have an available and validated VM cluster network before you can create a VM cluster.
   e. **VM Cluster Type:**

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

      - **Exadata Database:** Standard Database VM with no restrictions, suitable for all workloads.
      - **Exadata Database-Developer:** Developer Database VM with restrictions, suitable for application development only.

   f. **Configure:**
      - **DB Servers:**
        – Click **Change DB Servers** for VM placement to allocate VM resources.
        – On the **Change DB Servers** dialog:
          VM Cluster Type - Exadata Database: Select a minimum of one database server for VM placement. If you require a high availability database service that remains available during maintenance and unplanned outages, select at least two database servers. Maximum resources available for allocation per VM are based on the number of database servers selected.

          VM Cluster Type - Exadata Database-Developer: Select one database server for VM placement. Only one database server may be selected.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> * DB Servers, which already have 8 VMs running on them are not available for selection.
> * When calculating maximum local storage resources across selected DB Servers, the reserved local storage needed by the system to host a VM is deducted from the DB Server with the least resources.
> For example, if the local storage available across selected DB servers is 823 GB for DB Server 3 and 813 GB for DB Server 4, then the minimum across selected servers is 813 GB and the maximum available for resource allocation is 813 GB - 184 GB (reserved local storage for hosting VM on X8M DB servers) = 629 GB.
>
> For more information, see *Estimating How Much Local Storage You Can Provision to Your VMs*.

  - Click **Save**.

- **Specify the OCPU (ECPUs for X11M) count per VM:** Specify the OCPU (ECPUs for X11M) count to be provisioned for each VM in this cluster. The minimum value is 2 OCPUs per VM or 8 ECPUs per VM for X11M (for a live VM condition), unless you are specifying zero OCPUs or 0 ECPUs for X11M (for a shutdown VM condition).
If you specify a value of zero, then the VM cluster virtual machines are all shut down at the end of the cluster creation process. In this case, you can later start the virtual machines by scaling the OCPU (ECPUs for X11M) resources. See *Using the Console to Scale the Resources on a VM Cluster*.

  The OCPU (ECPUs for X11M) count for the whole VM Cluster will be calculated automatically based on the per-VM OCPU (ECPUs for X11M) count you have specified and the number of physical Database Servers configured for the VM Cluster.

  OCPU: An Oracle Compute Unit (OCPU) provides CPU capacity equivalent of one physical core of an Intel Xeon processor with hyperthreading enabled. Each OCPU corresponds to two hardware execution threads, known as vCPUs.

  See, *Oracle Platform as a Service and Infrastructure as a Service – Public Cloud Service DescriptionsMetered & Non-Metered*.

  ECPU: An ECPU is an abstracted measure of compute resources. ECPUs are based on the number of cores elastically allocated from a pool of compute and storage servers.

- **Requested OCPU (ECPUs for X11M) count for the VM Cluster:** Displays the total number of CPU cores allocated to the VM cluster based on the value you specified in the **Specify the OCPU (ECPUs for X11M) count per VM** field. This field is not editable.
- **Specify the memory per VM (GB):** Specify the memory for each individual VM. The value must be a multiple of 1 GB and is limited by the available memory on the Exadata infrastructure.
- **Requested memory for the VM Cluster (GB):** Displays the total amount of memory allocated to the VM cluster based on the value you specified in the **Specify the memory per VM (GB)** field. This field is not editable.

Copyright © 2008, 2026, Oracle and/or its affiliates.

g. **VM Storage:** This section lets you specify storage options for the VM image and VM backups. The VM image is the total storage needed by all file systems during cluster provisioning. VM backups are managed by Oracle and store the configured VM's backup data. Previously, both were stored on local database server storage. Now, you can select Exascale storage for the VM image, VM backups, or both.

> **Note**
>
> VM clusters using ASM for databases can also use Exascale for VM storage.

- **VM file system storage:** Specify the storage for the VM Image (VM file system storage).
  - **Database server (default):** Uses local database server storage for all VM file systems.
  - **Exascale storage:** Uses Exascale storage for all VM file systems.
- **VM file system storage capacity per VM (GB):** Displays the total storage required to provision the VM. This field is not editable. This value is calculated by adding the storage allocated to each guest VM file systems, such as `/u01`, `/u02`, `/(root)`, `/var` etc. To view current allocation details, click **Change file system configuration**. You can update the storage allocation for any file system by entering the new value within the displayed storage limits. Click **Update**. The VM file system storage capacity per VM (GB) field now displays the updated value.
- **Total VM file system storage capacity across VM cluster (GB):** Displays the total storage required to provision the VM cluster. This field is not editable. This value is calculated by multiplying the number of nodes by the VM file system storage required per VM.
- **VM backup storage:** Specify the storage for the Oracle-managed VM backups.
  - **Database server (default):** Uses local database server storage for VM backups.
  - **Exascale:** Uses Exascale storage for VM backups.
- **Backup storage per VM (GB):** Displays the total storage required for a single VM backup. This field is hidden, If you choose Exascale as the VM backup storage.
- **Backup storage across VM cluster:** Total backup storage across VM cluster (GB): If you choose Database server for VM backups, it displays the total storage required across the VM cluster for one VM backup whereas if you choose Exascale for VM backups, it displays the total storage required across the VM cluster for two VM backups.

h. **Exadata Storage:** The following settings define how the Exadata storage is configured for use with the VM cluster. The storage type once selected cannot be changed later on once the VM cluster is provisioned with the desired storage type. You have two options to choose: Automatic storage type (ASM) and Exascale. For more information about Exascale storage type, see *Using the Console to Create an Exascale VM Cluster*.

Automatic Storage Management (ASM)

- **Usable Exadata Storage (TB):** Specify the size for each individual VM. The minimum recommended size is 2 TB.
- **Storage allocation to ASM disk groups:** You have following options to choose:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Allocate Oracle chosen default storage sizes:** Selecting this option allocates Oracle-defined default storage sizes to the DATA, RECO, and SPARSE ASM disk groups.
  - **Allocate storage for Exadata snapshots:** Check this option to create a sparse disk group, which is required to support Exadata snapshot functionality. Exadata snapshots enable space-efficient clones of Oracle databases that can be created and destroyed very quickly and easily.
  - **Allocate Storage for local backups:** Check this option to configure the Exadata storage to enable local database backups. If you select this option, more space is allocated to the RECO disk group to accommodate the backups. If you do not select this option, you cannot use local Exadata storage as a backup destination for any databases in the VM cluster.

Table 5-17 Storage Allocation

Storage Allocation	DATA Disk Group	RECO Disk Group	SPARSE Disk Group
Exadata Snapshots: **No** Enable Backups on Local Exadata Storage: **No**	80%	20%	0% (The SPARSE disk group is not created.)
Exadata Snapshots: **No** Enable Backups on Local Exadata Storage: **Yes**	40%	60%	0% (The SPARSE disk group is not created.)
Allocate Storage for Exadata Snapshots: **Yes** Enable Backups on Local Exadata Storage: **No**	60%	20%	20%
Allocate Storage for Exadata Snapshots: **Yes** Enable Backups on Local Exadata Storage: **Yes**	35%	50%	15%

- **Allocate custom storage sizes:** Select this option to specify custom storage sizes for the DATA, RECO, and SPARSE ASM disk groups, subject to system-defined limits. These values override the Oracle-selected default allocations. The following minimum disk group size requirements apply:
  - DATA disk group: Minimum of 10% of total storage.
  - RECO disk group: Minimum of 10% of total storage.
  - SPARSE disk group: Minimum of 0%. If you specify a non-zero value, the minimum is 10%. The maximum is subject to Exadata system limits.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> After provisioning, you cannot reduce a disk group's size below its current used space plus a reserved capacity equal to 15% of the disk group's total size. This restriction applies in addition to the requirements described above.

i. **Select version:**

- **Choose the Oracle Grid Infrastructure release:** From the drop-down menu, choose either 19c or 26ai as the Grid Infrastructure release.

> **Note**
>
> The selected Oracle Grid Infrastructure release determines which Oracle Database versions are supported on the VM cluster.
>
> You cannot provision an Oracle Database version that is newer than the selected Grid Infrastructure version. For example, to provision a 26ai database, you must choose the 26ai Grid Infrastructure release.

> **Note**
>
> To provision a VM cluster with Grid Infrastructure 26ai, the following minimum software versions are required:
>
> – Exadata Guest VM running Exadata System Software 23.1.8
> – Exadata Infrastructure running Exadata System Software 23.1.x

- The **Display All Available Grid Infrastructure Versions** switch is OFF by default. When it is OFF, the system automatically uses the latest Grid Infrastructure version during VM cluster provisioning.
  When toggled ON, the drop-down menu displays the latest four supported Grid Infrastructure versions, allowing you to select your preferred version from the list.
- **Choose an Exadata guest version:**
  - **Exadata infrastructure with Oracle Linux 7 and Exadata image version 22.1.10.0.0.230422:**
    * The **Change image** button is not enabled.
    * The Oracle Grid Infrastructure version defaults to 19.0.0.0.0.
    * The Exadata guest version will be the same as that of the host OS.
  - **Exadata infrastructure with Oracle Linux 8 and Exadata image version 23.1.3.0.0.230613:**
    * The Exadata guest version defaults to the latest (23.1.3.0).
    * The Oracle Grid Infrastructure version defaults to 19.0.0.0.0
    * The **Change image** button is enabled.
    * Click **Change image**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The resulting Change image panel displays the list of available major versions of Exadata image (23.1.3.0 and 22.1.3.0).

The most recent release for each major version is indicated by "(latest)".

* Slide **Display all available versions**.
  Six past versions including the latest versions of Exadata images 23.1.3.0 and 22.1.3.0 are displayed.
* Choose a version.
* Click **Save Changes**.

j. **Add SSH Key:** Specify the public key portion of an SSH key pair that you want to use to access the VM cluster virtual machines. You can upload a file containing the key, or paste the SSH key string.

   To provide multiple keys, upload multiple key files or paste each key into a separate field. For pasted keys, ensure that each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.

k. **Choose a license type:**

   - **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.

     > **Note**
     >
     > BYOL is not available for Exadata Database-Developer VM Cluster type.

   - **License Included:** Select this option to subscribe to Oracle Database software licenses as part of Exadata Database Service on Cloud@Customer.

l. **Diagnostics Collection:**

   By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes. For more information, see *Getting Started with Events*.

   > **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.

   - **Enable Diagnostic Events:** Allow Oracle to collect and publish critical, warning, error, and information events to me.
   - **Enable Health Monitoring:** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
   - **Enable Incident Logs and Trace Collection:** Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

     All three checkboxes are selected by default. You can leave the default settings as is or clear the checkboxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information** >> **Diagnostics Collection**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **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).

m. **Show Advanced Options:**

- **Time zone:** The default time zone for the Exadata Infrastructure is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

> **Note**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, then select the **Select another time zone** option, select a **Region** or **country**, and then select the corresponding **Time zone**.
>
> If you do not see the region or country you want, then select **Miscellaneous**, and then select an appropriate **Time zone**.

- **Cloud Automation Update:** Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster. Set the start time for cloud automation updates.

> **Note**
>
> Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable early access for cloud tools update: VM clusters designated for early access receive updates 1-2 weeks before they are available to other systems. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period: Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> – The freeze period can extend for a maximum of 45 days from the start date.
>
> – Oracle automation will automatically apply updates with critical security fixes (CVSS >= 9) even during a configured freeze period.

- **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.

6. Optionally, you can save the resource configuration as a stack.
   - **To save the resource configuration as a Stack:**
     a. Click **Save as Stack**.
     b. In the resulting Save as Stack dialog, provide the following details:
        i. **Name:** (Optional) Provide an easy to remember descriptive name.
        ii. **Description:** (Optional) Enter a short description.
        iii. **Compartment:** Select a compartment where this Stack will reside.
        iv. **Tags:** Add tags.
     c. Click **Save**.
        After saving the Stack, the system displays a banner with a link to the saved Stack.
     d. Click the link to open the Stack in the Resource Manager Service console. See, *Resource Manager and Terraform*.
   - **To view the details of a Stack:**
     a. Open the navigation menu. Under **Developer Services**, click **Resource Manager**.
     b. Click **Stacks**.
     c. Click the name of the Stack that you want to view details.
        Or, click the Actions menu (three dots), and select the **View stack details** option.
7. Click **Create VM Cluster**.

   The VM Cluster Details page is now displayed. While the creation process is running, the state of the VM cluster is **Pending**. When the VM cluster creation process completes, the state of the VM cluster changes to **Available**.

   The Exadata Database Storage section on the VM Cluster Details page shows the type of storage configured, which, in this case, is ASM.

Related Topics

- Oracle Exadata Database Service on Cloud@Customer Service Description
  Learn how you can leverage the combined capabilities of Oracle Exadata and Oracle Cloud Infrastructure with Oracle Exadata Database Service on Cloud@Customer
- Using the Console to Scale the Resources on a VM Cluster
  Starting in Oracle Exadata Database Service on Cloud@Customer, you can scale up or down multiple resources at the same time. You can also scale up or down resources one at a time.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.
- Estimating How Much Local Storage You Can Provision On Your VMs
- Resource Tags
- Oracle PaaS/IaaS Cloud Service Description documents
- Oracle Platform as a Service and Infrastructure as a Service – Public Cloud Service DescriptionsMetered & Non-Metered
- Getting Started with Events
- Overview of Database Service Events
- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Resource Manager and Terraform

Using the Console to Create an Exascale VM Cluster

To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

To create an Exascale VM cluster, ensure that you that have:

- Active Exadata infrastructure available to host the VM cluster.
- A validated VM cluster network available for the VM cluster to use.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Exadata VM Clusters**.
4. Click **Create Exadata VM Cluster**.
5. Provide the requested information on the Create Exadata VM Cluster page:
   a. **Select a compartment:** From the list of available compartments, choose the compartment that you want to contain the VM cluster.
   b. **Provide the display name:** The display name is a user-friendly name that you can use to identify the VM cluster. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the VM cluster.

Copyright © 2008, 2026, Oracle and/or its affiliates.

c. **Select Exadata Infrastructure:** From the list, choose the Exadata infrastructure to host the VM cluster. You are not able to create a VM cluster without available and active Exadata infrastructure.

d. **Select a VM Cluster Network:** From the list, choose a VM cluster network definition to use for the VM cluster. You must have an available and validated VM cluster network before you can create a VM cluster.

e. **VM Cluster Type**

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

- **Exadata Database:** Standard Database VM with no restrictions, suitable for all workloads.
- **Exadata Database-Developer:** Developer Database VM with restrictions, suitable for application development only.

f. **Configure:**

- **DB Servers:**
  - Click **Change DB Servers** for VM placement to allocate VM resources.
  - On the **Change DB Servers** dialog:
    VM Cluster Type - Exadata Database: Select a minimum of one database server for VM placement. If you require a high availability database service that remains available during maintenance and unplanned outages, select at least two database servers. Maximum resources available for allocation per VM are based on the number of database servers selected.

    VM Cluster Type - Exadata Database-Developer: Select one database server for VM placement. Only a single database server may be selected.

    > **Note**
    >
    > * DB Servers, which already have 8 VMs running on them are not available for selection.
    > * When calculating maximum local storage resources across selected DB Servers, the reserved local storage needed by the system to host a VM is deducted from the DB Server with the least resources.
    > For example, if the local storage available across selected DB servers is 823 GB for DB Server 3 and 813 GB for DB Server 4, then the minimum across selected servers is 813 GB and the maximum available for resource allocation is 813 GB - 184 GB (reserved local storage for hosting VM on X8M DB servers) = 629 GB.
    >
    > For more information, see *Estimating How Much Local Storage You Can Provision to Your VMs*.

  - Click **Save**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Specify the OCPU (ECPUs for X11M) count per VM:** Specify the OCPU (ECPUs for X11M) count to be provisioned for each VM in this cluster. The minimum value is 2 OCPUs per VM or 8 ECPUs per VM for X11M (for a live VM condition), unless you are specifying zero OCPUs or 0 ECPUs for X11M (for a shutdown VM condition).
  If you specify a value of zero, then the VM cluster virtual machines are all shut down at the end of the cluster creation process. In this case, you can later start the virtual machines by scaling the OCPU (ECPUs for X11M) resources. See *Using the Console to Scale the Resources on a VM Cluster*.

  The OCPU (ECPUs for X11M) count for the whole VM Cluster will be calculated automatically based on the per-VM OCPU (ECPUs for X11M) count you have specified and the number of physical Database Servers configured for the VM Cluster.

  OCPU: An Oracle Compute Unit (OCPU) provides CPU capacity equivalent of one physical core of an Intel Xeon processor with hyperthreading enabled. Each OCPU corresponds to two hardware execution threads, known as vCPUs.

  See, *Oracle Platform as a Service and Infrastructure as a Service – Public Cloud Service DescriptionsMetered & Non-Metered*.

  ECPU: An ECPU is an abstracted measure of compute resources. ECPUs are based on the number of cores elastically allocated from a pool of compute and storage servers.
- **Requested OCPU (ECPUs for X11M) count for the VM Cluster:** Displays the total number of CPU cores allocated to the VM cluster based on the value you specified in the **Specify the OCPU (ECPUs for X11M) count per VM** field. This field is not editable.
- **Specify the memory per VM (GB):** Specify the memory for each individual VM. The value must be a multiple of 1 GB and is limited by the available memory on the Exadata infrastructure.
- **Requested memory for the VM Cluster (GB):** Displays the total amount of memory allocated to the VM cluster based on the value you specified in the **Specify the memory per VM (GB)** field. This field is not editable.

g. **VM Storage:** This section lets you specify storage options for the VM image and VM backups. The VM image is the total storage needed by all file systems during cluster provisioning. VM backups are managed by Oracle and store the configured VM's backup data. Previously, both were stored on local database server storage. Now, you can select Exascale storage for the VM image, VM backups, or both.

> **Note**
>
> VM clusters using ASM for databases can also use Exascale for VM storage.

- **VM file system storage:** Specify the storage for the VM Image (VM file system storage).
  - **Database server (default):** Uses local database server storage for all VM file systems.
  - **Exascale storage:** Uses Exascale storage for all VM file systems.
- **VM file system storage capacity per VM (GB):** Displays the total storage required to provision the VM. This field is not editable. This value is calculated by

Copyright © 2008, 2026, Oracle and/or its affiliates.

adding the storage allocated to each guest VM file systems, such as `/u01`, `/u02`, `/(root)`, `/var` etc. To view current allocation details, click **Change file system configuration**. You can update the storage allocation for any file system by entering the new value within the displayed storage limits. Click **Update**. The VM file system storage capacity per VM (GB) field now displays the updated value.

- **Total VM file system storage capacity across VM cluster (GB):** Displays the total storage required to provision the VM cluster. This field is not editable. This value is calculated by multiplying the number of nodes by the VM file system storage required per VM.
- **VM backup storage:** Specify the storage for the Oracle-managed VM backups.
  - **Database server (default):** Uses local database server storage for VM backups.
  - **Exascale:** Uses Exascale storage for VM backups.
- **Backup storage per VM (GB):** Displays the total storage required for a single VM backup. This field is hidden, If you choose Exascale as the VM backup storage.
- **Backup storage across VM cluster:** Total backup storage across VM cluster (GB): If you choose Database server for VM backups, it displays the total storage required across the VM cluster for one VM backup whereas if you choose Exascale for VM backups, it displays the total storage required across the VM cluster for two VM backups.

h. **Storage type:** When creating a VM cluster, you must choose storage type for the database. You can select one of two options: Automatic Storage Management (ASM) or Exascale. This choice is permanent and cannot be changed after the VM is provisioned. For more information about ASM storage type, see *Using the Console to Create an ASM VM Cluster*.

> **Note**
>
> Minimum Requirements for Configuring Exascale Storage
>
> - Supported on Exadata Infrastructure Model X8M or later.
> - Requires Exadata System Software release 24.1 or later.
> - Requires Oracle Grid Infrastructure version 26ai (24.3)
> - Supports Oracle Database versions 26ai (23.4) or later.
> - Storage servers (cells) must be running release 25.1.9 or later, or 25.2.2 or later.
>   If storage servers are running release 25.2.1, the Exascale feature must be disabled.
>
> Exascale option will be disabled If the minimum requirement is not met.

Exascale database storage vault:

- **Create new storage vault:** Choose this option to create a new Exascale database storage vault during VM cluster provisioning.
  - **Storage vault name:** Enter a descriptive name for the vault. Click the **change compartment** link and choose a compartment if you want to create this vault in a different compartment.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Storage capacity for databases:** Enter the storage capacity for the databases within the minimum and maximum values displayed on the screen.

> **Note**
>
> If additional space is needed beyond the maximum shown, the Exascale capacity must be increased. For more information, see *Using the Console to Scale an Exascale Storage Vault*.

- **Select existing storage vault:** Select a vault that resides in the compartment of your choice.

i. **Select version:**

> **Note**
>
> Only Oracle Database 26ai can be provisioned on the Exascale VM cluster.

- **Choose the Oracle Grid Infrastructure release:** The Oracle Grid Infrastructure release defaults to 26ai.

> **Note**
>
> The selected Oracle Grid Infrastructure release determines which Oracle Database versions are supported on the VM cluster.
>
> You cannot provision an Oracle Database version that is newer than the selected Grid Infrastructure version. For example, to provision a 26ai database, you must choose the 26ai Grid Infrastructure release.

> **Note**
>
> To provision a VM cluster with Grid Infrastructure 26ai, the following minimum software versions are required:
>
> - Exadata Guest VM running Exadata System Software 23.1.8
> - Exadata Infrastructure running Exadata System Software 23.1.x

- The **Display All Available Grid Infrastructure Versions** switch is OFF by default. When it is OFF, the system automatically uses the latest Grid Infrastructure version during VM cluster provisioning.
  When toggled ON, the drop-down menu displays the latest four supported Grid Infrastructure versions, allowing you to select your preferred version from the list.
- **Choose an Exadata guest version:**
  - The Exadata guest version defaults to the latest (24.1.6.0)
  - The Oracle Grid Infrastructure version defaults to 26ai
  - The **Change image** button is enabled.
  - Click **Change image**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The resulting Change image panel displays the list of available major versions of Exadata image (24.1.6.0 and later).

The most recent release for each major version is indicated by "(latest)".

- Slide **Display all available versions**.
  Six past versions including the latest versions of Exadata images 24.1.6.0 and later are displayed.
- Choose a version.
- Click **Save Changes**.

j. **Add SSH Key:** Specify the public key portion of an SSH key pair that you want to use to access the VM cluster virtual machines. You can upload a file containing the key, or paste the SSH key string.

   To provide multiple keys, upload multiple key files or paste each key into a separate field. For pasted keys, ensure that each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.

k. **Choose a license type:**

   - **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.

     > **Note**
     >
     > BYOL is not available with Exadata Database-Developer VM Cluster type.

   - **License Included:** Select this option to subscribe to Oracle Database software licenses as part of Exadata Database Service on Cloud@Customer.

l. **Diagnostics Collection:**

   By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes. For more information, see *Getting Started with Events*.

   > **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.

   - **Enable Diagnostic Events:** Allow Oracle to collect and publish critical, warning, error, and information events to me.
   - **Enable Health Monitoring:** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
   - **Enable Incident Logs and Trace Collection:** Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

     All three checkboxes are selected by default. You can leave the default settings as is or clear the checkboxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information** >> **Diagnostics Collection**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **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).

m. **Show Advanced Options:**

- **Time zone:** The default time zone for the Exadata Infrastructure is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

> **Note**
>
> If you want to set a time zone other than UTC or the browser-detected time zone, then select the **Select another time zone** option, select a **Region** or **country**, and then select the corresponding **Time zone**.
>
> If you do not see the region or country you want, then select **Miscellaneous**, and then select an appropriate **Time zone**.

- **Cloud Automation Update:** Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster. Set the start time for cloud automation updates.

> **Note**
>
> Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable early access for cloud tools update: VM clusters designated for early access receive updates 1-2 weeks before they are available to other systems. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period: Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> – The freeze period can extend for a maximum of 45 days from the start date.
>
> – Oracle automation will automatically apply updates with critical security fixes (CVSS >= 9) even during a configured freeze period.

- **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.

6. Optionally, you can save the resource configuration as a stack.
   - **To save the resource configuration as a Stack:**
     a. Click **Save as Stack**.
     b. In the resulting Save as Stack dialog, provide the following details:
        i. **Name:** (Optional) Provide an easy to remember descriptive name.
        ii. **Description:** (Optional) Enter a short description.
        iii. **Compartment:** Select a compartment where this Stack will reside.
        iv. **Tags:** Add tags.
     c. Click **Save**.
        After saving the Stack, the system displays a banner with a link to the saved Stack.
     d. Click the link to open the Stack in the Resource Manager Service console. See, *Resource Manager and Terraform*.
   - **To view the details of a Stack:**
     a. Open the navigation menu. Under **Developer Services**, click **Resource Manager**.
     b. Click **Stacks**.
     c. Click the name of the Stack that you want to view details.
        Or, click the Actions menu (three dots), and select the **View stack details** option.
7. Click **Create VM Cluster**.

   The VM Cluster Details page is now displayed. While the creation process is running, the state of the VM cluster is **Pending**. When the VM cluster creation process completes, the state of the VM cluster changes to **Available**.

   The Exadata Database Storage section on the VM Cluster Details page shows the type of storage configured, which, in this case, is Exascale.

Related Topics

- Oracle Exadata Database Service on Cloud@Customer Service Description
  Learn how you can leverage the combined capabilities of Oracle Exadata and Oracle Cloud Infrastructure with Oracle Exadata Database Service on Cloud@Customer
- Using the Console to Scale the Resources on a VM Cluster
  Starting in Oracle Exadata Database Service on Cloud@Customer, you can scale up or down multiple resources at the same time. You can also scale up or down resources one at a time.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.
- Estimating How Much Local Storage You Can Provision On Your VMs
- Resource Tags
- Oracle PaaS/IaaS Cloud Service Description documents
- Oracle Platform as a Service and Infrastructure as a Service – Public Cloud Service DescriptionsMetered & Non-Metered
- Getting Started with Events
- Overview of Database Service Events
- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Resource Manager and Terraform

VM Cloud Automation Software Update Management

VM Cloud Automation software updates to Oracle-managed agents and tools on guest VMs ensure access to the latest Oracle Cloud features and support ongoing operational efficiency. Environments may depend on specific versions for custom scripts and workflows, or require updates to align with business cycles and critical operations.

With these capabilities, customers benefit from greater control and flexibility over their update process, gaining the ability to minimize risk by proactively testing updates and ensuring operational stability by protecting critical business windows from disruption. This approach supports the stability of business operations and the efficient adoption of Oracle Cloud enhancements.

Customers can control the scheduling and application of updates to VM Cloud Automation software in several ways:

- **Daily update scheduling:** Configure specific periods during the day when guest VMs poll for and apply updates, aligning software maintenance with preferred maintenance windows and operational requirements.
- **Phased rollout and early testing:** Assign specific clusters to receive updates first, enabling phased rollouts and allowing updates to be evaluated in non-production environments before wider deployment.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Freeze periods:** Define windows during which updates are paused to ensure uninterrupted performance during business-critical operations.

All updates are applied online, maintaining continuity for running guest VMs and their databases.

- What to Expect During VM Cloud Automation Software Updates
- Manage VM Cloud Automation Update Preferences
- Configure the freeze period
- Cancel Cloud Automation Update preference while the freeze period is in effect
- Re-enable Cloud Automation Update

What to Expect During VM Cloud Automation Software Updates

Handling Background Activities

When an update is scheduled, Oracle ensures that updates are applied only when the VM Cluster is not engaged in background activities that could be incompatible with the update process. If the VM Cluster is performing tasks such as configuration changes, software updates, backup jobs, or other background workflows during the configured update timeslot, the automation update will be deferred. Updates will be retried on the next available day during your defined timeslot, provided no conflicting activities are occurring at that time.

Update Delivery and Early Access

VM Cloud Automation software updates are delivered to your VM Cluster according to your configured schedule once an update is published by Oracle and becomes generally available. For customers who want to validate updates before deploying them more broadly, an option exists to enable early access for specific test clusters. To enable early access, set the early access flag in your cloud automation update configuration. A VM Cluster with early access enabled will receive and apply updates one week ahead of general availability, allowing you to thoroughly test and certify new updates in a non-production environment before production deployment.

This approach helps ensure your VM Clusters remain stable and reliable, while providing opportunities to proactively validate changes and minimize the risk of operational disruptions.

Manage VM Cloud Automation Update Preferences

Cloud Automation Update

Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster.

Set the start time for cloud automation updates.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable Early Access for Cloud Tools Update

VM clusters designated for early access receive updates 1-2 weeks before they are available to other clusters. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period

Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

> **Note**
>
> - The freeze period can extend for a maximum of 45 days from the start date.
> - Oracle automation will automatically apply updates with critical security fixes (CVSS >= 9) even during a configured freeze period.

Configure the freeze period

You can always update the freeze period to extend up to 45 days from the freeze period start date.

1. On the Exadata VM Cluster Details page, click **Actions** and then select **Manage cloud automation update**.
2. On the resulting Manage Cloud Automation Update page, move the **Configure freeze period** slider to enable.

   > **Note**
   >
   > The freeze period can be extended for a maximum of 45 days from the freeze period start date.

3. Click **Save**.

Cancel Cloud Automation Update preference while the freeze period is in effect

1. On the Exadata VM Cluster Details page, click **Actions** and then select **Manage cloud automation update**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

2. On the resulting Manage Cloud Automation Update page, move the **Configure freeze period** slider to disable.
3. Click **Save**.
4. On the Cancel freeze period dialog, click **Cancel freeze period** to confirm

   Oracle automation requires at least 7 days to apply any pending updates before a new freeze period can be configured. You can set up a new freeze period starting 7 days from the cancellation date of the previous freeze period.

   Cloud Automation Update will be applied during the time period when freeze period is not configured.

Re-enable Cloud Automation Update

1. On the Exadata VM Cluster Details page, click **Actions** and then select **Manage cloud automation update**.

   (or)

   Click the Cloud Automation Update **Edit** link in the Version section.
2. On the resulting Manage Cloud Automation Update page, move the **Configure freeze period** slider to enable.

> **Note**
>
> Once the freeze period expires (the end date has been reached) or the freeze period is canceled, Oracle automation requires 7 days to apply any pending updates before the freeze period is re-enabled again on the cluster.

3. Click **Save**.

Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection

You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

> **Note**
>
> - You are opting in with the understanding that the list of events, metrics, and log files collected can change in the future. You can opt-out of this feature at any time.
> - Oracle may add more metrics in the future, but if you have already chosen to collect metrics, you need not update your opt-in value. It will remain enabled/disabled based on your current preference.
> - If 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Exadata 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**.
6. Click **Edit**.

   Edit Diagnostics Collection Settings window is displayed.
7. Select or clear the checkboxes and then click **Save**.

Related Topics

- Overview of Automatic Diagnostic Collection
  By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- Incident Logs and Trace Files
  This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- Health Metrics
  Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Using the Console to Add VMs to a Provisioned Cluster

To add virtual machines to a provisioned cluster, use this procedure.

Once the VM cluster is upgraded to Exadata Database Service Guest VM OS 23.1, you will be able to add a new VM or a new database server to this VM cluster if Exadata Cloud@Customer Infrastructure is running an Exadata System Software version 22.1.16 and later.

> **Note**
>
> Upgrade to Exadata System Software 23.1 for Exadata Cloud@Customer Infrastructure will be available with February 2023 update cycle.

Consider reviewing the points below that will assist you in adding VMs to a provisioned cluster.

- The same Guest OS Image version running on the existing provisioned VMs in the cluster is used to provision new VMs added to extend the VM cluster. However, any customizations made to the Guest OS Image on the existing VMs must be manually applied to the newly added VM.
- For VM clusters running a Guest OS Image version older than a year, you must update the Guest OS Image version before adding a VM to extend the cluster.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- For databases not part of a Data Guard configuration, only databases that are running on all VMs in the existing cluster will be added to the newly provisioned VM. Any database running on a subset of VMs will not extend automatically to run on the newly added VM.

When you attempt to add a VM to a VM cluster, you might encounter the error `[FATAL] [INS-32156] Installer has detected that there are non-readable files in oracle home.` To resolve the issue, follow the steps outlined in *Adding a VM to a VM Cluster Fails* before you try adding a cluster node.

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

   VM Clusters is selected by default.

2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.

3. Click the name of a VM cluster where you want to add virtual machines.

4. In the VM Cluster Details page, under **Resources**, click **Virtual Machines**, and then click **Add Virtual Machines**.

5. On the **Add Virtual Machines** dialog, select additional DB servers on which to add the VM.

   You cannot unselect existing DB Servers. The maximum resources available per VM get updated based on the newly added DB servers.

   DB Server Statuses include **In this VM cluster**, **Network not configured**, **Hosting a VM**, **Available to add**, and **Insufficient resources**. You can only add DB servers with the **Available to add** selection eligibility.

   DB servers that do not have a network configured are not available to add. To configure the network, edit the VM Cluster Network of the associated infrastructure. For more information, see *Using the Console to Add Another DB Server to the VM Cluster Network*.

6. Select the DB servers with the **Available to add** selection eligibility and then click **Add**.

   The statuses of the DB servers change to Allocated.

> **Note**
>
> You cannot remove an allocated DB server.

To extend the database instance for Data Guard-enabled databases for the newly added VMs, see *Nodelist is not Updated for Data Guard-Enabled Databases*.

Related Topics

- Using the Console to Add Another DB Server to the VM Cluster Network
  To a DB server to the VM Cluster Network, complete this procedure.
- Adding a VM to a VM Cluster Fails
- Nodelist is not Updated for Data Guard-Enabled Databases

Using the Console to View a List of DB Servers on an Exadata Infrastructure

To view a list of database server hosts on an Oracle Exadata Cloud@Customer system, use this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Infrastructure**, click **Exadata Infrastructure**.
3. In the list of Exadata Infrastructures, click the display name of the infrastructure you wish to view details.
4. Under **Resources**, click **DB Servers**.
5. In the list of DB Servers, click the name of the DB Server that you wish to view details.

   DB Server lists VMs from each cluster hosted on them along with resources allocated to them.

Using the Console to Remove a VM from a VM Cluster

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

> **Note**
>
> Terminating a VM from a cluster requires the removal of any database which is part of a Data Guard configuration (either primary or standby) from the VM to proceed with the terminate flow. For more information on manual steps, see My Oracle Support note *2811352.1*.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to scale the CPU resources.
3. Click **Exadata VM Clusters**.
4. Click the name of the VM cluster for which you want to remove a virtual machine.
5. Under **Resources**, click **Virtual Machines**.
6. In the list of virtual machines, click the **Actions** icon (three dots) for a virtual machine, and then click **Terminate**.
7. On the Terminate Virtual Machine dialog, enter the name of the virtual machine, and then click **Terminate**.

   VM removed from the cluster. **VM Cluster Details** page displays the updated resource allocation details under **VM Cluster Resource Allocation**.

Related Topics

- https://support.oracle.com/rs?type=doc&id=2811352.1

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 **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to update the license type.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 Exadata Database Service on Cloud@Customer.

   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.

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

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Exadata VM Clusters**.
4. Click the name of the VM cluster that you want to add SSH key(s).
5. In the VM Cluster Details page, click **Add SSH Keys**.
6. In the ADD SSH Keys dialog, choose any one of the methods:

   - **Generate SSH key pair:** Select this option if you want the Control Plane to generate public/private key pairs for you.
     Click **Save Private Key** and **Save Public Key** to download and save SSH Key pair.
   - **Upload SSH key files:** Select this option to upload the file that contains SSH Key pair.
   - **Paste SSH keys:** Select this option to paste the SSH key string.
     To provide multiple keys, click **Another SSH Key**. For pasted keys, ensure that each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.
7. Click **Save Changes**.

Related Topics

- Managing Key Pairs on Linux Instances

Using the Console to Scale the Resources on a VM Cluster

Starting in Oracle Exadata Database Service on Cloud@Customer, you can scale up or down multiple resources at the same time. You can also scale up or down resources one at a time.

Scale down resources under the following circumstances:

- **Use Case 1:** If you have allocated all of the resources to one VM cluster, and if you want to create multiple VM clusters, then there wouldn't be any resources available to allocate to the new clusters. Therefore, scale down the resources as needed to then create additional VM clusters.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Use Case 2:** If you want to allocate different resources based on the workload, then scale down or scale up accordingly. For example, you may want to run nightly batch jobs for reporting/ETL and scale down the VM once the job is over.

You can scale down the following resources in any combinations:

- OCPU (ECPUs for X11M)
- Memory
- Local storage
- Exadata storage

Each scaling operation can take several minutes to complete. The time for each operation will vary based on activity in the system, but as a general rule, most operations should be completed within 15 minutes for a quarter rack, 20 minutes for a half rack, and 30 minutes for a full or larger rack. Performing multiple OCPU (ECPUs for X11M) scaling operations over a short period of time can lengthen the time for completion. Although online, OCPU (ECPUs for X11M) scaling is not implemented on all VMs in parallel so as to detect and protect from any anomalies before they affect the entire system. Memory and Local Storage scaling require a VM reboot and are performed one VM at a time in a rolling manner.

If you run multiple scale down operations, then each operation is performed serially. For example, if you scale memory and local storage from the Console, then the system will first scale memory, and when that operation completes, it will scale storage. The time to complete all operations will be the sum of the time to complete individual operations.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to scale the CPU resources.
3. Click **Exadata VM Clusters**.
4. Click the name of the VM cluster for which you want to scale the CPU resources.

   The VM Cluster Details page displays information about the selected VM cluster.
5. Click **Scale VM Resources**.
6. In the dialog box, adjust any or all of the following:

   - **OCPU (ECPUs for X11M) Count:**
     The OCPU (ECPUs for X11M) Count value must be a multiple of the number of virtual machines so that every virtual machine has the same number of CPU cores enabled.

     If you set the OCPU (ECPUs for X11M) Count to zero, then the VM cluster virtual machines are all shut down. If you change from a zero setting, then the VM cluster virtual machines are all started. Otherwise, modifying the number of enabled CPU cores is an online operation, and virtual machines are not rebooted because of this operation. See also *System Configuration*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you have explicitly set the `CPU_COUNT` database initialization parameter, that setting is not affected by modifying the number of CPU cores that are allocated to the VM cluster. Therefore, if you have enabled the Oracle Database instance caging feature, the database instance does not use extra CPU cores until you alter the `CPU_COUNT` setting. If `CPU_COUNT` is set to `0` (the default setting), then Oracle Database continuously monitors the number of CPUs reported by the operating system and uses the current count.

- **Memory:**
  Specify the memory for each individual VM. The value must be a multiple of 1 GB and is limited by the available memory on the Exadata infrastructure.

  When you scale up or down the memory, the associated virtual machines are rebooted in a rolling manner one virtual machine at a time to minimize the impact on the VM cluster.

- **Local file system size:**
  Specify the size for each individual VM. The value must be a multiple of 1 GB and is limited by the available size of the file system on the Exadata infrastructure.

  When you scale up or down the local file system size, the associated virtual machines are rebooted in a rolling manner one virtual machine at a time to minimize the impact on the VM cluster.

  a. Click **Show advanced options** to view additional local file systems configuration options.

  b. Resize the `/`, `/u01`, `/tmp`, `/var`, `/var/log`, `/var/log/audit`, and `/home` file systems as needed.

  > **Note**
  >
  > – You can only expand these file systems and cannot reduce the size once expanded.
  >
  > – Due to backup partitions and mirroring, the `/` and `/var` file systems will consume twice the space they were allocated, which is indicated in the read-only **Total allocated storage for / (GB) due to mirroring** and **Total allocated storage for /var (GB) due to mirroring** fields.

  c. After creating the VM Cluster, check the Exadata Resources section on the Exadata Infrastructure Details page to check the file size allocated to the local storage (`/u02`) and local storage (additional file systems).

  Reserved local storage per VM (GB): Displays the size reserved internally for root file systems, Oracle Grid Infrastructure Homes, and diagnostic logs.

- **Usable Exadata storage size:**
  Specify the total amount of Exadata storage that is allocated to the VM cluster. This storage is allocated evenly from all of the Exadata Storage Servers. The minimum recommended size is 2 TB.

  You may reduce the Exadata storage allocation for a VM cluster. However, you must ensure that the new amount covers the existing contents, and you should also allow for anticipated data growth.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> When you downsize, each VM must retain 15% free space (3,4 storage servers) or 9% free space (5 or more storage servers).

Modifying the Exadata storage allocated to the VM cluster is an online operation. Virtual machines are not rebooted because of this operation.

7. Click **Save**.

Related Topics

- System Configuration Options for Oracle Exadata Database Service on Cloud@Customer
  Oracle Exadata Database Service on Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X11M.

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

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

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
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 nodes, click the **Actions** icon (three dots) for a node, and then click one of the following actions:
   a. **Start:** Restarts a stopped node. After the node is restarted, the **Stop** action is enabled.
   b. **Stop:** Shuts down the node. After the node is stopped, the **Start** action is enabled.
   c. **Reboot:** Shuts down the node, and then restarts it.

Using the Console to Check the Status of a VM Cluster Virtual Machine

Review the health status of a VM cluster virtual machine.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
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 **Exadata VM Clusters**.
4. Click the name of the VM cluster that contains the virtual machine that you are interested in.

   The VM Cluster Details page displays information about the selected VM cluster.
5. In the **Resources** list, click **Virtual Machines**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The list of virtual machines displays. For each virtual machine in the VM cluster, the name, state, and client IP address are displayed.

6. In the node list, find the virtual machine that you are interested in and check its state.

   The color of the icon and the associated text it indicates its status.

   - **Available:** Green icon. The node is operational.
   - **Starting:** Yellow icon. The node is starting because of a start or reboot action in the Console or API.
   - **Stopping:** Yellow icon. The node is stopping because of a stop or reboot action in the Console or API.
   - **Stopped:** Yellow icon. The node is stopped.
   - **Failed:** Red icon. An error condition prevents the continued operation of the virtual machine.

Using the Console to Move a VM Cluster to Another Compartment

To change the compartment that contains your VM cluster on Oracle Exadata Database Service on Cloud@Customer, 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 **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to move.
3. Click **Exadata 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**, and then select **Move resource**.
6. In the resulting dialog, choose the new compartment for the VM cluster, and click **Move Resource**.

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 **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to terminate.
3. Click **Exadata VM Clusters**.
4. Click the name of the VM cluster that you want to terminate.

   The VM Cluster Details page displays information about the selected VM cluster.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. Click **More actions**, and then select **Terminate**.
6. In the resulting dialog, enter the name of the VM cluster, and click **Terminate** to confirm the action.

Using the API to Manage Oracle Exadata Database Service on Cloud@Customer VM Clusters

Review the list of API calls to manage your Oracle Exadata Database Service on Cloud@Customer VM cluster networks and VM clusters.

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 Cloud@Customer VM cluster networks and VM clusters:

VM cluster networks:

- `GenerateRecommendedVmClusterNetwork`
- `CreateVmClusterNetwork`
- `DeleteVmClusterNetwork`
- `GetVmClusterNetwork`
- `ListVmClusterNetworks`
- `UpdateVmClusterNetwork`
- `ValidateVmClusterNetwork`

VM clusters:

- `CreateVmCluster`
- `DeleteVmCluster`
- `GetVmCluster`
- `ListVmClusters`
- `UpdateVmCluster`

For the complete list of APIs, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- GenerateRecommendedVmClusterNetwork
- CreateVmClusterNetwork
- DeleteVmClusterNetwork
- GetVmClusterNetwork
- ListVmClusterNetworks
- UpdateVmClusterNetwork

Copyright © 2008, 2026, Oracle and/or its affiliates.

- ValidateVmClusterNetwork
- CreateVmCluster
- DeleteVmCluster
- GetVmCluster
- ListVmClusters
- UpdateVmCluster
- Database Service API

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**
>
> The use of the serial console feature requires Exadata Infrastructure version 22.1.10 or higher for 22.X users and version 23.1.1 or higher for 23.X users. The serial console feature will be available on any new VM Clusters created immediately but will only be available on previously existing VM Clusters after the next quarterly maintenance cycle. 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 for meeting 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 Database**, click **Exadata Database Service on Cloud@Customer**.
2. Under **Region**, select the region that you want to associate with the Oracle Exadata infrastructure.
3. Under **Infrastructure**, click **Exadata Infrastructure**.
4. Click the name of the infrastructure that you are interested in.
5. In the resulting Infrastructure Details page, go to the Version section to the find the DB Server version installed.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Required IAM Policies
  An administrator must grant you secure access to the virtual machine console on the Exadata Database Service on Cloud@Customer system through an IAM policy.
- Prerequisites
  You must install an SSH client and create SSH key pairs.
- Create the Virtual Machine Serial 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 Exadata Database Service on Cloud@Customer 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 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.

- Ports to Open for Control Plane Connectivity
- Install an SSH Client and a Command-line Shell (Microsoft Windows)
- Create SSH Key Pairs
- Sign in to a Virtual Machine From the Serial Console
- Connect Through Firewalls

Ports to Open for Control Plane Connectivity

Ensure that the firewall rules are correct so that the Control Plane Server (CPS) can reach the required OCI endpoints. For more information, see Table 3-2

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

- Git for Windows
- Windows Subsystem for Linux

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

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

Create SSH Key Pairs

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

- Create the SSH key Pair for Linux
- Create the SSH Key Pair for Windows Using PuTTY

Create the SSH key Pair for Linux

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

1. Open a shell or terminal for entering the commands.
2. At the prompt, enter `ssh-keygen` and provide a name for the key when prompted. Optionally, include a passphrase.

   The keys will be created with the default values: RSA keys of 2048 bits.

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

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

Argument	Description
`-t rsa`	Use the RSA algorithm.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Argument	Description
`-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.

> **Note**
>
> Ensure that you are using the latest version of PuTTY, see http://www.putty.org.

1. Find `puttygen.exe` in the PuTTY folder on your computer, for example, `C:\Program Files (x86)\PuTTY`. Double-click `puttygen.exe` to open it.
2. Specify a key type of SSH-2 RSA and a key size of 2048 bits:
   - In the **Key** menu, confirm that the default value of **SSH-2 RSA key** is selected.
   - For the **Type of key to generate**, accept the default key type of **RSA**.
   - Set the **Number of bits in a generated key** to 2048 if not already set.
3. Click **Generate**.
4. To generate random data in the key, move your mouse around the blank area in the PuTTY window.

   When the key is generated, it appears under **Public key for pasting into OpenSSH authorized_keys file**.
5. A **Key comment** is generated for you, including the date and timestamp. You can keep the default comment or replace it with your own more descriptive comment.
6. Leave the **Key passphrase** field blank.
7. Click **Save private key**, and then click **Yes** in the prompt about saving the key without a passphrase.

   The key pair is saved in the PuTTY Private Key (PPK) format, which is a proprietary format that works only with the PuTTY tool set.

Copyright © 2008, 2026, Oracle and/or its affiliates.

You can name the key anything you want, but use the `ppk` file extension. For example, `mykey.ppk`.

8. Select all of the generated key that appears under Public key for pasting into OpenSSH `authorized_keys` file, copy it using **Ctrl + C**, paste it into a text file, and then save the file in the same location as the private key.

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

You can name the key anything you want, but for consistency, use the same name as the private key and a file extension of `pub`. For example: `mykey.pub`.

9. Write down the names and location 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

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

For more information about generating SSH key pair, see Create the SSH Key Pair for Windows Using PuTTY
Do the following on the Create serial console access window:

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

Sign in to a Virtual Machine From the Serial Console

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

Connect Through Firewalls

Copyright © 2008, 2026, Oracle and/or its affiliates.

If the client you will use to access the serial console is behind a firewall, you must ensure that this client is able to reach the required endpoint in order 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.exacc.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.

> **Note**
>
> 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 **Exadata Database Service on Cloud@Customer**.
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 **Click Create console connection**.
5. In the resulting Create serial console access window, you have three options for adding the SSH key
   - **Generate a key pair for me**: You can have Oracle Cloud Infrastructure generate an SSH key pair to use. If you are using PowerShell or PuTTY to connect to the instance from a Windows client, you cannot use the generated SSH key pair without first converting it to a `.ppk` file.
   - **Upload public key file**: Browse to a public key file on your computer. If you followed the steps in *Creating SSH Key Pairs* in the Prerequisites section to create a key pair, use this option to navigate to the `.pub` file.
   - **Paste public key**: Paste the content of your public key file into the text box.
6. Click **Create**.

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

Related Topics

- Create SSH Key Pairs

Make an SSH Connection to the Serial Console

Copyright © 2008, 2026, Oracle and/or its affiliates.

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
- Connect from Mac OS X and Linux Operating Systems
- Connect from Windows Operating Systems

Validate Server Host Keys

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

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

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

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

Connect from Mac OS X and Linux Operating Systems

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

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

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 ProxyCommand, 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 username 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-C@C 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.

Related Topics

- Validate Server Host Keys
- Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems

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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you are connecting to the instance from a Windows client using PowerShell, `plink.exe` is required. `plink.exe` is the command link connection tool included with PuTTY. You can install PuTTY or install `plink.exe` separately. For more information, see *Installing an SSH Client and a Command-line Shell (Windows)*.

To connect to the serial console on Microsoft Windows

1. On the Virtual Machine details page in the Oracle Cloud Infrastructure Console, under **Resources**, click **Console connection**.
2. Click the Actions menu (three dots).

   Depending on which SSH client you are using, do one of the following:

   - If you are using Windows PowerShell, click **Copy serial console connection for Windows**.
   - If you are using OpenSSH, click **Copy serial console connection for Linux/Mac**.

   > **Note**
   >
   > The copied connection string for Windows contains the parameter `-i` specifying the location of the private key file. The default value for this parameter in the connection string references an environment variable that might not be configured on your Windows client, or it might not represent the location where the private key file is saved. Verify the value specified for the `-i` parameter and make any required changes before proceeding to the next step.

3. Paste the connection string copied from the previous step into a text file so that you can add the file path to the private key file.
4. In the text file, replace `$env:homedrive$env:homepath\oci\console.ppk` with the file path to the `.ppk` file on your computer. This file path appears twice in the string. Replace it in both locations.

   > **Note**
   >
   > For PuTTY versions 0.82 and above, add parameter `-legacy-stdio-prompts` to the first call to `plink`.

5. Paste the modified connection string into the PowerShell window or your OpenSSH client, and then press **Enter** to connect to the console.
6. If prompted, validate and accept the fingerprint of the server host key.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

8. Reboot your virtual machine.

   You do not need to enter a username 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-C@C VM Cluster Details page.

   b. Under **Resources**, click **Virtual Machines**.

   c. Select **Reboot** from the Actions menu (three dots) for the virtual machine that you want to reboot.

- To create a connection using the SSH key pair generated using the OCI Console

Related Topics

- Windows PowerShell
- Install an SSH Client and a Command-line Shell (Microsoft Windows)
- Validate Server Host Keys

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 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)

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

To convert a generated .key private key file

1. Open PuTTYgen.
2. Click **Load**, and select the private key generated when you created the instance.

   The extension for the key file is `.key`.
3. Click **Save private key**.
4. Specify a name for the key.

   The extension for the new private key is `.ppk`.
5. Click **Save**.

Using Cloud Shell to Connect to the Serial Console

You can connect to the serial console quickly and easily using the Cloud Shell integration. Cloud Shell is a web browser-based terminal accessible from the Console. The Cloud Shell integration automatically creates the instance console connection and a temporary SSH key. The only prerequisite for connecting to the serial console from Cloud Shell is granting users the correct permissions. For an introductory walkthrough of using Cloud Shell, see *Using Cloud Shell*.

> **Note**
>
> - By default, Cloud Shell limits network access to OCI internal resources in your tenancy home region only unless you have enabled the Cloud Shell managed Public Network. Your administrator must configure an Identity policy to enable Cloud Shell Public Network. For more information, see Cloud Shell Networking.
> - You cannot concurrently connect to more than one DB node using Cloud Shell. As an example, if you have an open connection to *DBnode1* and want to connect to *DBnode2*, you must first exit the active Cloud Shell from *DBnode1* and then establish a connection to *DBnode2*.
> - Ensure that the firewall rules are correct so that the Control Plane Server (CPS) can reach the required OCI endpoints. For more information, see Table 3-2

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Required IAM Policies
  An administrator must grant you secure access to the virtual machine console on the Exadata Database Service on Cloud@Customer system through an IAM policy.

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 **Exadata Database Service on Cloud@Customer**.
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

> **Note**
>
> To access the serial console and to use console history, firewall rules must be configured so that the Control Plane Server (CPS) can access the necessary OCI endpoints. Please review Table 3-2 details for Object Storage and VM console connectivity requirements.

You can capture and display recent serial console data for a Virtual Machine. The data includes configuration messages that occur when the Virtual Machine boots, such as kernel and BIOS messages, and is useful for checking the status of the Virtual Machine or diagnosing and troubleshooting problems.

The console history captures up to a megabyte of the most recent serial console data for the specified Virtual Machine. Note that the raw console data, including multi-byte characters, is captured.

The console history is a point-in-time record. To troubleshoot a malfunctioning Virtual Machine using an interactive console connection, use a serial console connection.

- Managing Console History Data

Managing Console History Data

You can use the Console or API to manage console history captures. Console history lets you see serial output from your Virtual Machine without having to connect to the instance remotely. The console history can be used to audit previous access and actions taken with the serial console.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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
  Review the list of API calls to manage console history data.

Using the Console to Capture the Console History

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
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 **Exadata Database Service on Cloud@Customer**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.

   Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. Click the name of the history that you're interested in.
6. On the resulting window, click **Download** to download a copy of the console history.

Using the Console to View Console History Captures

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 **Exadata Database Service on Cloud@Customer**.
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 **Exadata Database Service on Cloud@Customer**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.

   Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. On the console history list, for the console history capture that you want to view, click the **Actions** menu, and then click **Delete**.
6. On the confirmation dialog, click **Delete console history**.

Using the API to Manage the Console History Data

Review the list of API calls to manage console history data.

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

For the complete list of APIs, see Database Service API.

Use the following API operations to manage the console history data.

- To capture the console history, use the **createDbNodeConsoleHistory** method.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- 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
- To edit the system configuration files
- To add or reset the SSH key for the opc user
- To reset the password for the opc user

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 © 2008, 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.

Related Topics

- Default User Accounts for Oracle Exadata

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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> You can edit the file and remove your previous key if you want to. However, make sure to keep the cloud automation keys to prevent cloud automation from breaking.

```
echo '<contents of public key file>' >> authorized_keys
```

5. Restart the instance by running the following command:

   ```
   /usr/sbin/reboot -f
   ```

To reset the password for the opc user

1. From the Bash shell, run the following command to load the SElinux policies to preserve the context of the files you are modifying.

   This step is necessary to sign in to your instance using SSH and the Console.

   ```
   /usr/sbin/load_policy -i
   ```

2. Run the following command to remount the root partition with read/write permissions:

   ```
   /bin/mount -o remount, rw /
   ```

3. Run the following command to reset the password for the `opc` user:

   ```
   sudo passwd opc
   ```

4. Restart the instance by running the following command:

   ```
   sudo reboot -f
   ```

> **Note**
>
> Setting a `root` password would be an acceptable alternative to setting an `opc` password.

Exiting the Virtual Machine Serial Console Connection

- To exit the serial console connection
- To delete the serial console connection for a Virtual Machine

To exit the serial console connection

When using SSH, the ~ character at the beginning of a new line is used as an escape character.

Copyright © 2008, 2026, Oracle and/or its affiliates.

1. To exit the serial console, enter:

   ~.

2. To suspend the SSH session, enter:

   ~^z

   The ^ character represents the **CTRL** key.

3. To see all the SSH escape commands, enter:

   ~?

To delete the serial console connection for a Virtual Machine

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Click the VM Cluster that you're interested in.
3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.

   Under **Resources**, **Console connection** is selected by default.

4. Click the Actions menu, and then click **Delete**. Confirm when prompted.

Manage Software Images

- Creation and Storage of Software Images
  Software images are resources within your tenancy that you create before provisioning or updating a DB system, Exadata Cloud Infrastructure instance, Database Home, database, or Grid Infrastructure.
- Using a Software Image with an Oracle Exadata Database Service on Cloud@Customer System
  Create, save, and reuse a Software Image.
- Using the Console to View the List of Software Images
- Using the Console to Create a Database Software Image
  To create an Oracle Database Software Image with the console, use this procedure.
- Using the Console to Create a Grid Infrastructure Software Image
  To create a Grid Infrastructure Software Image with the console, use this procedure.
- Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Database Home or Oracle Grid Infrastructure Home
  OPatch utility enables you to apply the interim patches to Oracle Database Home or Oracle Grid Infrastructure Home. You can find the `opatch` utility in the `$ORACLE_HOME/Opatch` directory.
- Using the Console to Delete a Software Image
- Using the Console to View the Update Information of a Software Image
  To view the Oracle Database version, update information (RU level), and included one-off (interim) patches of a database software image, use the following instructions:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Move a Software Image to a Different Compartment
  Follow these steps to move a Software Image to a different compartment of your choice:
- Using the API for Managing Software Images
  Review the list of API calls to manage Software Images.

Creation and Storage of Software Images

Software images are resources within your tenancy that you create before provisioning or updating a DB system, Exadata Cloud Infrastructure instance, Database Home, database, or Grid Infrastructure.

There are two types of software image resources:

- **Grid Infrastructure software image:** Grid Infrastructure software images are resources containing Oracle Grid Infrastructure software used to update Oracle Grid Infrastructure. Grid Infrastructure software images are either Oracle-published software releases or custom software images created by the customer that include the desired Grid Infrastructure release updates (GIRU) and additional one-off (interim) patches.
- **Database software image:** Database software images are resources containing Oracle Database software used to provision and update Oracle Databases and Oracle Database Homes. Database software images are either Oracle-published software releases or custom software images created by the customer that include the desired Database release updates (DBRU) and additional one-off (interim) patches.

There is no limit on the number of software images you can create in your tenancy, and you can create your images with any Oracle Database software or Oracle Grid Infrastructure version and update supported in Oracle Cloud Infrastructure.

Software images are automatically stored in Oracle-managed Object Storage and can be viewed and managed in the Oracle Cloud Infrastructure Console. Software images are regional-level resources but they can be accessed from any region within your tenancy.

> **Note**
>
> The software images incur Object Storage usage costs.

Related Topics

- About Regions and Availability Domains
- Using the Console to Create a Database Software Image
  To create an Oracle Database Software Image with the console, use this procedure.

Using a Software Image with an Oracle Exadata Database Service on Cloud@Customer System

Create, save, and reuse a Software Image.

Creating a Software Image enables you to:

- Create custom Database and Grid Infrastructure images based on Software Images, RU, and one-off (interim) patches.
- Save a custom image automatically to Object Storage as a resource.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Provision an Oracle Database home or Oracle Database with the desired RU and one-off (interim) patches.
- Update the Database Home and Grid Infrastructure created using the Software Image.
- Clone Software Image to another service in the Data Guard creation process.

> **Note**
>
> The Software Images are created and managed by the customer and they are available for use until explicitly deleted.

Related Topics

- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to perform patching operations on Exadata database virtual machines and Database Homes by using the Console, API, or the CLI.
- Use Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer
  Learn to configure and manage Data Guard Groups in your VM cluster.

Using the Console to View the List of Software Images

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Software Images**.
   The resulting Software images page displays the list of custom software images, which includes details such as Image type (Database, Grid infrastructure), and version.

Using the Console to Create a Database Software Image

To create an Oracle Database Software Image with the console, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Cloud@Customer**.
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.
9. Optionally, you can enter a comma-delimited list of one-off (interim) patch numbers.

Copyright © 2008, 2026, Oracle and/or its affiliates.

10. Optionally, you can upload an Oracle Home inventory file from an existing Oracle Database. For more information, see Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Home.
11. Click **Show Advanced Options** to add **tags** to your database software image. 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 software image**.

Related Topics

- Release Update Introduction and FAQ (Doc ID 2285040.1)
- Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Database Home or Oracle Grid Infrastructure Home
  OPatch utility enables you to apply the interim patches to Oracle Database Home or Oracle Grid Infrastructure Home. You can find the `opatch` utility in the `$ORACLE_HOME/Opatch` directory.
- Resource Tags

Using the Console to Create a Grid Infrastructure Software Image

To create a Grid Infrastructure Software Image with the console, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Software Images**.
3. Click **Create Software Image**.
4. In the resulting Create software image page, click **Grid Infrastructure Software image**.
5. In the **Display name** field, provide a display name for your image. Avoid entering confidential information.
6. Choose your **Compartment**.
7. Choose a **Grid Infrastructure release**.
8. Choose the **Grid Infrastructure version** for your image.
9. Optionally, you can enter a comma-delimited list of one-off (interim) patch numbers.
10. Optionally, you can upload a Grid Infrastructure Home inventory file from an existing Oracle Grid Infrastructure. For more information, see Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Home.
11. Click **Show Advanced Options** to add **tags** to your database software image. 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 software image**.

Related Topics

- Release Update Introduction and FAQ (Doc ID 2285040.1)

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Database Home or Oracle Grid Infrastructure Home
  OPatch utility enables you to apply the interim patches to Oracle Database Home or Oracle Grid Infrastructure Home. You can find the `opatch` utility in the `$ORACLE_HOME/Opatch` directory.
- Resource Tags

Using the OPatch lsinventory Command to Verify the Patches Applied to an Oracle Database Home or Oracle Grid Infrastructure Home

OPatch utility enables you to apply the interim patches to Oracle Database Home or Oracle Grid Infrastructure Home. You can find the `opatch` utility in the `$ORACLE_HOME/Opatch` directory.

Using the `lsinventory` command provided by OPatch, you can create a file that lists the interim patches applied to an Oracle Database Home or Oracle Grid Infrastructure Home. This file can then be uploaded to the OCI Console during the creation of a custom software Image to add the exact set of patches used by the source Oracle Database Home or Oracle Grid Infrastructure Home to the list of patches included in the software image. You can find the `opatch` utility in the `$ORACLE_HOME/Opatch` directory. The following example shows how to use the `lsinventory` command to create the `lsinventory` file:

1. Run the `opatch lsinventory` command to get the list of interim patches applied.

```
$ORACLE_HOME/OPatch/opatch lsinventory
Oracle Interim Patch Installer version 12.2.0.1.21
Copyright (c) 2021, Oracle Corporation. All rights reserved.

Oracle Home : /u02/app/oracle/product/19.0.0.0/dbhome_2
Central Inventory : /u01/app/oraInventory
from : /u02/app/oracle/product/19.0.0.0/dbhome_2/oraInst.loc
OPatch version : 12.2.0.1.21
OUI version : 12.2.0.7.0
Log file location : /u02/app/oracle/product/19.0.0.0/dbhome_2/cfgtoollogs/
opatch/opatch2021-01-21_09-22-45AM_1.log

Lsinventory Output file location : /u02/app/oracle/product/19.0.0.0/
dbhome_2/cfgtoollogs/opatch/lsinv/lsinventory2021-01-21_09-22-45AM.txt


Oracle Interim Patch Installer version 12.2.0.1.41
Copyright (c) 2024, Oracle Corporation.  All rights reserved.


Oracle Home       : /u01/app/oracle/product/19.0.0.0/gridhome_1
Central Inventory : /u01/app/oraInventory
   from           : /u01/app/oracle/product/19.0.0.0/gridhome_1/oraInst.loc
OPatch version    : 12.2.0.1.41
OUI version       : 12.2.0.7.0
Log file location : /u01/app/oracle/product/19.0.0.0/gridhome_1/
cfgtoollogs/opatch/opatch2024-04-19_19-24-22PM_1.log

Lsinventory Output file location : /u01/app/oracle/product/19.0.0.0/
gridhome_1/cfgtoollogs/opatch/lsinv/lsinventory2024-04-19_19-24-22PM.txt
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

2. Use the `lsinventory` output file to extract the additional interim patches applied to a specific Oracle Home.

Using the Console to Delete a Software Image

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Software images**.
3. In the list of software images, find the image you want to delete and click the action 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 Console to View the Update Information of a Software Image

To view the Oracle Database version, update information (RU level), and included one-off (interim) patches of a database software image, use the following instructions:

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Software images**.
3. In the list of software images, find the image you want to view and click on the display name of the image.
4. On the Software Image details page for your selected image, details about the image are displayed:
   - The **General Information** section includes details such as Image type.
   - The **Patch Information** section includes details such as release, version, and available interim patches.
   - The **One-Off Patches** field displays the number of one-off (interim) patches included in the image, if any. The count includes all patches specified when creating the image (including patches listed in `lsinventory`).
     To view all the patches, click **View all items**. The resulting View all One-Off patches window displays the list of all one-off patches (interim). Click **Close** to close the window.

Using the Console to Move a Software Image to a Different Compartment

Follow these steps to move a Software Image to a different compartment of your choice:

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Cloud@Customer**.
2. Under **Resources**, click **Software images**.
3. In the list of software images, find the image you want to move and click the action icon (three dots) at the end of the row.
4. Click **Move resource**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. In the resulting Move resource to a different compartment dialog, choose a compartment.
6. Click **Move resource**

Using the API for Managing Software Images

Review the list of API calls to manage 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".

Use these API operations to manage Database Software Images:

- `CreateDatabaseSoftwareImage`
- `ListDatabaseSoftwareImages`
- `GetDatabaseSoftwareImage`
- `DeleteDatabaseSoftwareImage`
- `ChangeDatabaseSoftwareImageCompartment`

For the complete list of APIs, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- CreateDatabaseSoftwareImage
- ListDatabaseSoftwareImages
- GetDatabaseSoftwareImage
- DeleteDatabaseSoftwareImage
- ChangeDatabaseSoftwareImageCompartment
- Database Service API

Create Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer System

Learn to create Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer.

- About Creating Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer 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.
- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the API to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home, review the list of API calls.

About Creating Oracle Database Homes on an Oracle Exadata Database Service on Cloud@Customer 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 with Oracle Exadata Database Service on Cloud@Customer
  Learn to use the `dbaascli` utility on Oracle Exadata Database Service on Cloud@Customer.

Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer

To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

> **Note**
>
> Minimum requirements for provisioning a Database 26ai home:
>
> - Grid Infrastructure 26ai
> - Exadata Guest VM running Exadata System Software 23.1.8

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

   VM Clusters is selected by default.
2. Choose your Compartment.

   A list of VM Clusters is displayed for the chosen Compartment.

Copyright © 2008, 2026, Oracle and/or its affiliates.

3. In the list of VM clusters, click the VM cluster on which you want to create the Database Home.
4. Under **Resources**, click **Database Homes**.
5. Click **Create Database Home**.
6. In the Create Database Home dialog, enter the following:
   - **Database Home display name**: The display name for the Database Home.
   - **Unified Auditing:** Select this check box to enable Unified Auditing framework.

     > **Note**
     >
     > You cannot disable Unified Auditing after provisioning the Database Home.

     – **For Oracle Database versions lower than 12.1:** You cannot use the Unified Auditing framework and should instead use the Traditional Audit, the legacy Oracle Database audit framework.

     – **For Oracle Database versions 12.1 or higher:** You can enable Unified Auditing from the OCI Console. For Oracle Database versions 12.1 or higher but lower than version 26ai, the **Unified Auditing** check box is not selected by default. However, it is selected by default for Oracle AI Database 26ai.
     Unified Auditing field in the **General Information** section on the **Database Home Details** page displays if Unified Auditing is **Enabled** or **Disabled**.

   - **Database image**: Determines what Oracle Database version is used for the database. You can mix database versions on the Exadata VM Cluster, but not editions. By default, the latest Oracle-published database software image is selected.
     Click **Change Database Image** to use an older Oracle-published image or a custom Database Software Image that you have created in advance, then choose one of the options below on the **Select a Database Software Image** window.

     Oracle Provided Database Software Images: These images contain generally available versions of Oracle Database software.

     Custom Database Software Images: These images are created by your organization and contain customized configurations of software updates and patches.

     Use the **Select a compartment** and **Select a Database version** selectors to limit the list of custom Database Software Images to a specific compartment or Oracle Database software major release version.

     Use the **Region** filter to access the software images created in a different region.

     After choosing a software image, click **Select** to return to the Create Database Home dialog.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

A Database Software Image will not be available for Database Home creation if:

- The database version of Database Software Image is out of support. For example, Database Software Images created using 11.2.0.4 will not be available for Database Home provisioning after 31-Dec-2022.
- The Exadata model should support the PSU/RU version of the Database Software Image. For example, for the 19c release, the X8M-2 model supports RU version 19.4 and greater.
- Only Database Software Images created specifically in the context of Exadata Cloud@Customer service can be used while provisioning and patching Database Homes within the Exadata Cloud@Customer service.
- The Database Software Image is not in `Available` state, that is, `Deleted` or is being `Updated`.

- **Show Advanced Options**

  You have the option to configure advanced options.

  - **Tags:** (Optional) You can choose to apply 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 if you should apply tags, then skip this option (you can apply tags later) or ask your administrator.

Note that after Home install, patch to the latest if the latest patch is available.

7. Click **Create**.

   When the Database Home creation is complete, the status changes from `Provisioning` to `Available`.

Related Topics

- Resource Tags

Using the API to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer

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 Cloud@Customer, 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

Copyright © 2008, 2026, Oracle and/or its affiliates.

Manage Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems

Learn to manage Oracle Database homes on Oracle Exadata Database Service on Cloud@Customer.

- About Managing Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems
  You can delete or view information about Oracle Database Homes (referred to as **Database Homes** in Oracle Cloud Infrastructure) by using the Oracle Cloud Infrastructure Console, the API, or the CLI.
- Manage Database Home Using the Console
  Use the OCI console to manage the various operations needed on a Database Home.
- Using the API to Manage Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  Review the list of API calls to manage Oracle Database home.
- Differences Between Managing Resources with dbaascli and the Database API
  Learn how Oracle Oracle Exadata Database Service on Cloud@Customer automatically synchronizes `dbaascli` utility parameters and Database Service API properties of Oracle Database instance and Oracle Database Home.

About Managing Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems

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.

To find out how to delete or view information about Database Homes manually, See "Using the `dbaascli` Utility on Oracle Exadata Database Service on Cloud@Customer."

> **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.

Related Topics

- Using the dbaascli Utility with Oracle Exadata Database Service on Cloud@Customer
  Learn to use the `dbaascli` utility on Oracle Exadata Database Service on Cloud@Customer.

Manage Database Home Using the Console

Use the OCI console to manage the various operations needed on a Database Home.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to View Information About an Oracle Database Home
  To view the configuration details of an Oracle Database home, use this procedure.
- Using the Console to Delete an Oracle Database Home
  To delete an Oracle Database home with the Console, use this procedure.

Using the Console to View Information About an Oracle Database Home

To view the configuration details of an Oracle Database home, use this procedure.

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

   VM Clusters is selected by default.
2. Click **Compartment**, and select your compartment.

   A list of VM Clusters is displayed for the compartment you selected.
3. In the list of VM clusters, click the VM cluster that contains the Database Home in which you are interested.
4. Under **Resources**, click **Database Homes**.
5. In the list of Database Homes, find the Database Home you want to view, and then click the Database Home name to display details about it.

Using the Console to Delete an Oracle Database Home

To delete an Oracle Database home with the Console, use this procedure.

You cannot delete an Oracle Database home that contains databases. Before you can delete an Oracle Database home, you must first terminate the databases to empty the Database Home. You can terminate a database by using the Console.

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

   VM Clusters is selected by default.
2. Select **Compartment**, and choose the Compartment that you want to view.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the VM cluster that contains the Database Home that you want to delete.
4. Under **Resources**, click **Database Homes**.
5. In the list of Database Homes, find the Database Home that you want to delete, and click thje Database Home name to display details about it.
6. On the Database Home Details page, click **Delete**.

Related Topics

- Using the Console to Terminate a Database
  You can terminate a database and thereby remove the terminated database from the Cloud Control Plane.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the API to Manage Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer

Review the list of API calls to manage Oracle Database home.

For information about using the API and signing requests, see "REST APIs" and "Security Credentials". For information about SDKs, see "Software Development Kits and Command Line Interface".

Use these API operations to manage Database Homes:

- `ListDbHomes`
- `GetDbHome`
- `DeleteDbHome`

For the complete list of APIs, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- ListDbHomes
- GetDbHome
- DeleteDbHome
- Database Service API

Differences Between Managing Resources with dbaascli and the Database API

Learn how Oracle Oracle Exadata Database Service on Cloud@Customer automatically synchronizes `dbaascli` utility parameters and Database Service API properties of Oracle Database instance and Oracle Database Home.

You can manage resources using host tooling such as `dbaascli` or `dbaasapi`, or tools based on the Database service API (including the Oracle Cloud Infrastructure Console, SDKs, and the API itself). For the management operations discussed in this topic, if you perform them using `dbaascli` or `dbaasapi`, then the updates are not visible in tools based on the Database Service API until the next synchronization operation, which happens every 10 minutes.

The following table lists operations that you perform using `dbaascli` or `dbaasapi`, and maps these operations to the related Database Service API parameters:

Operation	Host Tooling Parameters (dbaascli / dbaasapi)	Database Service API Parameters
Creating Database Home	`name`, `DatabaseVersion`, `dbHomeLocation`, `createTime`	`displayName`, `dbVersion`, `dbHomeLocation`, `timeCreated`
Updating Database Home	`DatabaseVersion`, `dbHomeLocation`	`dbVersion`, `dbHomeLocation`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	Host Tooling Parameters (dbaascli / dbaasapi)	Database Service API Parameters
Creating Database	`dbName`, `dbUniqueName`, `pdbName`, `characterSet`, `NlsCharacterSet`, `dbClass`, `createTime`	`dbName`, `dbUniqueName`, `pdbName`, `characterSet`, `ncharacterSet`, `dbType`, `timeCreated`
Updating Database	`dbHomeId`, `dbUniqueName`, `dbClass`	`dbHomeId`, `dbUniqueName`, `dbType`

If you terminate a Database Home using the `dbaascli` or `dbaasapi`, the status of the Database Home is displayed as `Terminated` in the Database Service REST API based tools. If you terminate a database, the status of the database is displayed as `Failed`.

Manage Databases on Oracle Exadata Database Service on Cloud@Customer

- Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer
  Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer.
- Oracle Database Releases Supported by Oracle Exadata Database Service on Cloud@Customer
  Learn about the versions of Oracle Database that Oracle Exadata Database Service on Cloud@Customer supports.
- About Provisioning and Configuring Oracle Databases on Oracle Exadata Database Service on Cloud@Customer
  Learn about provisioning and configuring Oracle Database on Oracle Exadata Database Service on Cloud@Customer
- Using the Console to Manage Databases on Oracle Exadata Database Service on Cloud@Customer
  To create or terminate a database, complete procedures using the Oracle Exadata console.
- Using the API to Manage Oracle Database Components
  Use various API features to help manage your databases on Oracle Exadata Database Service on Cloud@Customer.
- Changing the Database Passwords
  To change the SYS password, or to change the TDE wallet password, use this procedure.
- Manage Pluggable Databases on Oracle Exadata Database Service on Cloud@Customer
  Learn to manage pluggable databases on Oracle Exadata Database Service on Cloud@Customer.
- Connect to an Oracle Database using Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users
  You can configure Oracle Exadata Database Service on Cloud@Customer 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Authenticating and Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Cloud@Customer
  An Oracle Database can be configured for Microsoft Azure users of Microsoft Entra ID to connect using single sign-on authentication.

Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer

Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer.

Before you can create and use an Oracle Database on Exadata Database Service on Cloud@Customer, you must:

- Provision Exadata Database Service on Cloud@Customer infrastructure
- Configure a VM cluster
- Create any required backup destinations

You can create one or more databases on each Oracle Exadata Database Service on Cloud@Customer system. Other than the storage and processing limits of your Oracle Exadata system, there is no maximum for the number of databases that you can create. By default, databases on Exadata Database Service on Cloud@Customer use Oracle Database Enterprise Edition - Extreme Performance. This edition provides all the features of Oracle Database Enterprise Edition, plus all of the database enterprise management packs, and all of the Enterprise Edition options, such as Oracle Database In-Memory, and Oracle Real Application Clusters (Oracle RAC). If you use your own Oracle Database licenses, then your ability to use various features is limited by your license holdings. TDE Encryption is required for all cloud databases. All new tablespaces will automatically be enabled for encryption.

> **Note**
>
> Custom locale objects (such as language, territory, character set, and collation) require deploying custom locale data files to both the database and storage servers, which is not supported in Exadata Database Service on Cloud@Customer.

Oracle Database Releases Supported by Oracle Exadata Database Service on Cloud@Customer

Learn about the versions of Oracle Database that Oracle Exadata Database Service on Cloud@Customer supports.

Exadata Database Service on Cloud@Customer supports the following Oracle Database software releases:

- Oracle AI Database 26ai
- Oracle Database 19c (19.x)
- Oracle Database 12c Release 2 (12.2.0.1) (requires a valid Upgrade Support contract)
- Oracle Database 12c Release 1 (12.1.0.2) (requires a valid Upgrade Support contract)
- Oracle Database 11g Release 2 (11.2.0.4) (requires a valid Upgrade Support contract)

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

Related Topics

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

About Provisioning and Configuring Oracle Databases on Oracle Exadata Database Service on Cloud@Customer

Learn about provisioning and configuring Oracle Database on Oracle Exadata Database Service on Cloud@Customer

Each Oracle Database is configured as follows:

- When you provision a database, you can associate it with a backup destination, and enable automatic backups.
- When a database is provisioned an archivelog maintenance job is added to the `crontab` for the database.
  - If the database is not enabled for backups, then the archivelog job will maintain FRA space by deleting Archive Redo Logs older than 24 hours.
  - If the database is enabled for backups, then the archivelog job will backup archivelogs that have not been backed up. Once an archived log is backed up, it will be purged when older than 24 hours.
- Each database is configured with Oracle Real Application Clusters (Oracle RAC) database instances running on every node in the virtual machine (VM) cluster.
- Each database is created in an Oracle home, which uses a separate set of Oracle binaries in a separate Oracle home location.
- Each database is configured with default instance parameter settings. While the defaults are reasonable for many cases, you should review the instance parameter settings to ensure that they meet your specific application needs.
  In particular, review the Oracle Database system global area (SGA) and program global area (PGA) instance parameter settings, especially if your VM cluster supports multiple databases. Also, ensure that the sum of all Oracle Database memory allocations never exceeds the available physical memory on each virtual machine.
  - 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.
  - The database initialization parameter `INMEMORY_FORCE` is set to `CELLMEMORY_LEVEL` for all 19.8 and later databases created via the cloud automation. This setting will enable the Exadata Columnar Cache feature, which dramatically speeds up analytic queries. It is available for 19.8 and later databases and no In Memory license is required when running in Exadata Cloud. For more information, see INMEMORY_FORCE
- Exadata Database Service will only create databases with 8K block size. This parameter cannot be changed.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Each database using Oracle Database 12c Release 1 or a later release is configured as a container database (CDB). One pluggable database (PDB) is created inside the CDB. By default:
  - The first PDB is configured with a local PDB administration user account, named `PDBADMIN`.
  - The `PDBADMIN` user account is initially configured with the same administration password as the `CDB SYS` and `SYSTEM` users.
  - The `PDBADMIN` user account is initially configured with basic privileges assigned through two roles; `CONNECT` and `PDB_DBA`. However, for most practical administrative purposes you must assign extra privileges to the `PDBADMIN` user account, or to the `PDB_DBA` role.

  You can use native Oracle Database facilities to create extra PDBs, and to manage all of your PDBs. The `dbaascli` utility also provides a range of convenient PDB management functions, including the ability to create nonCDB databases (release 19.x and earlier).

> **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.

Using the Console to Manage Databases on Oracle Exadata Database Service on Cloud@Customer

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

- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.
- Using the Console to Manage SYS User and TDE Wallet Passwords
  Learn to manage administrator (SYS user) and TDE wallet passwords.
- Using the Console to Migrate Encryption Keys from Oracle Wallet to Other Encryption Methods (Oracle Key Vault, External Keystore)
  Learn to migrate encryption keys between different encryption methods.
- Using the Console to Rotate Encryption Keys
  Learn how to rotate encryption keys. Note that key rotation is only available for databases that do not use Oracle Wallet as the encryption method.
- Using the Console to Move a Database to Another Database Home
  Learn to move a database to another Database Home.
- To run datapatch on a container database (CDB)
  To run datapatch on a container database, use this procedure.
- To run datapatch on a pluggable database (PDB)
  To run datapatch on a pluggable database, use this procedure.
- Using the Console to Terminate a Database
  You can terminate a database and thereby remove the terminated database from the Cloud Control Plane.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Create a Database

To create an Oracle Database with the console, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
   VM Clusters is selected by default.
2. Choose your **Compartment**.
   A list of VM Clusters is displayed for the chosen Compartment.
3. Click the name of a VM cluster where you want to create the database.
   In the VM Cluster Details page, under **Resources**, **Databases** is selected by default.
4. Click **Create Database**.
   (or)
   a. Click **Database Homes**.
   b. Click the name of the Database Home where you want to create the database.
   c. Click **Create Database**.
5. Provide the requested information in the Create Database page:

> **Note**
>
> You cannot modify the `db_name`, `db_unique_name`, and SID prefix after creating the database.

- **Provide the database name**: Specify a user-friendly name that you can use to identify the database. The database name must contain only the permitted characters. Review the following guidelines when selecting a database name.
  - maximum of 8 characters
  - contain only alphanumeric characters
  - begin with an alphabetic character
  - cannot be part of first 8 characters of a `db_unique_name` on the VM cluster
  - unique within a VM cluster
  - **DO NOT** use `grid` because `grid` is a reserved name
  - **DO NOT** use `ASM` because `ASM` is a reserved name
- **Provide a unique name for the database**: Optionally, specify a unique name for the database. This attribute defines the value of the `db_unique_name` database parameter. The value is case insensitive.

  The `db_unique_name` must contain only the permitted characters. Review the following guidelines when selecting a database name.
  - maximum of 30 characters
  - can contain alphanumeric and underscore (_) characters
  - begin with an alphabetic character
  - unique across the fleet/tenancy

Copyright © 2008, 2026, Oracle and/or its affiliates.

If a unique name is not provided, then the `db_unique_name` defaults to the following format *<db_name>_<3 char unique string>_<region-name>*.

If you plan to configure the database for backup to a Recovery Appliance backup destination, then the unique database name must match the name that is configured in the Recovery Appliance.

- **Select a database version**: From the list, choose the Oracle Database software release that you want to deploy.
- **Provide the name of the first PDB**: (Optional) Specify the name for the first PDB. A PDB is created with the database.
  To avoid potential service name collisions when using Oracle Net Services to connect to the PDB, ensure that the PDB name is unique across the entire VM cluster. If you do not provide the name of the first PDB, then a system-generated name is used.
- **Database Home**: Select an existing Database Home or create one as applicable. Note that this field is not available when you create a Database from the Database Home details page.
  - **Select an existing Database Home**: If one or more Database Homes already exist for the database version you have selected, then this option is selected by default. And, you will be presented with a list of Database Homes. Select a Database Home from the list.
  - **Create a new Database Home**: If no Database Homes exist for the database version you have selected, then this option is selected by default.
    a. Enter **Database Home display name**.
    b. Click **Change Database Image** to select your software version.
    Select a **Database Software Image** window is displayed.
    c. Select an **Image Type**, **Oracle Provided Database Software Images**, or **Custom Database Software Images**.
    If you choose **Oracle Provided Database Software Images**, then you can use the **Display all available version** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

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

- **Provide the administration password**: Provide and confirm the Oracle Database administration password. This password is used for administration accounts and functions in the database, including:
  - The password for the Oracle Database `SYS` and `SYSTEM` users.
  - The Transparent Data Encryption (TDE) Keystore password.

  For Oracle Database 12c Release 1 or later releases, the password for the PDB administration user in the first PDB (`PDBADMIN`) must be nine to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. In addition, the password must

Copyright © 2008, 2026, Oracle and/or its affiliates.

not contain the name of the tenancy or any reserved words, such as `Oracle` or `Table`, regardless of casing.

  - **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.

- **Backup Destination Type:** Select a backup destination for the database. From the list, choose an option:
  - **Object Storage**: Select to store backups in an Oracle-managed object storage container on Oracle Cloud Infrastructure.
    To use this option, your Oracle Exadata Cloud@Customer system must have egress connectivity to Oracle Cloud Infrastructure Object Storage.
  - **NFS**: Select to store backups in one of your previously defined backup destinations that use Network File System (NFS) storage. For more information, refer to the information about backup destinations in this publication.
    If you select this option, then you must also choose from the list of NFS **Backup Destinations**.
  - **Recovery Appliance**: Select to store backups in one of your previously defined backup destinations that use Oracle Zero Data Loss Recovery Appliance. Refer to the information about backup destination options in this document.
    If you select Oracle Zero Data Loss Recovery Appliance as your backup option, then you must also:
    * Choose from the list of appliance **Backup Destinations**.
    * Choose from the **VPC User** list, which contains the list of virtual private catalog (VPC) user names that are defined in the Oracle Zero Data Loss Recovery Appliance backup destination.
    * Provide the **Password** for the VPC user.

    > **Note**
    >
    > If you select a backup destination, then you cannot change a backup location after the database is created.

  - **Enable automatic backups:** Specify the settings for backing up the database to Object Storage, Network File System, or Zero Data Loss Recovery Appliance (ZDLRA).

    > **Note**
    >
    > You can change this setting after database creation.

    * **Object Storage**
      * **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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

        - **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.
      - **Backup Retention Period:** You can choose one of the following preset retention periods.
        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.
      - **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 as per backup retention period:** Select this option if you want to retain database backups for the entire period defined in the Backup retention period 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.
    - **Network File System (NFS)**
      - **Compartment:** Choose the compartment where the NFS backup destination exists.
      - **Backup Destinations:** Choose from the list of NFS backup destinations.
      - **NFS (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.
        - **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.
      - **NFS (L1):** Incremental backup scheduling time (UTC): Specify the time window when the incremental backups start.
      - **Backup Retention Period:** You can choose one of the following preset retention periods.

Copyright © 2008, 2026, Oracle and/or its affiliates.

NFS 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.

* **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 as per backup retention period:** Select this option if you want to retain database backups for the entire period defined in the Backup retention period 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.

* **Recovery Appliance**
  * **Compartment:** Choose the compartment where the Recovery Appliance backup destination exists.
  * **Backup Destinations:** Choose from the list of Recovery Appliance backup destinations.
  * **VPC User:** Provide a VPC user name and password for connecting to the Zero Data Loss Recovery Appliance.
  * **Recovery Appliance (L0):**
    * **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.
  * **Recovery Apliance (L1):** Incremental backup scheduling time (UTC): Specify the time window when the incremental backups start.
  * **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 Backup retention period 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.

- **Key Management**: Choose an encryption option, **Oracle Wallet**, **Oracle Key Vault**, or **External HSM**. 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

To set the TDE wallet password manually, uncheck this option and enter the **TDE wallet password** and confirm it.

- **Oracle Key Vault**: Select the compartment where you have created the Key Store, and then select the Key Store. As part of the CDB creation, a new wallet is created for the CDB in Oracle Key Vault (OKV). Also, a TDE Master Key is generated for the CDB and added to the wallet in OKV.

> **Note**
>
> * Container Databases (CDB) and pluggable databases (PDB) only support 256-bit Hardware Security Module (HSM) Vault keys.
> * Validate OKV Key encryption post restart: OKV TDE Maser Key is validated every time you start or restart your CBD.
> * Start or restart fails if the key is not validated. Work requests and life cycle states indicate the reason for failure.
> * View OKV keys post database restore: When you restore a CDB, the master key associated with that backup is restored as well.
> * Enable CDB backups to capture wallet name: CDB backups information about the wallet associated with the backup.
> * OKV Wallet or TDE Master Key on CDB deletion: If you delete a CDB, then the wallet and TDE Master Key remain in OKV and will not be deleted.

- **External Keystore:**
  - **TDE wallet password:** Enter the TDE wallet password.
  - **External Keystore credential:** Enter user credentials in the format your external keystore provider supports.

- (Optional) Select **Show Advanced Options**. From this window, you can select the following options:
  - **Provide the Oracle SID prefix**:

> **Note**
>
> Entering a SID prefix is only available for 12.1 databases and above.

Optionally, specify the Oracle SID prefix for the database. The instance number is automatically appended to the SID prefix to become the `instance_name` database parameter. If not provided, then the SID prefix defaults to the `db_name`.

Review the following guidelines when selecting a database name:

* maximum of 12 characters
* contain only alphanumeric characters
* begin with an alphabetic character
* unique in the VM cluster

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Real-time data protection: Real-time protection is the continuous transfer of redo changes from a protected database to ZDLRA. This reduces data loss and provides a recovery point objective (RPO) near 0.
  Once Real-time data protection is enabled, you can see the status in the Backup section of the Database details page.
- **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`.
- **Tags**: (Optional) You can choose to apply tags. If you have permissions to create a resource, 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, refer to information about 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.

6. Click **Create**.

> **Note**
>
> You can now:
>
> - Create or delete a CDB while a Data Guard setup is running on another database within the same Oracle home, and vice versa.
> - Create or delete a CDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.
> - Create or delete a CDB while concurrently creating or deleting a PDB within the same Oracle home, and vice versa.
> - Create or delete a CDB concurrently on different databases within the same Oracle home.
> - Create or delete a CDB while simultaneously updating VM Cluster tags.

Related Topics

- Resource Tags
- Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer
  Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Cloud@Customer.
- Manage Keys Using an External Keystore
  Review the use cases and implementation details of an external keystore.

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 **Exadata Database Service on Cloud@Customer**
2. Choose your **Compartment** that contains the VM cluster that hosts the database that you want to change passwords.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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.

Using the Console to Migrate Encryption Keys from Oracle Wallet to Other Encryption Methods (Oracle Key Vault, External Keystore)

Learn to migrate encryption keys between different encryption methods.

> **Note**
>
> - Migration requires the credentials for the encryption key, which must already be present on the external keystore.
> - Key migration works seamlessly in a Data Guard environment.
> - You cannot migrate keys if other databases in the same VM cluster are already using different encryption methods, such as an external keystore or Oracle Key Vault (OKV).
> - You will encounter an issue during PDB creation or cloning operations on databases migrated to an external keystore. The fix for this issue will be included in the January 2025 Release Update (RU) for database versions 19c and 26ai. Until the RU becomes available, Oracle recommends applying the one-off patch 36930984 to address the issue.

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

   VM Clusters is selected by default.
2. Select the compartment containing the VM cluster that hosts the database for which you are migrating encryption keys.

Copyright © 2008, 2026, Oracle and/or its affiliates.

3. Click the name of the VM cluster that contains the database that you want to migrate encryption keys.
4. In the Resources list of the VM Cluster Details page, click **Databases**.
5. Click the name of the database that you want to migrate encryption keys to a difference encryption method.

   The Database Details page displays information about the selected database.
6. In the **Encryption** section, click the **Migrate key** link.
7. On the resulting Migrate Key page, change the key encryption method and click **Save changes**.

Related Topics

- Manage Keys Using an External Keystore
  Review the use cases and implementation details of an external keystore.

Using the Console to Rotate Encryption Keys

Learn how to rotate encryption keys. Note that key rotation is only available for databases that do not use Oracle Wallet as the encryption method.

> **Note**
>
> The "Rotate key" functionality applies only to databases that are not using Oracle Wallet.

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

   VM Clusters is selected by default.
2. Select the compartment containing the VM cluster that hosts the database for which you are rotating encryption keys.
3. Click the name of the VM cluster that contains the database that you want to rotate encryption keys.
4. In the Resources list of the VM Cluster Details page, click **Databases**.
5. Click the name of the database that you want to rotate encryption keys.

   The Database Details page displays information about the selected database.
6. In the **Encryption** section, click the **Rotate key** link.
7. On the resulting Rotate Key dialog, click **Rotate**.

Related Topics

- Manage Keys Using an External Keystore
  Review the use cases and implementation details of an external keystore.

Using the Console to Move a Database to Another Database Home

Learn to move a database to another Database Home.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

VM Clusters 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. In the Resources list of the VM Cluster Details page, click **Databases**.
5. Click the name of the database that you want to move.

   The Database Details page displays information about the selected database.
6. Click **More actions**, and then click **Move to Another Home**.
7. On the resulting Move Database to Another Home page:
   - 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.

   - Click **Advanced options**, choose one of the following **Datapatch options**.
     - **Run datapatch on the database**: This is the default option, which applies the datapatch operation to the selected database.
     - **Skip datapatch on the database**: Select this option to skip the datapatch execution on the selected database.
     - **Skip datapatch for closed PDBs only**: Select this option to skip datapatch on only those PDBs that are in a closed state during the update.
8. Click **Move**.

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 `Updating`. With the default "**Run Datapatch on the database**" option, the datapatch operation is automatically executed on the database during the move process, completing post-patch SQL actions for all patches, including one-off patches. 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 run datapatch on a container database (CDB)

To run datapatch on a container database, use this procedure.

> **Note**
>
> This option is visible only if the datapatch operation is skipped on the container database, causing its version to be inconsistent than that of the Database Home. This discrepancy is also highlighted by the version mismatch banner at the top of the Database Details page.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

   VM Clusters is selected by default.
2. Choose your **Compartment** that contains the VM cluster that hosts the database that you want to move.
3. Click **Exadata VM Clusters**.
4. In the list of VM clusters, click the name of the VM cluster that contains the database on which you wan to run datapatch.
5. On the resulting VM Cluster details page, click **Databases**.
6. Click the name of the database on which you want to run datapatch.
7. Click **More actions**, then click **Run datapatch**.
8. On the Run datapatch page, click the **Skip datapatch on closed PDBs** toggle button to skip the datapatch operation on all associated closed PDBs.
9. Click **Run datapatch**.

To run datapatch on a pluggable database (PDB)

To run datapatch on a pluggable database, use this procedure.

> **Note**
>
> This option is visible only if the datapatch operation is skipped on the pluggable database, causing its version to be inconsistent than that of the Database Home. This discrepancy is also highlighted by the version mismatch banner at the top of the Database Details page.

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

   VM Clusters is selected by default.
2. Choose your **Compartment** that contains the VM cluster that hosts the database that you want to move.
3. Click **Exadata VM Clusters**.
4. 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. On the VM cluster details page, in the databases table, click the name of the database that contains the pluggable database you want to run datapatch.
6. On the resulting database details page, click **Pluggable Databases**.
7. Click the name of the PDB where you want to run datapatch.
8. Click **More actions**, then click **Run datapatch**.
9. On the Run datapatch dialog, review the message and then click **Run datapatch**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Terminate a Database

You can terminate a database and thereby remove the terminated database from the Cloud Control Plane.

Terminating a database removes it from the Cloud Control Plane. In the process, all of the associated data files and backups are destroyed.

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

   VM Clusters is selected by default.
2. Choose your **Compartment** that contains the VM cluster that hosts the database that you want to terminate.
3. Click the name of the VM cluster that contains the database that you want to terminate.
4. In the Resources list of the VM Cluster Details page, click **Databases**.
5. Click the name of the database that you want to terminate.

   The Database Details page displays information about the selected database.
6. Click **More actions**, and then select **Terminate**.
7. In the resulting dialog, enter the name of the database, and then click **Terminate Database** to confirm the action.

> **Note**
>
> You can now:
>
> - Terminate a CDB when Data Guard setup is running on another database in the same Oracle home, and vice versa.
> - Create or delete a CDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.

Using the API to Manage Oracle Database Components

Use various API features to help manage your databases on Oracle Exadata Database Service on Cloud@Customer.

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 following API operations to manage various database components.

Database homes:

- `CreateDbHome`
- `DeleteDbHome`
- `GetDbHome`
- `ListDbHomes`

Databases:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `CreateDatabase`
- `GetDatabase`
- `ListDatabases`
- `UpdateDatabase`
- `UpdateDatabaseDetails`

Nodes:

- `GetDbNode`
- `List DbNodes`

Use `UpdateDatabase` to move a database to a different Database Home, thereby updating the database to the same version as the target Database Home.

For the complete list of APIs, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- CreateDbHome
- DeleteDbHome
- GetDbHome
- ListDbHomes
- CreateDatabase
- GetDatabase
- ListDatabases
- UpdateDatabase
- UpdateDatabaseDetails
- GetDbNode
- List DbNodes
- Database Service API

Changing the Database Passwords

To change the SYS password, or to change the TDE wallet password, use this procedure.

The password that you specify in the **Database Admin Password** field when you create a new Exadata Database Service on Cloud@Customer 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 © 2008, 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 the SYS Password for an Exadata Database Service on Cloud@Customer Database

1. Log onto the Exadata Database Service on Cloud@Customer virtual machine as `opc`.
2. Run the following command:

   ```
   sudo dbaascli database changepassword --dbname database_name --user SYS
   ```

To Change Database Passwords in a Data Guard Environment

1. Run the following command on the primary database:

   ```
   dbaascli database changePassword –dbName <dbname> --user SYS --
   prepareStandbyBlob true --blobLocation <location to create the blob file>
   ```

2. Copy the blob file created to all the standby databases and update the file ownership to `oracle` user.
3. Run the following command on all the standby databases:

   ```
   dbaascli database changePassword –dbName <dbname> --user SYS --
   standbyBlobFromPrimary <location of copies the blob file>
   ```

To Change the TDE Wallet Password for an Exadata Database Service on Cloud@Customer Database

1. Log onto the Exadata Database Service on Cloud@Customer virtual machine as `opc`.
2. Run the following command:

   ```
   sudo dbaascli tde changepassword --dbname database_name
   ```

Manage Pluggable Databases on Oracle Exadata Database Service on Cloud@Customer

Learn to manage pluggable databases on Oracle Exadata Database Service on Cloud@Customer.

- Pluggable Database Operations
  You can create and manage pluggable databases (PDBs) in Oracle Exadata Database Service on Cloud@Customer systems using the Console and APIs.
- Concurrently Create or Delete Pluggable Databases (PDBs)
  You can now create or delete up to 10 PDBs concurrently, even when the container database (CDB) is in an updating state. However, you cannot create or delete PDBs while

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

- Cost and Usage Attribution for Pluggable Databases (PDBs)

Pluggable Database Operations

You can create and manage pluggable databases (PDBs) in Oracle Exadata Database Service on Cloud@Customer systems 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*.

Oracle 19c or later databases created in a virtual machine include an initial PDB that you can access from the CDB's **Database Details** page in the Console. Using the Console or APIs, you can start, stop, clone, and delete the PDB. You can also create additional PDBs in the container database. You can monitor all PDB operations performed using the Console or APIs using the *work request* generated by the operation.

- **Backup**
  You can take a backup of the PDB optionally during create, clone, or relocate operations when the CDB is configured with the auto-backup feature. The PDB backup destination will always be the same as CDB, and the backups cannot be accessed directly or created on demand. Oracle recommends immediately backing up the PDB after you create or clone it. This is because the PDB will not be recoverable until the next daily auto-backup completes successfully, leading to a possible data loss.
- **Restore**
  - **Base Database Service / Oracle Exadata Database Service on Dedicated Infrastructure:**
    - **In place restore:** You can restore a PDB within the same CDB to last known good state or to a specified timestamp.
    - **Out of place restore:** You can restore a PDB by creating a database (CDB) from the backup, then selecting a PDB or a subset of them you want to restore on the new database.
  - **Oracle Exadata Database Service on Cloud@Customer:**
    - **In place restore:** You can restore a PDB within the same CDB to last known good state and specified timestamp.
    - **Out of place restore:** It's not available.

  You can perform an in-place restore when you want to move a PDB back to a specified state or time. Both the CDB and PDB must be up and running and only one PDB can be restored at a time.
  - If you have multiple PDBs in your CDB and want to restore multiple of them to the same CDB, then you could restore each individual PDB, one PDB at a time, from the CDB backup.
  - When the CDB is down, you could restore the complete CDB and all the PDBs in that CDB will also be restored.
  - You could either restore the database to the specified timestamp or to its last known good state.
- **Relocate**

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

- Across compartments, VM clusters, DB system (for BaseDB only), or VCNs (not applicable to ExaDB-C@C). If two different VCNs are used, then both VCNs must be peered before relocating.
- To the same or a higher database version.

During relocate, the PDB will be removed from the source CDB and moved to the destination CDB that is up and running. In a Data Guard association, a PDB relocated to the primary will be synchronized with the standby as well.

- **Clone**
  A clone is an independent and complete copy of the given database as it existed at the time of the cloning operation. You can create clones of your PDB within the same CDB or a different CDB and refresh the cloned PDB.

  The following types of clones are supported:

  - **Local clone:** A copy of the PDB is created within the same CDB.
  - **Remote clone:** A copy of the PDB is created in a different CDB.
    You can perform a remote clone of a PDB from one CDB to another CDB within the same availability domain (AD):
    * Across compartments, VM clusters, DB system (for BaseDB only), or VCNs (not applicable to ExaDB-C@C). If two different VCNs are used, then both VCNs must be peered before cloning.
    * To the same or a higher database version.
  - **Refreshable clone:** A copy of the PDB is created in a different CDB, and you will be able to refresh the cloned PDB.

    You can perform a refreshable clone of a PDB from one CDB to another CDB within the same availability domain (AD):
    * Across compartments, VM clusters, DB system (for BaseDB only), or VCNs (not applicable to ExaDB-C@C). If two different VCNs are used, then both VCNs must be peered before cloning.
    * To the same or a higher database version.
- **Refreshable Clone**
  A refreshable clone enables you to keep your remote clone updated with the source PDB. You can only refresh while the PDB is in mount mode. The only open mode you can have is read-only and refresh cannot be done while it is in read-only mode.

  - A database link user credential is required for creating a refreshable clone.
  - Clone, relocate, and in-place restore operations are not supported in the refreshable clone. Relocate and in-place restore operations are not supported in the source, and the source can only be deleted after disconnecting or deleting the refreshable clone.
  - In a Data Guard association, a refreshable clone cannot be created on standby, but it can be created on the primary. However, the primary will not be synced to the standby.

  > **Note**
  >
  > A PDB in standby cannot be used as the source for a refreshable PDB.

- **Convert Refreshable PDB to Regular PDB**

Copyright © 2008, 2026, Oracle and/or its affiliates.

You can convert a refreshable PDB to a regular PDB by disconnecting the refreshable clone (destination PDB) from the source PDB at any time. If the refresh PDB is in a Data Guard association, when it is converted to a regular PDB the PDB will be synced to the standby as part of the conversion process.

- **Open Modes**
  On the Console, you can see the open modes of a PDB, such as read-write, read-only, and mounted. If the PDB status is the same across all nodes, the system displays the same status for all PDBs. If the PDB statuses are different across the nodes, the system displays a message indicating on which nodes the PDBs are opened in read-write mode. You cannot change the open mode of a PDB through the API or Console. However, you can start or stop a PDB. Starting the PDB will start it in read-write mode. Stopping the PDB will close it and it will remain in mount mode.
- Limitations for Pluggable Database Management
  Review the list of limitations in managing PDBs.
- Create a Pluggable Database
  You can create a PDB from the OCI Console, or with the pluggable database APIs.
- Manage a Pluggable Database
  To start, stop, clone, and delete a PDB, use these procedures.

Related Topics

- Multitenant Architecture
- Work Requests

Limitations for Pluggable Database Management

Review the list of limitations in managing PDBs.

- 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 with Oracle Exadata Database Service on Cloud@Customer* 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 with Oracle Exadata Database Service on Cloud@Customer* for more information.

Create a Pluggable Database

You can create a PDB from the OCI Console, or with the pluggable database APIs.

- Using the Console to Create a Pluggable Database
  To create a pluggable database with the console, use this procedure.
- Using the Console to Relocate a Pluggable Database
  To relocate a pluggable database with the console, use this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the API to Create a Pluggable Database
  Use various API features to help create your pluggable databases on Oracle Exadata Database Service on Cloud@Customer.

Using the Console to Create a Pluggable Database

To create a pluggable database with the console, use this procedure.

> **Note**
>
> If the databases are created directly on Guest VM, the attributed usage data would be delayed.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
   VM Clusters is selected by default.
2. Choose your **Compartment**.
   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of cloud VM clusters, click the name of the cluster in which you want to create the PDB, and then click its name to display the database details page.
4. In the lower-left corner of the database details page, click **Pluggable Databases**.
   A list of existing PDBs in this database is displayed.
5. Click **Create Pluggable Database**.
   The Create Pluggable Database dialog box is displayed.
6. In the Create Pluggable Database dialog box, enter the following:
   - **Enter PDB Name:** Enter a name for the PDB. The name must begin with an alphabetic character and can contain a maximum of 30 alphanumeric characters.
   - **Unlock my PDB Admin Account:**
     - To enter the administrator's password, check this check box.
       - **PDB Admin Password:** Enter PDB admin password. The password must contain:
         - a minimum of 9 and a maximum of 30 characters
         - at least two uppercase characters
         - at least two lowercase characters
         - at least two special characters. The valid special characters are underscore (_), a pound or hash sign (#), and dash (-). You can use two of the same characters or any combination of two of the same characters.
         - at least two numeric characters (0 - 9)
       - **Confirm PDB Admin Password:** Enter the same PDB Admin password in the confirmation field.
     - To skip entering the administrator's password, uncheck this check box. If you uncheck this check box, then the PDB is created but you cannot use it. To use the PDB, you must reset the administrator password.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> When you create a new PDB, a local user in the PDB is created as the administrator and granted the `PDB_DBA` role locally to the administrator.

To reset the password:

a. Connect to the container where your PDB exists using the SQL*Plus `CONNECT` statement.

```
SQL> show con_name;
CON_NAME
------------------------
CDB$ROOT
```

For more information, see *Administering a CDB* and *Administering PDBs* in the *Oracle® Multitenant Administrator's Guide*.

b. Find the administrator name of your PDB:

```
SQL> select grantee from cdb_role_privs where con_id = (select
con_id from cdb_pdbs where pdb_name = '<PDB_NAME>') and
granted_role = 'PDB_DBA';
```

c. Switch into your PDB:

```
SQL> alter session set container=<PDB_NAME>;
Session altered.


SQL> show con_name;
CON_NAME
------------------------
<PDB_NAME>
```

d. Reset the PDB administrator password:

```
SQL> alter user <PDB_Admin> identified by <PASSWORD>;
User altered.
```

- **TDE Wallet password of database:** Enter a wallet password for the CDB. This password has the same rules as the PDB Admin 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.

Copyright © 2008, 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.
  - **Encryption Key Store:** This shows if the encryption key for the CDB is managed by Oracle or by the customer. In the case of customer-manager keys, the name of the key store you have configured for CDB is displayed. You cannot edit this field.

7. Click **Create**.
   The system starts the creation process and opens the Work Request page for the new PDB. The Work request page shows the status of the creation process of the new PDB.

   By default, the Work Request details page shows the log messages created by the system. Click **Error Messages** or **Associated Resources** to see any error messages or associated resources for the creation process, in the Resources area on the left side of the page.

> **Note**
>
> - The numbers at the right side of the **Log Messages**, **Error Messages**, and **Associated Resources** links indicate how many of each item exists.
> - Create or delete a PDB while a Data Guard setup is running on another database within the same Oracle home, and vice versa.
> - Create or delete a PDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.
> - Create or delete a PDB concurrently on different databases within the same Oracle home.
> - Create or delete a PDB while simultaneously updating VM Cluster tags.

Related Topics

- Administering a CDB
- Administering PDBs

Using the Console to Relocate a Pluggable Database

To relocate a pluggable database with the console, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
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 cluster in which you want to create the PDB, and then click its name to display the database details page.
4. In the lower-left corner of the database details page, click **Pluggable Databases**.
   A list of existing PDBs in this database is displayed.
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**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

(or)

Click the Actions menu (three dots) and select **Relocate**.

6. In the resulting Relocate Pluggable Database window, enter the following:
   - **VM Cluster:** Use the menu to select the destination VM cluster.
   - **Destination database:** Use the menu to select an existing database where the PDB will be created. This database can be of the same version as the CDB the source PDB is in or of a higher version.
   - **New PDB name for the clone:** The name must begin with an alphabetic character and can contain up to 30 characters. To keep the PDB name the same, just re-enter the source PDB name.
   - **Database TDE wallet password:** Enter the TDE wallet password for the parent CDB of the source PDB.
   - **Unlock my PDB Admin Account:**
     - To enter the administrator's password, check this check box.
       - **PDB Admin Password:** Enter PDB admin password. The password must contain:
         - a minimum of 9 and a maximum of 30 characters
         - at least two uppercase characters
         - at least two lowercase characters
         - at least two special characters. The valid special characters are underscore (_), a pound or hash sign (#), and dash (-). You can use two of the same characters or any combination of two of the same characters.
         - at least two numeric characters (0 - 9)
       - **Confirm PDB Admin Password:** Enter the same PDB Admin password in the confirmation field.
     - To skip entering the administrator's password, uncheck this check box. If you uncheck this check box, then the PDB is created but you cannot use it. To use the PDB, you must reset the administrator password.

       > **Note**
       >
       > When you create a new PDB, a local user in the PDB is created as the administrator and granted the `PDB_DBA` role locally to the administrator.

       To reset the password:

       a. Connect to the container where your PDB exists using the SQL*Plus `CONNECT` statement.

       ```
       SQL> show con_name;
       CON_NAME
       ------------------------
       CDB$ROOT
       ```

       For more information, see *Administering a CDB* and *Administering PDBs* in the *Oracle® Multitenant Administrator's Guide*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

b. Find the administrator name of your PDB:

```
SQL> select grantee from cdb_role_privs where con_id = (select
con_id from cdb_pdbs where pdb_name = '<PDB_NAME>') and
granted_role = 'PDB_DBA';
```

c. Switch into your PDB:

```
SQL> alter session set container=<PDB_NAME>;
Session altered.


SQL> show con_name;
CON_NAME
------------------------
<PDB_NAME>
```

d. Reset the PDB administrator password:

```
SQL> alter user <PDB_Admin> identified by <PASSWORD>;
User altered.
```

- **TDE Wallet password of database:** Enter a wallet password for the CDB. This password has the same rules as the PDB Admin Password.
- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

> **Note**
>
> If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Advanced Options:**
  - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

7. Click **Relocate**.

> **Note**
>
> Relocate will incur downtime during the process and that the time required is based on the size of the PDB.

Related Topics

- Administering a CDB
- Administering PDBs

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Resource Tags

Using the API to Create a Pluggable Database

Use various API features to help create your pluggable databases on Oracle Exadata Database Service on Cloud@Customer.

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 this API operation to create pluggable databases on Oracle Exadata Database Service on Cloud@Customer systems.

- *CreatePluggableDatabase*

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- CreatePluggableDatabase
- Database Service API

Manage a Pluggable Database

To start, stop, clone, and delete a PDB, use these procedures.

- Using the Console to Start a Pluggable Database
  The PDB must be available and stopped to use this procedure.
- Using the Console to Stop a Pluggable Database
  The PDB must be available and running (started) to use this procedure.
- Using the Console to Delete a Pluggable Database
  The PDB must be available and stopped to use this procedure.
- Using the Console to Get Connection Strings for a Pluggable Database
  Learn how to get connection strings for the administrative service of a PDB. Oracle recommends connecting applications to an application service using the strings created for the application service.
- Clone a Pluggable Database (PDB)
  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.
- Restore a Pluggable Database (PDB)
  You can restore a PDB within the same CDB to last known good state and specified timestamp.
- Convert a Physical Standby Database to Snapshot Standby Database
  A snapshot standby database is a fully updateable standby database created by converting a physical standby database into a snapshot standby database.
- Convert a Snapshot Standby Database to Physical Standby Database
  Oracle recommends converting a snapshot standby database back to a physical standby database within 14 days.
- Create an Exascale PDB Snapshot

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Terminate an Exascale PDB Snapshot
- Create a Thick or Thin Local or Remote Clone from a PDB Snapshot
- Using the API to Manage Pluggable Databases
  Use various API features to help manage your pluggable databases on Oracle Exadata Database Service on Cloud@Customer.
- Using the API to Clone a Pluggable Database
  Clone a pluggable database (PDB) in the same database (CDB) as the source PDB or to a different database from the source PDB.

Using the Console to Start a Pluggable Database

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

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
2. Choose your **Compartment**.
   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the name of the VM cluster that contains the PDB you want to start, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to start.
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 start.
   The pluggable details page is displayed.
8. Click **Start**.
   The Start PDB dialog box is displayed.
9. Click **Start** to confirm the start operation.

Using the Console to Stop a Pluggable Database

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

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
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 stop, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to stop.
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 stop.
   The pluggable details page is displayed.
8. Click **Stop**.
   The Stop PDB dialog box is displayed.
9. Click **Stop** to confirm the stop operation.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Delete a Pluggable Database

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

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
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 delete, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to delete.
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 delete.
   The pluggable details page is displayed.
8. Click More Actions, and then choose **Delete**.
   The Delete PDB dialog box is displayed.
9. Click **Delete** to confirm the delete operation.

> **Note**
>
> You can now delete a PDB when Data Guard setup is running on another database in the same Oracle home, and vice versa.

Using the Console to Get Connection Strings for a Pluggable Database

Learn how to get connection strings for the administrative service of a PDB. Oracle recommends connecting applications to an application service using the strings created for the application service.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
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 get connections strings for, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to get connection strings.
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 get connection strings.
   The pluggable details page is displayed.
8. Click **PDB Connection**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

9. In the Pluggable Database Connection dialog, use the **Show** and **Copy** links to display and copy connection strings, as needed.
10. Click **Close** to exit the dialog.

Clone a Pluggable Database (PDB)

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.
  For more information about refreshable clones, see About Refreshable Clone PDBs..
- Using the Console to Create a Local Clone of a Pluggable Database (PDB)
  To create a local clone of your PDBs, follow this procedure.
- Using the Console to Create a Remote Clone of a Pluggable Database (PDB)
  To create a remote clone of your PDBs, follow this procedure.
- Using the Console to Create a Refreshable Clone of a Pluggable Database (PDB)
  To create a refreshable clone of your PDBs, follow this procedure.

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

To create a local clone of your PDBs, follow this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
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 clone, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to clone.
5. Click **Pluggable Databases** in the **Resources** section of the page.
   A list of existing PDBs in this database is displayed.
6. Click the name of the PDB that you want to clone.
   The pluggable details page is displayed.
7. Click **Clone**.
8. In the Clone PDB dialog box, enter the following:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Select clone type:** Select **Local clone** to create a copy of the source PDB to the same CDB.
- **VM Cluster:** Use the menu to select the source VM cluster.
- **Destination database:** This field is disabled.
- **New PDB name for the clone:** The name must begin with an alphabetic character and can contain up to 30 characters.
- **Database TDE wallet password:** Enter the TDE wallet password for the parent CDB of the source PDB.
- **Unlock my PDB Admin Account:**
  - To enter the administrator's password, check this check box.
    - **PDB Admin Password:** Enter PDB admin password. The password must contain:
      - a minimum of 9 and a maximum of 30 characters
      - at least two uppercase characters
      - at least two lowercase characters
      - at least two special characters. The valid special characters are underscore (_), a pound or hash sign (#), and dash (-). You can use two of the same characters or any combination of two of the same characters.
      - at least two numeric characters (0 - 9)
    - **Confirm PDB Admin Password:** Enter the same PDB Admin password in the confirmation field.
  - To skip entering the administrator's password, uncheck this check box. If you uncheck this check box, then the PDB is created but you cannot use it. To use the PDB, you must reset the administrator password.

> **Note**
>
> When you create a new PDB, a local user in the PDB is created as the administrator and granted the `PDB_DBA` role locally to the administrator.

To reset the password:

a. Connect to the container where your PDB exists using the SQL*Plus `CONNECT` statement.

```
SQL> show con_name;
CON_NAME
------------------------
CDB$ROOT
```

For more information, see *Administering a CDB* and *Administering PDBs* in the *Oracle® Multitenant Administrator's Guide*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

b. Find the administrator name of your PDB:

```
SQL> select grantee from cdb_role_privs where con_id = (select
con_id from cdb_pdbs where pdb_name = '<PDB_NAME>') and
granted_role = 'PDB_DBA';
```

c. Switch into your PDB:

```
SQL> alter session set container=<PDB_NAME>;
Session altered.


SQL> show con_name;
CON_NAME
------------------------
<PDB_NAME>
```

d. Reset the PDB administrator password:

```
SQL> alter user <PDB_Admin> identified by <PASSWORD>;
User altered.
```

- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

> **Note**
>
> If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Enable thin clone:** All clones are thin clones by default. For more information about thin clone, see Thin Cloning a Pluggable Database in the Oracle® Exadata Exascale User's Guide. If you specifically want a thick clone (full copy), you need to deselect this option.
- **Advanced Options:**
  - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

9. Click **Clone**.

Related Topics

- Resource Tags

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

To create a remote clone of your PDBs, follow this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
2. Choose your **Compartment**.

Copyright © 2008, 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 clone, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to clone.
5. Click **Pluggable Databases** in the **Resources** section of the page.
   A list of existing PDBs in this database is displayed.
6. Click the name of the PDB that you want to clone.
   The pluggable details page is displayed.
7. Click **Clone**.
8. 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.
   - **VM Cluster:** Use the menu to select the destination VM cluster.
   - **Destination database:** Use the menu to select an existing database where the PDB will be created. This database can be of the same version as the CDB the source PDB is in or of a higher version.
   - **New PDB name for the clone:** The name must begin with an alphabetic character and can contain up to 30 characters.
   - **Database TDE wallet password:** Enter the TDE wallet password for the parent CDB of the source PDB.
   - **Unlock my PDB Admin Account:**
     – To enter the administrator's password, check this check box.
       * **PDB Admin Password:** Enter PDB admin password. The password must contain:
         * a minimum of 9 and a maximum of 30 characters
         * at least two uppercase characters
         * at least two lowercase characters
         * at least two special characters. The valid special characters are underscore (_), a pound or hash sign (#), and dash (-). You can use two of the same characters or any combination of two of the same characters.
         * at least two numeric characters (0 - 9)
       * **Confirm PDB Admin Password:** Enter the same PDB Admin password in the confirmation field.
     – To skip entering the administrator's password, uncheck this check box. If you uncheck this check box, then the PDB is created but you cannot use it. To use the PDB, you must reset the administrator password.

       > **Note**
       >
       > When you create a new PDB, a local user in the PDB is created as the administrator and granted the `PDB_DBA` role locally to the administrator.

       To reset the password:

Copyright © 2008, 2026, Oracle and/or its affiliates.

a. Connect to the container where your PDB exists using the SQL*Plus `CONNECT` statement.

```
SQL> show con_name;
CON_NAME
------------------------
CDB$ROOT
```

For more information, see *Administering a CDB* and *Administering PDBs* in the *Oracle® Multitenant Administrator's Guide*.

b. Find the administrator name of your PDB:

```
SQL> select grantee from cdb_role_privs where con_id = (select
con_id from cdb_pdbs where pdb_name = '<PDB_NAME>') and
granted_role = 'PDB_DBA';
```

c. Switch into your PDB:

```
SQL> alter session set container=<PDB_NAME>;
Session altered.


SQL> show con_name;
CON_NAME
------------------------
<PDB_NAME>
```

d. Reset the PDB administrator password:

```
SQL> alter user <PDB_Admin> identified by <PASSWORD>;
User altered.
```

- **Source database SYS password:** Enter the database admin password.
- **Database link:** Enter the user name and password for the database link. Note that the user must be precreated in the source database. The DB link will be created in the destination using that username and password.

> **Note**
>
> If you do not provide database link information, cloud automation will create a database link using the common user. However, you can specify the database link information if you want cloud automation to use a specific database link.

- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Enable thin clone:** All clones are thin clones by default. For more information about thin clone, see Thin Cloning a Pluggable Database in the Oracle® Exadata Exascale User's Guide. If you specifically want a thick clone (full copy), you need to deselect this option.

> **Note**
>
> The thin clone option will be disabled (greyed out) if the source and target VM clusters do not share the same vault. In such cases, only thick clones are supported.

- **Advanced Options:**
  - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

9. Click **Clone**.

Related Topics

- Resource Tags

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

To create a refreshable clone of your PDBs, follow this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
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 clone, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to clone.
5. Click **Pluggable Databases** in the **Resources** section of the page.
   A list of existing PDBs in this database is displayed.
6. Click the name of the PDB that you want to clone.
   The pluggable details page is displayed.
7. Click **Clone**.
8. 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.
   - **VM Cluster:** Use the menu to select the destination VM cluster.

Copyright © 2008, 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.
- **Database TDE wallet password:** Enter the TDE wallet password for the parent CDB of the source PDB.
- **Unlock my PDB Admin Account:**
  - To enter the administrator's password, check this check box.
    - **PDB Admin Password:** Enter PDB admin password. The password must contain:
      - a minimum of 9 and a maximum of 30 characters
      - at least two uppercase characters
      - at least two lowercase characters
      - at least two special characters. The valid special characters are underscore (_), a pound or hash sign (#), and dash (-). You can use two of the same characters or any combination of two of the same characters.
      - at least two numeric characters (0 - 9)
    - **Confirm PDB Admin Password:** Enter the same PDB Admin password in the confirmation field.
  - To skip entering the administrator's password, uncheck this check box. If you uncheck this check box, then the PDB is created but you cannot use it. To use the PDB, you must reset the administrator password.

> **Note**
>
> When you create a new PDB, a local user in the PDB is created as the administrator and granted the `PDB_DBA` role locally to the administrator.

To reset the password:

a. Connect to the container where your PDB exists using the SQL*Plus `CONNECT` statement.

```
SQL> show con_name;
CON_NAME
------------------------
CDB$ROOT
```

For more information, see *Administering a CDB* and *Administering PDBs* in the *Oracle® Multitenant Administrator's Guide*.

b. Find the administrator name of your PDB:

```
SQL> select grantee from cdb_role_privs where con_id = (select
con_id from cdb_pdbs where pdb_name = '<PDB_NAME>') and
granted_role = 'PDB_DBA';
```


Copyright © 2008, 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.

> **Note**
>
> For a refreshable clone, the database link parameter is mandatory, and you must provide this information.

- **Enable thin clone:** All clones are thin clones by default. For more information about thin clone, see Thin Cloning a Pluggable Database in the Oracle® Exadata Exascale User's Guide. If you specifically want a thick clone (full copy), you need to deselect this option.

> **Note**
>
> The thin clone option will be disabled (greyed out) if the source and target VM clusters do not share the same vault. In such cases, only thick clones are supported.

- **Advanced Options:**
  - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

9. Click **Clone**.

- Using the Console to Refresh a Cloned Pluggable Database (PDB)
  To create a refresh a cloned PDB, follow this procedure.
- Using the Console to Convert a Refreshable Clone to a Regular Pluggable Database (PDB)
  To create a convert a refreshable clone to a regular PDB, follow this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Resource Tags

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

To create a refresh a cloned PDB, follow this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
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 refresh, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to refresh.
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 refresh.
   The pluggable details page is displayed.
8. Click **More Actions** and select **Refresh**.
9. In the resulting Refresh dialog box, click **Refresh** to confirm.

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

To create a convert a refreshable clone to a regular PDB, follow this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
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 convert to a regular PDB, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to convert to a regular PDB.
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 convert to a regular PDB.
   From the Pluggable Database details page, click **More Actions**, and then select **Convert to regular PDB**.

   (or)

   Click the Actions menu (three dots) and select **Convert to regular PDB**.
8. In the resulting Convert to regular PDB dialog, enter the following:
   - **Database TDE wallet password:** Enter the TDE wallet password for the parent 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **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.

9. Click **Convert**.

Restore a Pluggable Database (PDB)

You can restore a PDB within the same CDB to last known good state and specified timestamp.

- Using the Console to Perform an In-Place Restore of a Pluggable Database (PDB)
  To perform an in-place restore, follow this procedure.

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

To perform an in-place restore, follow this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud@Customer**. VM Clusters is selected by default.
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 restore, and then click its name to display the details page.
4. Under **Databases**, find the database containing the refreshable PDB you want to restore.
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 restore.
   From the Pluggable Database details page, click **More Actions**, and then select **Restore**.

   (or)

   Click the Actions menu (three dots) and select **Restore**.
8. 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.
9. Click **Restore**.

Convert a Physical Standby Database to Snapshot Standby Database

A snapshot standby database is a fully updateable standby database created by converting a physical standby database into a snapshot standby database.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

A snapshot standby database typically diverges from its primary database over time because redo data from the primary database is not applied as it is received. Local updates to the snapshot standby database cause additional divergence. The data in the primary database is fully protected however, because a snapshot standby can be converted back into a physical standby database at any time, and the redo data received from the primary is then applied.

A snapshot standby database provides disaster recovery and data protection benefits that are similar to those of a physical standby database. Snapshot standby databases are best used in scenarios where the benefit of having a temporary, updatable snapshot of the primary database justifies increased time to recover from primary database failures.

Note

- You will have to ensure there is sufficient space in the Fast Recovery Area. Oracle recommends converting a snapshot standby database back to a physical standby database in the shortest amount of time.
- Automatic backups are suspended on a snapshot standby. Archive logs from the production database will continue to be backed up to preserve recoverability.
- If you have performed a database upgrade while in the snapshot standby mode, you must ensure you use the correct Oracle Home when converting back to a physical standby database.

Before converting a physical standby database to a snapshot standby database, ensure the following conditions are met:

- It is recommended that you have more than one standby database in a Data Guard Group before converting a standby database to a snapshot standby database.
- Only one physical standby can be converted to a snapshot standby at a time. To convert multiple databases, perform the conversions sequentially.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click on the VM Cluster containing the databases you wish to view their roles in Data Guard associations.
4. In the **Databases** section under **Resources**, the role of each database in this VM Cluster is indicated in the Data Guard role column.
5. Choose a Standby database, select **Convert to snapshot standby** from the **More Actions** drop-down list.
6. Review the warning message on the Convert to snapshot standby window.
7. Click **Convert**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

After the conversion, the role of the standby database changes to **Snapshot Standby**, and the **Convert to Standby** option becomes available in the **More Actions** drop-down list.

Convert a Snapshot Standby Database to Physical Standby Database

Oracle recommends converting a snapshot standby database back to a physical standby database within 14 days.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click on the VM Cluster containing the databases you wish to view their roles in Data Guard associations.
4. In the **Databases** section under **Resources**, the role of each database in this VM Cluster is indicated in the Data Guard role column.
5. Choose a Snapshot standby database, select **Convert to Standby** from the **More Actions** drop-down list.

> **Note**
>
> If you convert your snapshot standby to physical standby database, all local updates to your snapshot standby database will be discarded and data from your primary database will be applied.

6. Click **Convert**.

> **Note**
>
> - You cannot perform a switchover operation on a snapshot standby database. You will need to convert the snapshot standby to a standby database before initiating a switchover.
> - Oracle managed automatic backups will be disabled for a snapshot standby database. Automatic backups if configured will resume upon conversion from snapshot standby to standby database.
> - While any database in the Data Guard configuration is in snapshot standby mode, there is no support to create additional PDBs on the primary or any of the standby databases.

Create an Exascale PDB Snapshot

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, locate the one with the storage type set to Exascale, where you want to create the Exascale PDB snapshot.
4. Click the name of the cluster.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. On the VM Cluster Details page, click the **Database** tab.
   A list of databases on this VM cluster is displayed.
6. Click the name of the database that contains the desired PDB.
7. On the Database Details page, click **Pluggable Databases**.
   A list of existing PDBs in this database is displayed.
8. Click the name of the pluggable database for which you want to create a snapshot.
9. On the Pluggable Database Details page, click **Create Exascale PDB snapshot**.
10. On the resulting Create Exascale PDB snapshot dialog, enter a descriptive name for the snapshot.
11. Click **Create**.
    The status of the workflow is displayed on the Pluggable Database Details page.

Terminate an Exascale PDB Snapshot

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
   VM Clusters is selected by default.
2. Choose your **Compartment**.
   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, locate the one with the storage type set to Exascale, where you have created the Exascale PDB snapshot.
4. Click the name of the cluster.
5. On the VM Cluster Details page, click the **Database** tab.
   A list of databases on this VM cluster is displayed.
6. Click the name of the database that contains the desired PDB.
7. On the Database Details page, click **Pluggable Databases**.
   A list of existing PDBs in this database is displayed.
8. Click the name of the pluggable database for which you have created a snapshot.
9. On the Pluggable Database Details page, click the **Exascale PDB snapshots** tab and identify the snapshot you want to terminate.
10. From the Actions menu (three dots), click **Terminate**.
11. On the resulting Terminate dialog, enter name of the snapshot, and then click **Terminate**.

Create a Thick or Thin Local or Remote Clone from a PDB Snapshot

> **Note**
>
> For more information about thin clones, see Thin Cloning a Pluggable Database in the Oracle® Exadata Exascale User's Guide.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
   VM Clusters is selected by default.
2. Choose your **Compartment**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

A list of VM Clusters is displayed for the chosen Compartment.

3. In the list of VM clusters, locate the one with the storage type set to Exascale, where you have created the Exascale PDB snapshot.
4. Click the name of the cluster.
5. On the VM Cluster Details page, click the **Database** tab.
   A list of databases on this VM cluster is displayed.
6. Click the name of the database that contains the desired PDB.
7. On the Database Details page, click **Pluggable Databases**.
   A list of existing PDBs in this database is displayed.
8. Click the name of the pluggable database for which you have created a snapshot.
9. On the Pluggable Database Details page, identify thesnapshot for which you want to create a thick or thin clone.
10. From the Actions menu (three dots), click **Clone from Snapshot**.
11. On the Clone pluggable database page, you can choose to create a local thick or thin clone.
    a. The source VM cluster name and CDB name is pre-populated under Destination section.
    b. Enter a new PDB name for the clone. The name must begin with an alphabetic character and can contain up to 30 characters.
    c. Enter the TDE wallet password for the parent CDB of the source PDB.
    d. **Unlock my PDB Admin Account:**
       - To enter the administrator's password, check this check box.
         - **PDB Admin Password:** Enter PDB admin password. The password must contain:
           * a minimum of 9 and a maximum of 30 characters
           * at least two uppercase characters
           * at least two lowercase characters
           * at least two special characters. The valid special characters are underscore (_), a pound or hash sign (#), and dash (-). You can use two of the same characters or any combination of two of the same characters.
           * at least two numeric characters (0 - 9)
         - **Confirm PDB Admin Password:** Enter the same PDB Admin password in the confirmation field.
       - To skip entering the administrator's password, uncheck this check box. If you uncheck this check box, then the PDB is created but you cannot use it. To use the PDB, you must reset the administrator password.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **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:
>
> i. Connect to the container where your PDB exists using the SQL*Plus CONNECT statement.
>
> ```
> SQL> show con_name;
> CON_NAME
> ------------------------
> CDB$ROOT
> ```
>
> For more information, see *Administering a CDB and Administering PDBs* in the *Oracle® Multitenant Administrator's Guide*.
>
> ii. 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';
> ```
>
> iii. Switch into your PDB:
>
> ```
> SQL> alter session set container=<PDB_NAME>;
> Session altered.
> SQL> show con_name;
> CON_NAME
> ------------------------
> <PDB_NAME>
> ```
>
> iv. Reset the PDB administrator password:
>
> ```
> SQL> alter user <PDB_Admin> identified by <PASSWORD>;
> User altered.
> ```

e. **Source database SYS password:** (Applies to remote clone) Enter the database admin password.

f. **Database link:** (Applies to remote clone) Enter the user name and password for the database link. Note that the user must be precreated in the source database. The DB link will be created in the destination using that username and password.

> **Note**
>
> If you do not provide database link information, cloud automation will create a database link using the common user. However, you can specify the database link information if you want cloud automation to use a specific database link.

g. **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 © 2008, 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.

h. **Enable thin clone:** All clones are thin clones by default. For more information about thin clone, see Thin Cloning a Pluggable Database in the Oracle® Exadata Exascale User's Guide. If you specifically want a thick clone (full copy), you need to deselect this option.

i. **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.

11. Click **Clone**.

Using the API to Manage Pluggable Databases

Use various API features to help manage your pluggable databases on Oracle Exadata Database Service on Cloud@Customer.

For information about using the API and signing requests, see *REST APIs* and *Security Credentials*. For information about SDKs, see *Software Development Kits and Command Line Interface*.

Use these APIs to manage pluggable databases on Exadata Cloud@Customer systems.

- *ListPluggableDatabases*
- *GetPluggableDatabase*
- *StartPluggableDatabase*
- *StopPluggableDatabase*
- *CreatePluggableDatabase*
- *DeletePluggableDatabase*
- *LocalclonePluggableDatabase*
- *RemoteclonePluggabledatabase*

For the complete list of APIs for the Database service, see *Database Service API*.

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- ListPluggableDatabases
- GetPluggableDatabase
- StartPluggableDatabase
- StopPluggableDatabase

Copyright © 2008, 2026, Oracle and/or its affiliates.

- CreatePluggableDatabase
- DeletePluggableDatabase
- LocalclonePluggableDatabase
- RemoteclonePluggabledatabase
- Database Service API

Using the API to Clone a Pluggable Database

Clone a pluggable database (PDB) in the same database (CDB) as the source PDB or to a different database from the source PDB.

For information about using the API and signing requests, see *REST APIs* and *Security Credentials*. For information about SDKs, see *Software Development Kits and Command Line Interface*.

Use these APIs to manage pluggable databases on Exadata Cloud@Customer systems.

- *LocalclonePluggableDatabase*
- *RemoteclonePluggabledatabase*

For the complete list of APIs for the Database service, see *Database Service API*.

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- LocalclonePluggableDatabase
- RemoteclonePluggabledatabase
- Database Service API

Concurrently Create or Delete Pluggable Databases (PDBs)

You can now create or delete up to 10 PDBs concurrently, even when the container database (CDB) is in an updating state. However, you cannot create or delete PDBs while a CDB is in an updating state if other operations—excluding PDB creation or deletion—are modifying its metadata or structure.

The maximum number of CDBs and PDBs that can be created on a cluster is determined by the available memory on the VMs. By default, for Oracle Exadata Database Service on Cloud@Customer, each CDB is allocated 12.6 GB of memory (7.6 GB for SGA and 5 GB for PGA) if the VM has more than 60 GB of total memory. For VMs with 60 GB or less, 6.3 GB is allocated (3.8 GB for SGA and 2.5 GB for PGA).

In addition, Grid Infrastructure or ASM typically consumes 2 to 4 GB of memory. You can adjust the VM's memory allocation if necessary. It's important to note that the memory allocated to a CDB is shared among its PDBs. If a CDB has insufficient memory, PDB creation will fail.

Additionally, PDBs can now be created or deleted concurrently in a Data Guard environment. There are no restrictions on concurrent PDB creation or deletion in non-Data Guard environments.

Note:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Concurrent creation of PDBs in an Oracle Database configured with Data Guard and using a file-based wallet for TDE encryption keys (Oracle-managed keys) is also supported, however Oracle will serially create the PDBs on the standby database server in a serial manner. However, you can delete them concurrently.
- A PDB cannot be deleted if its creation is still in progress.

Operations you can perform in parallel:

- Create or delete CDBs in a VM cluster.
- Create a CDB in a VM cluster while the deletion of another CDB is in progress within the same cluster.
- Delete a CDB in a VM cluster while the creation of another CDB is in progress within the same cluster.
- In a CDB, delete a PDB (PDB1) while another PDB (PDB2) is being created.
- Create or delete additional PDBs if the CDB is in an updating state.
- Scale OCPU while a PDB is being created or deleted.

Related Topics

- Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel

Cost and Usage Attribution for Pluggable Databases (PDBs)

> **Note**
>
> It is supported only on Oracle Databases 19c and higher running in a multitenant deployment.

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

Prerequisites:

- `dbaastools`: (minimum version) 24.3.2
  - To check the version of the `dbaastools` rpm on the guest VM, run:

    ```
    rpm -qa | grep dbaastools
    ```

  - To update the `dbaastools` rpm on the guest VM, run:

    ```
    dbaascli admin updateStack
    ```

    Confirm you have the minimum version of `dbaastools` needed after you update the `dbaastools` rpm by running the `rpm -qa | grep dbaastools` command.
- `dbcsagent` needs to be running on the guest VM. Minimum version of `dbcsagent` needed is 23.3.2.

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - To check the version of the `dbcsagent` on the guest VM, run:

    ```
    rpm -qa | grep dbcs-agent-update
    ```

  - You will need to open a service request on My Oracle Support to update the `dbcsagent` on the guest VM.
  - To check the status of the `dbcsagent`, run:

    ```
    systemctl status dbcsagent
    ```

    Run `systemctl start dbcsagent` if the `dbcsagent` is not in active (running) state.

    Check the status of the agent again to confirm that it is running.

- Network setup: The Control Plane Server must be able to establish a connection with the endpoint specified in the Metering and Monitoring section of Table 3-2 in the Network Requirements for Oracle Exadata Database Service on Cloud@Customer.
- Generate Attributed Cost Analysis Report for Pluggable Databases
  Follow the steps below to view the attributed costs based on CPU utilization for all pluggable databases within a VM Cluster.

Generate Attributed Cost Analysis Report for Pluggable Databases

Follow the steps below to view the attributed costs based on CPU utilization for all pluggable databases within a VM Cluster.

1. Open the navigation menu and click **Billing & Cost Management**. Under **Cost Management**, click **Cost Analysis**.
2. From **Reports**, select one of the predefined reports, or use the default **Costs by Service report**.
3. Make your preferred query adjustments.
   a. From **Start/End Date (UTC)**, select a time period.
   b. From **Granularity**, select **Daily or Monthly**.
   c. From **Show**, select **Attributed cost**.
   d. From **Filters**, select **Tag**.
   In the resulting **Tag** dialog, select **orcl-cloud** as the tag with the key **parent_resource_id_1** equal to the OCID of the VM Cluster.
   e. From **Grouping dimensions**, select the preferred grouping dimension. For example, **Resource OCID**.
   The VM Cluster OCID is the parent of the CDBs it contains, and the CDB OCID is the parent OCID of the PDBs it contains.
   f. Click **Apply** to apply the changes and reload the chart and table with the selected filters.
   The generated report will show the attributed costs for all the PDBs in the VM Cluster.
4. After you have made changes, the currently selected predefined report name from the **Reports** menu changes to (edited).
5. If you're done making changes and want to save a new report, click **Save** as new report.
6. In the **Save as new report** dialog, enter the report name in the **Name** field. Avoid entering confidential information..

Copyright © 2008, 2026, Oracle and/or its affiliates.

7. Click **Save**.
   A notification is displayed that your report has been saved, and the report is also selected in the **Reports** menu.
8. If you didn't already apply your custom report settings, click **Apply** to view your changes. The new saved report is now available for future selection from the **Reports** menu under **Saved Reports**.

   For more information about generating a PDB attributed cost analysis report, see Cost Analysis.

Connect to an Oracle Database using Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users

You can configure Oracle Exadata Database Service on Cloud@Customer 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 Cloud@Customer to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier), or SSO tokens.
- Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database
  Review the prerequisites for Identity and Access Management (IAM) authentication on an Oracle Database.
- Enabling the Database and Clients for IAM Integration
  Follow the appropriate link below to configure IAM users to access your database.

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database

Learn to enable an Oracle Database instance on Oracle Exadata Database Service on Cloud@Customer to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier), or SSO tokens.

- About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database
  You can enable an Oracle Database instance to use Oracle Cloud Infrastructure (IAM) authentication and authorization for users.
- Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password 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
  You can enable an Oracle Database instance to use Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO tokens.

Copyright © 2008, 2026, Oracle and/or its affiliates.

About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database

You can enable an Oracle Database instance to use Oracle Cloud Infrastructure (IAM) authentication and authorization for users.

> **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 IAM integration is 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 Cloud@Customer supports the following:

- *IAM Database Password Authentication*
- *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 Cloud@Customer, see *Authenticating and Authorizing IAM Users for Oracle DBaaS Databases* in the Oracle Database 19c Security Guide and Oracle AI Database 26ai Security Guide.

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password 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 user name and IAM database password. An IAM database password is a different password than the Oracle Cloud Infrastructure Console password. Using an IAM user with the 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 Working with IAM Database User Names and Passwords.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication

You can enable an Oracle Database instance to use Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO tokens.

To access the database using IAM tokens, 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 a custom, non-default domain.

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. For more information about the clients supported for this type of IAM database token usage, see Supported Client Drivers for IAM Connections.
- 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 the default or specified file location.
- 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.
- 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.
- IAM users and OCI applications can request a database token from IAM with several methods, including using an API key. See Configuring a Client Connection for SQL*Plus That Uses an IAM Token for an example. See Authenticating and Authorizing IAM Users for Oracle DBaaS Databases for a description of other methods such as using a delegation token within an OCI cloud shell.

In previous releases, you could only use the IAM username and database password to get a password verifier from IAM. Getting a token with these credentials is more secure than getting a password verifier because a password verifier is considered sensitive. Using a token means that you do not need to pass or use the verifier. Applications cannot pass a token that was retrieved by the IAM user name and password through the database client API. Only the database client can retrieve this type of token. A database client can only retrieve a database token using the IAM user name and IAM database password.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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.
- Configure a Network Connection to OCI
  Configure a network connection to OCI to be able to make calls to OCI IAM for the database instances on Oracle Exadata Database Service on Cloud@Customer to accept IAM database access tokens (`db-tokens`), or get IAM database password verifiers.
- Configure Proxy Settings
  Configure network proxy settings in your environment to allow the database to access OCI IAM. Replace the network proxy URL `http://www-proxy.example.com:80/` and the database name given in the example with yours.
- 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-C@C service instance must be stored under the `WALLET_ROOT` location. Create a TLS directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`.

Disable External Authentication Scheme

Review the prerequisites for enabling IAM user access to Oracle Database.

If the database is enabled for another external authentication scheme, verify that you want to use IAM on the Oracle Database instance. There can only be one external authentication scheme enabled at any given time.

If you want to use IAM and another external authentication scheme is enabled, you must first disable the other external authentication scheme.

Configure a Network Connection to OCI

Configure a network connection to OCI to be able to make calls to OCI IAM for the database instances on Oracle Exadata Database Service on Cloud@Customer to accept IAM database access tokens (`db-tokens`), or get IAM database password verifiers.

1. Consult your ExaDB-C@C administrator to determine the OCI region assigned to your ExaDB-C@C installation.
2. Determine the OCI IAM endpoint for that OCI region. For more information, see *Identity and Access Management Service API*.
3. Find the port number for **Identity Service for name resolution of Oracle operators**. For more information, see *Table 3-2 Ports to Open for Control Plane Connectivity* in *Network Requirements for Oracle Exadata Database Service on Cloud@Customer*.

   For example, if your OCI region is **Phoenix**, then open port **443** to `https://identity.us-phoenix-1.oci.oraclecloud.com`.
4. Configure your network to open this connection.

For more information on troubleshooting login failures, see Troubleshooting IAM Connections.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Identity and Access Management Service API
- Network Requirements for Oracle Exadata Database Service on Cloud@Customer

Configure Proxy Settings

Configure network proxy settings in your environment to allow the database to access OCI IAM. Replace the network proxy URL `http://www-proxy.example.com:80/` and the database name given in the example with yours.

1. Log in to the host operating system.
2. Set the proxy environment variables.

```
srvctl setenv database -db exampledbname -env "https_proxy=http://www-
proxy.example.com:80/"
srvctl setenv database -db exampledbname -env "http_proxy=http://www-
proxy.example.com:80/"
```

3. Stop the database and verify that the variables have been set:

```
$ srvctl stop database -db exampledbname
```

```
$ srvctl getenv database -db exampledbname
http_proxy=http://www-proxy.example.com:80/
https_proxy=http://www-proxy.example.com:80/
```

4. Restart the database.

```
$ srvctl start database -db exampledbname
```

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-C@C 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

Copyright © 2008, 2026, Oracle and/or its affiliates.

wallet or installed in the client system default certificate store (Windows and Linux only). When the session is established, the database client will check to see that the certificate sent over by the database server has been signed by the same root certificate.

A Well-Known Certificate Authority

Using a commonly known root certificate authority has some advantages in that the root certificate is most likely already stored in the client system default certificate store. There is no extra step for the client to store the root certificate if it is a common root certificate. The disadvantage is that this normally has a cost associated with it.

One-Way TLS

In the standard TLS session, only the server provides a certificate to the client to authenticate itself. The client doesn't need to have a separate client certificate to authenticate itself to the server (similar to how HTTPS sessions are established). While the database requires a wallet to store the server certificate, the only thing the client needs to have is the root certificate used to sign the server certificate.

Two-Way TLS (also called Mutual TLS, mTLS)

In mTLS, both the client and server have identity certificates that are presented to each other. In most cases, the same root certificate will have signed both of these certificates so the same root certificate can be used with the database server and client to authenticate the other certificate. mTLS is sometimes used to authenticate the user since the user identity is authenticated by the database server through the certificate. This is not necessary for passing IAM tokens but can be used when passing IAM tokens.

Client with a Wallet

A client wallet is mandatory when using mTLS to store the client certificate. However, the root certificate can be stored either in the same wallet or in the system default certificate store.

A Client without a Wallet

Clients can be configured without a wallet when using TLS under these conditions: 1) One-way TLS is being configured where the client does not have its own certificate and 2) the root certificate that signed the database server certificate is stored in the system default certificate store. The root certificate would most likely already be there if the server certificate is signed by a common certificate authority. If it's a self-signed certificate, then the root certificate would need to be installed in the system default certificate store to avoid using a client wallet.

For details on how to configure TLS between the database client and database server including the options described above, see Configuring Transport Layer Security Encryption in the *Oracle Database Security Guide*.

If you choose to use self-signed certificates and for additional wallet related tasks, Managing Oracle Database wallets and certificates in the *Oracle Database Security Guide*.

Enabling the Database and Clients for IAM Integration

Follow the appropriate link below to configure IAM users to access your database.

For complete details about the architecture for using IAM users on Oracle Exadata Database Service on Dedicated Infrastructure, see *Authenticating and Authorizing IAM Users for Oracle DBaaS Databases* in the Oracle Database 19c Security Guide and Oracle AI Database 26ai Security Guide.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Authenticating and Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Cloud@Customer

An Oracle Database can be configured for Microsoft Azure users of Microsoft Entra ID to connect using single sign-on authentication.

- About Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Cloud@Customer
  Users for Oracle Exadata Database Service on Cloud@Customer can be centrally managed in an MS-EI service.
- Configuring the Oracle Database for Microsoft Entra ID (MS-EI) Integration
  The MS-EI integration with the Oracle Database instance requires the database to be registered with MS-EI so that the database can request the MS-EI public key.

About Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Cloud@Customer

Users for Oracle Exadata Database Service on Cloud@Customer can be centrally managed in an MS-EI service.

The Oracle Database integration with MS-EI is supported for on-premises databases and most Oracle OCI DBaaS platforms.

The instructions for configuring MS-EI use the term "Oracle Database" to encompass these environments.

This type of integration enables the MS-EI user to access an Oracle Exadata Database Service on Cloud@Customer instance. MS-EI users and applications can log in with MS-EI Single Sign On (SSO) credentials to get an MS-EI OAuth2 access token to send to the database.

The administrator creates and configures the application registration (app registration) of the Oracle Exadata Database Service on Cloud@Customer instance with MS-EI. The administrator also creates application (app) roles for the database app registration in MS-EI, and assigns these roles to MS-EI users, groups, and applications. These app roles will be mapped to the database global schemas and global roles. An MS-EI principal that is assigned to an app role will be mapped to either a database global schema or database global role. An Oracle global schema can also be mapped exclusively to an MS-EI user. When the principal is a guest user or service principal, they can only be mapped to the database schema through an MS-EI app role. An Oracle global role can only be mapped to an MS-EI app role.

Tools and applications that are updated to support MS-EI tokens can authenticate users directly with MS-EI and pass the database access token to the Oracle Exadata Database Service on Cloud@Customer instance. You can configure existing database tools such as SQL*Plus to use an MS-EI token from a file location or get the token directly from MS-EI. When using a utility to get the token to pass to the database client driver through a file location, MS-EI tokens can be retrieved using tools like Microsoft PowerShell or Azure CLI and put into a file location. An MS-EI OAuth2 database access token is a bearer token with an expiration time. The Oracle Database client driver will ensure that the token is in a valid format and that it has not expired before passing it to the database. The token is scoped for the database. Assigned app roles for the Azure AD principal are included as part of the access token. The directory location for the MS-EI token should only have enough permission for the user to write the token file to the location and the database client to retrieve these files (for example, just

Copyright © 2008, 2026, Oracle and/or its affiliates.

read and write by the process user). Because the token allows access to the database, it should be protected within the file system.

MS-EI users can request a token as a client registered with MS-EI app registration by using methods such as the following:

- Entering the MS-EI credentials into an MS-EI authentication screen with or without multi-factor authentication

Oracle Exadata Database Service on Cloud@Customer supports the following MS-EI authentication flows:

- Interactive flow (authorization code), which is used when a browser can be used to enter credentials for the user
- Client credentials, which are for applications that connect as themselves (and not the end-user)
- On-Behalf-Of (OBO), where an application requests an access token on behalf of a logged-in user to send to the database
- ROPC is also supported for test and development environments

Oracle Exadata Database Service on Cloud@Customer accepts tokens representing the following MS-EI principals:

- MS-EI user, who is registered user in the MS-EI tenancy
- Guest user, who is registered as a guest user in the MS-EI tenancy
- Service, which is the registered application connecting to the database as itself with the client credential flow (connection pool use case)

Configuring the Oracle Database for Microsoft Entra ID (MS-EI) Integration

The MS-EI integration with the Oracle Database instance requires the database to be registered with MS-EI so that the database can request the MS-EI public key.

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

- Authenticating and Authorizing Microsoft Azure Active Directory Users for Oracle Databases in the Oracle Database 19c Security Guide.
- Authenticating and Authorizing Microsoft Azure Users for Oracle Databases in the Oracle AI Database 26ai Security Guide.
- Prerequisites for Microsoft Entra ID (MS-EI) Authentication
- Configure TLS to Use Microsoft Entra ID (MS-EI) tokens

Prerequisites for Microsoft Entra ID (MS-EI) Authentication

The MS-EI integration with the Oracle Database on Oracle Exadata Database Service on Cloud@Customer requires:

1. The Oracle Database to be version 19.18 or higher.
2. Connectivity to the database on TLS port 2484. Non TLS connections are not supported.
3. The Oracle Database to be registered with MS-EI.

Copyright © 2008, 2026, Oracle and/or its affiliates.

4. Users and applications that need to request an MS-EI token must also be able to have network connectivity to MS-EI. You may need to configure a proxy setting for the connection.
5. Outbound network connectivity to MS-EI so that the database can request the MS-EI public key.
    a. Ensure that the `ORACLE_SID`, `ORACLE_HOME`, and `PATH` variables of the database are set correctly.
    b. Run the following `srvctl` commands as `oracle` user on the operating system:
        - `$ srvctl setenv database -db <exampledbname> -env "https_proxy=http://www-proxy.example.com:80/"`
        - `$ srvctl setenv database -db <exampledbname> -env "http_proxy=http://www-proxy.example.com:80/"`
        - `$ srvctl stop database -db <exampledbname>`
        - `$ srvctl start database -db <exampledbname>`
        - `$ srvctl getenv database -db <exampledbname>`
          `http_proxy=http://www-proxy.example.com:80/`
          `https_proxy=http://www-proxy.example.com:80/`
          `TNS_ADMIN=<exampledbhomename>/network/admin/<exampledbname>`
    c. The HTTP proxy settings must also be set in the ExaDB-C@C networking. These settings are defined by your fleet administrator while creating the Exadata infrastructure as described in Using the Console to Provision Oracle Exadata Database Service on Cloud@Customer.

> **Note**
>
> - The network configuration including the HTTP Proxy can only be edited until the Exadata infrastructure is in **Requires Activation** state. Once it is activated, you cannot edit those settings.
> - Setting up an HTTP Proxy for an already provisioned Exadata infrastructure needs a Service Request (SR) in My Oracle Support. See Create a Service Request in My Oracle Support for details.

Related Topics

- Task 1: Create the service gateway

Configure TLS to Use Microsoft Entra ID (MS-EI) tokens

When sending MS-EI tokens from the database client to the database server, a TLS connection must be established. The TLS wallet with the database certificate for the ExaDB-C@C service instance must be stored under the `WALLET_ROOT` location. Create a `tls` directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

When configuring TLS between the database client and server there are several options to consider.

- Using a self-signed database server certificate vs a database server certificate signed by a commonly known certificate authority
- One-way TLS (TLS) vs Mutual or two-way TLS (mTLS)
- Client with or without a wallet

Self-Signed Certificate

Using a self-signed certificate is a common practice for internally facing IT resources since you can create these yourself and it's free. The resource (in our case, the database server) will have a self-signed certificate to authenticate itself to the database client. The self-signed certificate and root certificate will be stored in the database server wallet. For the database client to be able to recognize the database server certificate, a copy of the root certificate will also be needed on the client. This self-created root certificate can be stored in a client-side wallet or installed in the client system default certificate store (Windows and Linux only). When the session is established, the database client will check to see that the certificate sent over by the database server has been signed by the same root certificate.

A Well-Known Certificate Authority

Using a commonly known root certificate authority has some advantages in that the root certificate is most likely already stored in the client system default certificate store. There is no extra step for the client to store the root certificate if it is a common root certificate. The disadvantage is that this normally has a cost associated with it.

One-Way TLS

In the standard TLS session, only the server provides a certificate to the client to authenticate itself. The client doesn't need to have a separate client certificate to authenticate itself to the server (similar to how HTTPS sessions are established). While the database requires a wallet to store the server certificate, the only thing the client needs to have is the root certificate used to sign the server certificate.

Two-Way TLS (also called Mutual TLS, mTLS)

In mTLS, both the client and server have identity certificates that are presented to each other. In most cases, the same root certificate will have signed both of these certificates so the same root certificate can be used with the database server and client to authenticate the other certificate. mTLS is sometimes used to authenticate the user since the user identity is authenticated by the database server through the certificate. This is not necessary for passing IAM tokens but can be used when passing IAM tokens.

Client with a Wallet

A client wallet is mandatory when using mTLS to store the client certificate. However, the root certificate can be stored either in the same wallet or in the system default certificate store.

A Client without a Wallet

Clients can be configured without a wallet when using TLS under these conditions: 1) One-way TLS is being configured where the client does not have its own certificate and 2) the root certificate that signed the database server certificate is stored in the system default certificate store. The root certificate would most likely already be there if the server certificate is signed by a common certificate authority. If it's a self-signed certificate, then the root certificate would need to be installed in the system default certificate store to avoid using a client wallet.

For details on how to configure TLS between the database client and database server including the options described above, see:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Configuring Transport Layer Security Authentication in the Oracle Database 19c Security Guide.
- Configuring Transport Layer Security Encryption in the Oracle AI Database 26ai Security Guide.

If you choose to use self-signed certificates and for additional wallet related tasks, see:

- Managing Public Key Infrastructure (PKI) Elements in the Oracle Database 19c Security Guide.
- Configuring TLS Connection With a Client Wallet in the Oracle AI Database 26ai Security Guide.

Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer

Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Cloud@Customer.

- Backup Destinations
  A backup destination is a target for sending database backups. The target may be a local filesystem, Object store, NFS mount or ZDLRA.
- Oracle Database Backup Methods in Exadata Cloud
  Oracle Exadata Database Service on Cloud@Customer offers two approaches to configure and take backups as a recommended solution: Oracle Managed Backup and User Configured Backup.
- Configuring and Customizing Backups with dbaascli
  In addition to the console-based automated backup option, use dbaascli to configure and customize backups.
- Creating an On-Demand Backup by Using dbaascli
  You can use the `dbaascli` to create an on-demand backup of a complete database or an individual pluggable database (PDB):
- Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management
- Customizing Real Time Redo Transport (RTRT) Behavior for Recovery Appliance Backups
- Alternative Backup Methods
  Learn about alternative backup methods that are available in addition to the OCI Console.
- Standalone Backup

Backup Destinations

A backup destination is a target for sending database backups. The target may be a local filesystem, Object store, NFS mount or ZDLRA.

- About Managing Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  For backups, you can either use the Oracle Exadata Database Service on Cloud@Customer backup facility, or you can configure a backup location on a location you manage.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Prerequisites for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  To configure backup destinations on a Zero Data Loss Recovery Appliance location, or an NFS backup location, review the prerequisites.
- Using the Console for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  Learn how to use the console to create, edit, move, and terminate a backup destination for your infrastructure for Oracle Exadata Database Service on Cloud@Customer.
- Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Backup Destinations
  Review the list of API calls to manage your Oracle Exadata Database Service on Cloud@Customer backup destinations.

About Managing Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

For backups, you can either use the Oracle Exadata Database Service on Cloud@Customer backup facility, or you can configure a backup location on a location you manage.

Oracle Exadata Database Service on Cloud@Customer provides a backup facility, which you can configure individually on each database.

See: *Managing Databases on Exadata Cloud@Customer* and *Managing Database Backup and Recovery on Exadata Cloud@Customer*.

If you want to store backups on a Recovery Appliance, or on a network file storage (NFS) location that you manage, then you must first create a backup destination. Each backup destination defines the properties that are required to connect to the Recovery Appliance or NFS location, and each backup destination must be accessible in your data center from the VM cluster nodes.

The Oracle Exadata Database Service on Cloud@Customer backup facility can also store backups on Oracle Cloud Infrastructure object storage, or on local Exadata storage on your Oracle Exadata Database Service on Cloud@Customer system. However, you do not need to create a backup destination for any of these other locations. Instead, applicable options for backup to cloud object storage or local Exadata storage are available directly when you create a database.

> **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.

Related Topics

- Zero Data Loss Recovery Appliance
- Manage Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems
  Learn to manage Oracle Database homes on Oracle Exadata Database Service on Cloud@Customer.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console to Create a Backup Destination
  To create a backup destination, be prepared to provide values for the backup destination configuration.
- Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer
  Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Cloud@Customer.

Prerequisites for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

To configure backup destinations on a Zero Data Loss Recovery Appliance location, or an NFS backup location, review the prerequisites.

> **Note**
>
> Starting with dbaascli release 25.3.1.0.0, the Data Guard broker is mandatory when configuring Zero Data Loss Recovery Appliance as a backup destination in Data Guard-enabled databases.

- For a Zero Data Loss Recovery Appliance backup destination:
  - The appliance must be configured with a virtual private catalog (VPC) user, which is used for taking the backups.
  - The appliance must be configured with the unique database name of the database being backed up, and a mapping to the VPC user.
  - The appliance must be accessible from the Oracle Exadata Database Service on Cloud@Customer system using the Oracle Net Services connection string, which is provided by the Zero Data Loss Recovery Appliance administrator.
- For an NFS backup destination:
  - Oracle Exadata Database Service on Cloud@Customer non-autonomous databases:
    * You must mount the NFS server location to a local mount point directory on each node in the VM cluster.
    * The local mount point directory and the NFS server must be identical across all nodes in the cluster.
    * You must ensure that the NFS mount is maintained continuously on all of the VM cluster nodes.
    * The NFS-mounted file system must be readable and writable by the `oracle` operating system user on all of the VM cluster nodes.
  - Oracle Autonomous AI Database on Exadata Cloud@Customer:

    > **Note**
    >
    > Oracle Autonomous AI Database on Exadata Cloud@Customer supports NFS versions 3 and 4. Mounts with both versions enabled, however, default to version 4.

Copyright © 2008, 2026, Oracle and/or its affiliates.

* To ensure that the Autonomous VM cluster can access the NFS server over the (preferred) Backup Network, enter valid Backup Network IP addresses while configuring the VM Cluster Network.
  If the backup destination is not accessible over the backup network but accessible over the client network, the backup traffic will be routed over the client network.
* The NFS-mounted file system must be readable and writable by the `oracle` operating system user on all of the VM cluster nodes.
* If permissions are being controlled at the user level, then the `uid:gid` of the `oracle` user for the Autonomous VM cluster is 1001:1001.
* For NFSv4 servers that restrict port range, you must use the insecure option to enable clients other than an Administrator user to connect to the NFS server.

Using the Console for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

Learn how to use the console to create, edit, move, and terminate a backup destination for your infrastructure for Oracle Exadata Database Service on Cloud@Customer.

- Using the Console to Create a Backup Destination
  To create a backup destination, be prepared to provide values for the backup destination configuration.
- Using the Console to Edit a Backup Destination
  To edit a backup destination, be prepared to provide values for the backup destination configuration.
- Using the Console to Move a Backup Destination to Another Compartment
  To move a backup destination, be prepared to provide values for the backup destination configuration.
- Using the Console to Delete a Backup Destination
  To delete a backup destination, be prepared to provide values for the backup destination configuration.

Using the Console to Create a Backup Destination

To create a backup destination, be prepared to provide values for the backup destination configuration.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Backup Destinations**.
4. Click **Create Backup Destination**.
5. Provide the requested information in the **Create Backup Destination** page:
   a. Choose a compartment.

      From the list of available compartments, choose the compartment that you want to contain the backup destination.

   b. Name your backup destination.

      Specify a user-friendly name that you can use to identify the backup destination. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the backup destination.

Copyright © 2008, 2026, Oracle and/or its affiliates.

c. Choose either a Zero Data Loss Recovery Appliance or a network file system (NFS) backup destination.

> **Note**
>
> You can also set OCI Object Store as a backup destination. However, you cannot set it from this screen. You can configure OCI Object Store as a backup destination when creating a database. For more information, see *Backup Destination Type* in *Using the Console to Create a Database*.

Select **Recovery Appliance** or **Network Storage (NFS)**.

- If you select **Recovery Appliance**, then you must also specify the following for Zero Data Loss Recovery Appliance:
  - **Provide the Recovery Appliance connection string:** Specify the Oracle Net Services connection string that connects to the appliance. This information is typically provided by the Zero Data Loss Recovery Appliance administrator.

    > **Note**
    >
    > ADB customers must ensure to provide the connection string that connects to the recovery appliance in an Oracle "easy connect" string format, as shown below:
    >
    > `host:port/service name`
    >
    > Where *host* is the SCAN hostname of the Zero Data Loss Recovery Appliance.

  - **Provide the Virtual Private Catalog (VPC) Users:** Provide a VPC user name for connecting to the Zero Data Loss Recovery Appliance. You can specify multiple VPC user names in case you want to use the appliance as a backup destination for multiple databases. This information is typically provided by the Zero Data Loss Recovery Appliance administrator.
- If you select **Network Storage (NFS)**, then you must also specify the following:
  - **Self-mount for non-autonomous databases:**

    Provide the local NFS mount point path: Specify the local directory path on each VM cluster node where the NFS server location is mounted. The local directory path and the NFS server location must each be the same across all of the VM cluster nodes.
  - **Auto-mount for Autonomous Databases:**

    Use this destination for Autonomous Databases:
    - **NFS server:** Specify the IP address of the NFS server. Optionally, you can specify up to four IP addresses. These are just extra paths to the same NFS server. Note that Oracle supports only one NFS device for backups.
    - **NFS export share:** Specify the directory path where the exported file system is mounted.

d. Configure Advanced Options.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Tags:** (Optional) You can choose to apply 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, refer to information about 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.

6. Click **Create Backup Destination**.

   The Backup Destination Details page displays the newly created backup destination.

Related Topics

- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.
- Resource Tags

Using the Console to Edit a Backup Destination

To edit a backup destination, be prepared to provide values for the backup destination configuration.

You can only edit a backup destination if it is not currently associated with database.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to edit.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to edit.

   The Backup Destination Details page displays information about the selected backup destination.
5. Click **Edit**.
6. Use the **Edit Backup Destination** dialog to edit the backup destination attributes:

> **Note**
>
> You cannot edit a Backup Destination if there is already a database attached to it.

- If you are editing a Zero Data Loss Recovery Appliance backup destination:
  - **Provide the Recovery Appliance connection string:** Specify the Oracle Net Services connection string that connects to the Recovery Appliance. This information is typically provided by the Recovery Appliance administrator.
  - **Provide the Virtual Private Catalog (VPC) Users:** Provide a VPC user name for connecting to the Recovery Appliance. You can specify multiple VPC user names in case you want to use the Recovery Appliance as a backup destination for multiple databases. This information is typically provided by the Recovery Appliance administrator.
- If you are editing an NFS backup destination:
  - **Self-mount for non-autonomous databases:**

Copyright © 2008, 2026, Oracle and/or its affiliates.

Provide the local NFS mount point path: Specify the local directory path on each VM cluster node where the NFS server location is mounted. The local directory path and the NFS server location must each be the same across all of the VM cluster nodes.

- **Auto-mount for Autonomous Databases:**

  Use this destination for Autonomous Databases:

  * **NFS server:** Specify the IP address of the NFS server. Optionally, you can specify up to four IP addresses.
  * **NFS export share:** Specify the directory path where the exported file system is mounted.

7. Click **Save Changes**.

Using the Console to Move a Backup Destination to Another Compartment

To move a backup destination, be prepared to provide values for the backup destination configuration.

You can change the compartment that contains your backup destination by moving it.

When you move a backup destination, the compartment change does not affect other associated resources. These other resources, such as the associated databases, remain in their current compartment.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to move.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to move.

   The Backup Destination Details page displays information about the selected backup destination.
5. Click **Move Resource**.
6. In the resulting dialog, choose the new compartment for the backup destination and click **Move Resource**.

Using the Console to Delete a Backup Destination

To delete a backup destination, be prepared to provide values for the backup destination configuration.

Before you can delete a backup destination, you must ensure that it is not associated with any databases.

Deleting a backup destination:

- Does not remove any residual backups that are left in the backup destination
- Removes all references to the deleted backup destination from the Cloud Control Plane

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to delete.

Copyright © 2008, 2026, Oracle and/or its affiliates.

3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to delete.

   The Backup Destination Details page displays information about the selected backup destination.
5. Click **Delete**.
6. In the resulting dialog, enter the backup destination name and click **Delete Backup Destination** to confirm the action.

Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Backup Destinations

Review the list of API calls to manage your Oracle Exadata Database Service on Cloud@Customer backup destinations.

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 Cloud@Customer backup destinations:

- `CreateBackupDestination`
- `DeleteBackupDestination`
- `GetBackupDestination`
- `ListBackupDestination`
- `UpdateBackupDestination`
- `ChangeBackupDestinationCompartment`

For the complete list of APIs, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- CreateBackupDestination
- DeleteBackupDestination
- GetBackupDestination
- ListBackupDestination
- UpdateBackupDestination
- ChangeBackupDestinationCompartment
- Database Service API

Oracle Database Backup Methods in Exadata Cloud

Oracle Exadata Database Service on Cloud@Customer offers two approaches to configure and take backups as a recommended solution: Oracle Managed Backup and User Configured Backup.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Oracle Managed Backup
  Database backups are managed entirely by the Oracle Exadata Database Service on Cloud@Customer service based on a one-time configuration.
- User Configured Backup
  It is the user's responsibility to configure and execute backup operations using `dbaascli` according to their preferences.

Oracle Managed Backup

Database backups are managed entirely by the Oracle Exadata Database Service on Cloud@Customer service based on a one-time configuration.

Once configured, you need not perform any maintenance such as backup scheduling and deletion of the backups. Oracle manages the backups through well-defined workflows. Certain backup configuration parameters are not fully integrated with the Oracle Managed Backup workflow. If you want to set any of those parameters for the backups, then you can use `dbaascli database backup –configure` to set them. For more information, see *Configuring Database for Backup*.

- Backup Destinations
  A backup destination is a target for sending database backups. The target may be a local filesystem, Object store, NFS mount or ZDLRA.
- Automatic Backup and Recovery

Related Topics

- Configuring Database for Backup
  Once the backup destination is set up and available, you can use the dbaascli utility to configure the database with backup destination and associated configuration parameters, for example, Backup Retention Recovery Window, Backup Scheduling, Archivelog Scheduling, and so on.

Backup Destinations

A backup destination is a target for sending database backups. The target may be a local filesystem, Object store, NFS mount or ZDLRA.

- About Managing Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  For backups, you can either use the Oracle Exadata Database Service on Cloud@Customer backup facility, or you can configure a backup location on a location you manage.
- Prerequisites for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  To configure backup destinations on a Zero Data Loss Recovery Appliance location, or an NFS backup location, review the prerequisites.
- Using the Console for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer
  Learn how to use the console to create, edit, move, and terminate a backup destination for your infrastructure for Oracle Exadata Database Service on Cloud@Customer.
- Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Backup Destinations
  Review the list of API calls to manage your Oracle Exadata Database Service on Cloud@Customer backup destinations.

Copyright © 2008, 2026, Oracle and/or its affiliates.

About Managing Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

For backups, you can either use the Oracle Exadata Database Service on Cloud@Customer backup facility, or you can configure a backup location on a location you manage.

Oracle Exadata Database Service on Cloud@Customer provides a backup facility, which you can configure individually on each database.

See: *Managing Databases on Exadata Cloud@Customer* and *Managing Database Backup and Recovery on Exadata Cloud@Customer*.

If you want to store backups on a Recovery Appliance, or on a network file storage (NFS) location that you manage, then you must first create a backup destination. Each backup destination defines the properties that are required to connect to the Recovery Appliance or NFS location, and each backup destination must be accessible in your data center from the VM cluster nodes.

The Oracle Exadata Database Service on Cloud@Customer backup facility can also store backups on Oracle Cloud Infrastructure object storage, or on local Exadata storage on your Oracle Exadata Database Service on Cloud@Customer system. However, you do not need to create a backup destination for any of these other locations. Instead, applicable options for backup to cloud object storage or local Exadata storage are available directly when you create a database.

> **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.

Related Topics

- Zero Data Loss Recovery Appliance
- Manage Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems
  Learn to manage Oracle Database homes on Oracle Exadata Database Service on Cloud@Customer.
- Using the Console to Create a Backup Destination
  To create a backup destination, be prepared to provide values for the backup destination configuration.
- Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer
  Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Cloud@Customer.

Prerequisites for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

To configure backup destinations on a Zero Data Loss Recovery Appliance location, or an NFS backup location, review the prerequisites.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Starting with dbaascli release 25.3.1.0.0, the Data Guard broker is mandatory when configuring Zero Data Loss Recovery Appliance as a backup destination in Data Guard-enabled databases.

- For a Zero Data Loss Recovery Appliance backup destination:
  - The appliance must be configured with a virtual private catalog (VPC) user, which is used for taking the backups.
  - The appliance must be configured with the unique database name of the database being backed up, and a mapping to the VPC user.
  - The appliance must be accessible from the Oracle Exadata Database Service on Cloud@Customer system using the Oracle Net Services connection string, which is provided by the Zero Data Loss Recovery Appliance administrator.
- For an NFS backup destination:
  - Oracle Exadata Database Service on Cloud@Customer non-autonomous databases:
    * You must mount the NFS server location to a local mount point directory on each node in the VM cluster.
    * The local mount point directory and the NFS server must be identical across all nodes in the cluster.
    * You must ensure that the NFS mount is maintained continuously on all of the VM cluster nodes.
    * The NFS-mounted file system must be readable and writable by the `oracle` operating system user on all of the VM cluster nodes.
  - Oracle Autonomous AI Database on Exadata Cloud@Customer:

    > **Note**
    >
    > Oracle Autonomous AI Database on Exadata Cloud@Customer supports NFS versions 3 and 4. Mounts with both versions enabled, however, default to version 4.

    * To ensure that the Autonomous VM cluster can access the NFS server over the (preferred) Backup Network, enter valid Backup Network IP addresses while configuring the VM Cluster Network.
      If the backup destination is not accessible over the backup network but accessible over the client network, the backup traffic will be routed over the client network.
    * The NFS-mounted file system must be readable and writable by the `oracle` operating system user on all of the VM cluster nodes.
    * If permissions are being controlled at the user level, then the `uid:gid` of the `oracle` user for the Autonomous VM cluster is 1001:1001.
    * For NFSv4 servers that restrict port range, you must use the insecure option to enable clients other than an Administrator user to connect to the NFS server.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console for Backup Destinations for Oracle Exadata Database Service on Cloud@Customer

Learn how to use the console to create, edit, move, and terminate a backup destination for your infrastructure for Oracle Exadata Database Service on Cloud@Customer.

- Using the Console to Create a Backup Destination
  To create a backup destination, be prepared to provide values for the backup destination configuration.
- Using the Console to Edit a Backup Destination
  To edit a backup destination, be prepared to provide values for the backup destination configuration.
- Using the Console to Move a Backup Destination to Another Compartment
  To move a backup destination, be prepared to provide values for the backup destination configuration.
- Using the Console to Delete a Backup Destination
  To delete a backup destination, be prepared to provide values for the backup destination configuration.

Using the Console to Create a Backup Destination

To create a backup destination, be prepared to provide values for the backup destination configuration.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Backup Destinations**.
4. Click **Create Backup Destination**.
5. Provide the requested information in the **Create Backup Destination** page:
   a. Choose a compartment.

      From the list of available compartments, choose the compartment that you want to contain the backup destination.

   b. Name your backup destination.

      Specify a user-friendly name that you can use to identify the backup destination. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the backup destination.

   c. Choose either a Zero Data Loss Recovery Appliance or a network file system (NFS) backup destination.

      > **Note**
      >
      > You can also set OCI Object Store as a backup destination. However, you cannot set it from this screen. You can configure OCI Object Store as a backup destination when creating a database. For more information, see *Backup Destination Type* in *Using the Console to Create a Database*.

      Select **Recovery Appliance** or **Network Storage (NFS)**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- If you select **Recovery Appliance**, then you must also specify the following for Zero Data Loss Recovery Appliance:
  - **Provide the Recovery Appliance connection string:** Specify the Oracle Net Services connection string that connects to the appliance. This information is typically provided by the Zero Data Loss Recovery Appliance administrator.

    > **Note**
    >
    > ADB customers must ensure to provide the connection string that connects to the recovery appliance in an Oracle "easy connect" string format, as shown below:
    >
    > *host*:*port*/*service name*
    >
    > Where *host* is the SCAN hostname of the Zero Data Loss Recovery Appliance.

  - **Provide the Virtual Private Catalog (VPC) Users:** Provide a VPC user name for connecting to the Zero Data Loss Recovery Appliance. You can specify multiple VPC user names in case you want to use the appliance as a backup destination for multiple databases. This information is typically provided by the Zero Data Loss Recovery Appliance administrator.
- If you select **Network Storage (NFS)**, then you must also specify the following:
  - **Self-mount for non-autonomous databases:**

    Provide the local NFS mount point path: Specify the local directory path on each VM cluster node where the NFS server location is mounted. The local directory path and the NFS server location must each be the same across all of the VM cluster nodes.
  - **Auto-mount for Autonomous Databases:**

    Use this destination for Autonomous Databases:
    - **NFS server:** Specify the IP address of the NFS server. Optionally, you can specify up to four IP addresses. These are just extra paths to the same NFS server. Note that Oracle supports only one NFS device for backups.
    - **NFS export share:** Specify the directory path where the exported file system is mounted.

d. Configure Advanced Options.

- **Tags:** (Optional) You can choose to apply 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, refer to information about 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.

6. Click **Create Backup Destination**.

   The Backup Destination Details page displays the newly created backup destination.

Related Topics

- Using the Console to Create a Database
  To create an Oracle Database with the console, use this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Resource Tags

Using the Console to Edit a Backup Destination

To edit a backup destination, be prepared to provide values for the backup destination configuration.

You can only edit a backup destination if it is not currently associated with database.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to edit.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to edit.

   The Backup Destination Details page displays information about the selected backup destination.
5. Click **Edit**.
6. Use the **Edit Backup Destination** dialog to edit the backup destination attributes:

   > **Note**
   >
   > You cannot edit a Backup Destination if there is already a database attached to it.

   - If you are editing a Zero Data Loss Recovery Appliance backup destination:
     - **Provide the Recovery Appliance connection string:** Specify the Oracle Net Services connection string that connects to the Recovery Appliance. This information is typically provided by the Recovery Appliance administrator.
     - **Provide the Virtual Private Catalog (VPC) Users:** Provide a VPC user name for connecting to the Recovery Appliance. You can specify multiple VPC user names in case you want to use the Recovery Appliance as a backup destination for multiple databases. This information is typically provided by the Recovery Appliance administrator.
   - If you are editing an NFS backup destination:
     - **Self-mount for non-autonomous databases:**

       Provide the local NFS mount point path: Specify the local directory path on each VM cluster node where the NFS server location is mounted. The local directory path and the NFS server location must each be the same across all of the VM cluster nodes.
     - **Auto-mount for Autonomous Databases:**

       Use this destination for Autonomous Databases:

       * **NFS server:** Specify the IP address of the NFS server. Optionally, you can specify up to four IP addresses.
       * **NFS export share:** Specify the directory path where the exported file system is mounted.
7. Click **Save Changes**.

Using the Console to Move a Backup Destination to Another Compartment

Copyright © 2008, 2026, Oracle and/or its affiliates.

To move a backup destination, be prepared to provide values for the backup destination configuration.

You can change the compartment that contains your backup destination by moving it.

When you move a backup destination, the compartment change does not affect other associated resources. These other resources, such as the associated databases, remain in their current compartment.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to move.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to move.

   The Backup Destination Details page displays information about the selected backup destination.
5. Click **Move Resource**.
6. In the resulting dialog, choose the new compartment for the backup destination and click **Move Resource**.

Using the Console to Delete a Backup Destination

To delete a backup destination, be prepared to provide values for the backup destination configuration.

Before you can delete a backup destination, you must ensure that it is not associated with any databases.

Deleting a backup destination:

- Does not remove any residual backups that are left in the backup destination
- Removes all references to the deleted backup destination from the Cloud Control Plane

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to delete.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to delete.

   The Backup Destination Details page displays information about the selected backup destination.
5. Click **Delete**.
6. In the resulting dialog, enter the backup destination name and click **Delete Backup Destination** to confirm the action.

Using the API to Manage Oracle Exadata Database Service on Cloud@Customer Backup Destinations

Review the list of API calls to manage your Oracle Exadata Database Service on Cloud@Customer backup destinations.

Copyright © 2008, 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".

Use these API operations to manage Oracle Exadata Database Service on Cloud@Customer backup destinations:

- `CreateBackupDestination`
- `DeleteBackupDestination`
- `GetBackupDestination`
- `ListBackupDestination`
- `UpdateBackupDestination`
- `ChangeBackupDestinationCompartment`

For the complete list of APIs, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- CreateBackupDestination
- DeleteBackupDestination
- GetBackupDestination
- ListBackupDestination
- UpdateBackupDestination
- ChangeBackupDestinationCompartment
- Database Service API

Automatic Backup and Recovery

- About Managing Database Backup for Oracle Exadata Database Service on Cloud@Customer
  Learn how configure backup when creating the database on Oracle Exadata Database Service on Cloud@Customer.
- Using the Console to Configure and Manage Backup and Recovery
  Learn how to use the console to view a list of available backups, edit backup settings, and restore a database for Oracle Exadata Database Service on Cloud@Customer.
- Using the API to Manage Database Backup and Recovery
  Learn how to use the API to manage database backup and recovery with Oracle Exadata Database Service on Cloud@Customer.

About Managing Database Backup for Oracle Exadata Database Service on Cloud@Customer

Learn how configure backup when creating the database on Oracle Exadata Database Service on Cloud@Customer.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Oracle Exadata Database Service on Cloud@Customer provides automatic database backup facilities that use Oracle Recovery Manager (RMAN). When you create a database on Oracle Exadata Database Service on Cloud@Customer, you can specify a backup destination and enable automatic backups. For more information, refer to the information in this publication about managing backup destinations for Oracle Exadata Database Service on Cloud@Customer.

After database creation, you can also:

- View a list of available backups.
- Enable or disable automatic backups.
- Edit backup settings.
- Restore a database.

You can perform these operations by using either the Console, or the API.

Automatic database backups are configured as follows:

- Automatic backups are scheduled daily. The automatic backup process can run at any time within the daily backup window, which is between midnight and 6:00 AM in the time zone of the virtual machine (VM) cluster that hosts the database.
- Automatic backups use a combination of full (RMAN level 0) and incremental (RMAN level 1) database backups:
  - For backups to a Zero Data Loss Recovery Appliance, after an initial full backup is performed, Zero Data Loss Recovery Appliance creates and validates virtual full backups from each daily incremental backup.
  - For backups to NFS, or OSS, the default interval between level 0 backups is seven days. The default level 0 day is Sunday.
- The retention period defines the period for which automatic backups are maintained:
  - For backups to Zero Data Loss Recovery Appliance, the retention policy that is implemented in the appliance controls the retention period.
  - For backups to Oracle Cloud Infrastructure Object Storage, or to an NFS backup destination, you can choose one of the following preset retention periods: 7 days, 14 days, 30 days, 45 days, or 60 days. The default retention period is 30 days.
- By default, Oracle Database runs in `ARCHIVELOG` mode, and archived redo log files are backed up every 30 minutes. .
- Regardless of the backup destination, backups of user data are encrypted by default.

While a backup is in progress, Oracle recommends that you avoid performing actions that could interfere with availability, such as restarting virtual machines, or applying patches. If an automatic backup operation fails, then the backup is deferred until the next day's backup window.

When required, you can restore Oracle Database to:

- The latest available restore point.
- A specific point in time by providing a time stamp.
- An Oracle Database System Change Number (SCN).

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> The backup and recovery facilities described in this topic cater only for database backup and recovery, which includes Oracle Database data files, log files, control files, and the server parameter (SP) file. You are responsible for backing up other files on your virtual machines. In particular, Oracle strongly recommends that you back up the Transparent Data Encryption (TDE) keystore (wallet). Without the TDE keystore, the Oracle Database backups are effectively useless, because you cannot read the data contained in the backup.

> **Note**
>
> If TAG based recovery fails with error `ORA-01152`, then use Recovery Manager (RMAN) directly to complete the recovery.
>
> If the server parameter file (SPFILE) recovery fails for local configuration using `dbaascli`, then use Recovery Manager (RMAN) directly to complete the recovery.

Using the Console to Configure and Manage Backup and Recovery

Learn how to use the console to view a list of available backups, edit backup settings, and restore a database for Oracle Exadata Database Service on Cloud@Customer.

- Viewing a List of Available Backups with the Console
  To view a list of available backups with Oracle Exadata Database Service on Cloud@Customer, complete this procedure.
- Editing Backup Settings with the Console
  To edit backup destinations, change backup schedules and other backup administration, you can use with the Oracle Exadata Database Service on Cloud@Customer console.
- Restoring a Database with the Console
  To restore a database to a point in time, to a system change number (SCN), or to the latest backup, use the Oracle Exadata Database Service on Cloud@Customer Console.

Viewing a List of Available Backups with the Console

To view a list of available backups with Oracle Exadata Database Service on Cloud@Customer, complete this procedure.

> **Note**
>
> Only the managed backups are synced to the Console. If you configure backups directly in the backend, then they are not synced to the Console. This is an expected behavior and Oracle has no plans to change this behavior.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that hosts the database in which you are interested.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that hosts the database in which you are interested.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. In the **Resources** list of the VM Cluster Details page, click **Databases**.
6. Click the name of the database in which you are interested.

   The Database Details page displays information about the selected database, which includes a list of the available backups.

Editing Backup Settings with the Console

To edit backup destinations, change backup schedules and other backup administration, you can use with the Oracle Exadata Database Service on Cloud@Customer console.

Use this procedure to change the available backup settings:

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that hosts the database for which you want to edit backup settings.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that hosts the database for which you want to edit backup settings.
5. In the **Resources** list of the VM Cluster Details page, click **Databases**.
6. Click the name of the database for which you want to edit backup settings.
   The Database Details page displays information about the selected database.
7. Click **Actions**, and then select **Configure automatic backups**.
8. Your current backup configuration determines the changes that you can make in the **Configure automatic backups** panel, as follows:
   - If automatic backups are not configured, then you can use the following settings to define the backup configuration for the database:
     - **Backup Destination Type:** From the list, choose an option.
       * **Local** Select to store backups locally in the Exadata Storage Servers on your Exadata Database Service on Cloud@Customer system.
         This option is available only if you enabled backups on local Exadata storage in the VM cluster that you want to host the database.
       * **Object Storage** Select to store backups in an object storage container managed by Oracle on Oracle Cloud Infrastructure.
         To use this option, your Exadata Database Service on Cloud@Customer system must have egress connectivity to Oracle Cloud Infrastructure Object Storage.
       * **NFS** Select to store backups in one of your previously defined backup destinations that uses Network File System (NFS) storage. See "Managing Backup Destinations for Exadata Database Service on Cloud@Customer".
         If you select this option, then you must also choose from the list of NFS **Backup Destinations**.
       * **Recovery Appliance** Select to store backups in one of your previously defined backup destinations that uses Oracle Zero Data Loss Recovery Appliance. See *Managing Backup Destinations for Exadata Database Service on Cloud@Customer*.
         If you select this option, then you must also provide the following information:
         * Choose **Backup Destinations** from the list of **Recovery Appliance** .

Copyright © 2008, 2026, Oracle and/or its affiliates.

* Choose from the **VPC User** list, which contains the list of virtual private catalog (VPC) user names that are defined in the **Recovery Appliance** backup destination.
* Provide the **Password** for the VPC user.

> **Note**
>
> If you select a backup destination, then you cannot change it later.

For more information on customizing Real Time Redo Transport (RTRT) behavior, see *Customizing Real Time Redo Transport (RTRT) Behavior for Recovery Appliance Backups*

– **Enable automatic backups:** Select this option to enable daily backups using the policy for automatic backups.

– **Backup retention period:** Select this option to choose one of the options for the length of time that automatic backups are retained.
For backups to local Exadata storage, you can choose a retention period of 7 days, or 14 days. The default retention period is 7 days.

For backups to Oracle Cloud Infrastructure Object Storage or to an NFS backup destination, you can choose one of the following preset retention periods: 7 days, 14 days, 30 days, 45 days, or 60 days. The default retention period is 30 days.

This option does not apply to Recovery Appliance backup destinations. For backups to Oracle Zero Data Loss Recovery Appliance, the retention policy that is implemented in the appliance controls the retention period.

- If automatic backups were previously configured, then you can make the following changes:
  – For Oracle Zero Data Loss Recovery Appliance backup destinations, you can update the **Password** for the virtual private catalog (VPC) user that is used to access the appliance.
  – For backup destinations that do not use Oracle Zero Data Loss Recovery Appliance, you can update the **Backup retention period** for automatic backups:
    * For backups to local Exadata storage, you can choose a retention period of 7 days or 14 days. The default retention period is 7 days.
    * For backups to Oracle Cloud Infrastructure Object Storage, or to an NFS backup destination, you can choose one of the following preset retention periods: 7 days, 14 days, 30 days, 45 days, or 60 days. The default retention period is 30 days.
    * For backups to Oracle Zero Data Loss Recovery Appliance, the retention policy that is implemented in the appliance controls the retention period.
  – You can set the option to **Enable automatic backups**. Select this option to enable automatic database backups. Deselect this option to suspend automatic database backups.

9. Click **Save**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Backup Destinations
  A backup destination is a target for sending database backups. The target may be a local filesystem, Object store, NFS mount or ZDLRA.
- Customizing Real Time Redo Transport (RTRT) Behavior for Recovery Appliance Backups

Restoring a Database with the Console

To restore a database to a point in time, to a system change number (SCN), or to the latest backup, use the Oracle Exadata Database Service on Cloud@Customer Console.

> **Note**
>
> You can only restore a 26ai database in the same DB home (in-place restore).

Use the following procedure to restore a database:

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that hosts the database that you want to restore.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that hosts the database that you want to restore.
5. In the **Resources** list of the VM Cluster Details page, click **Databases**.
6. Click the name of the database that you want to restore.
   The Database Details page displays information about the selected database.
7. Click **Restore Database**.
8. In the resulting dialog box, select one of the following options, and click **Restore Database**:
   - **Restore to latest**: The database is restored and recovered with zero, or least possible, data loss.
   - **Restore to a timestamp**: The database is restored and recovered to the specified timestamp.
   - **Restore to SCN**: The database is restored and recovered to the specified Oracle Database System Change Number (SCN). The specified SCN must be valid otherwise the operation fails.

> **Note**
>
> Backup fails after a point in time restore to a timestamp or SCN on NFS storage. Wait for 10 minutes or so before proceeding with the backup.

Using the API to Manage Database Backup and Recovery

Learn how to use the API to manage database backup and recovery with Oracle Exadata Database Service on Cloud@Customer.

Copyright © 2008, 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".

Use these API operations to manage database backup and recovery:

- `GetBackup`
- `ListBackups`
- `RestoreDatabase`
- `UpdateDatabase` - To enable and disable automatic backups.

For the complete list of APIs, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- GetBackup
- ListBackups
- RestoreDatabase
- UpdateDatabase
- Database Service API

User Configured Backup

It is the user's responsibility to configure and execute backup operations using `dbaascli` according to their preferences.

Backups using `dbaascli` are not recognized as Oracle managed backups. Use the `dbaascli database backup`, `dbaascli pdb backup`, `dbaascli database recover`, and `dbaascli pdb recover` commands to configure and perform various backup and recover operations.

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.

- How to Backup Using dbaascli for User Configured Backup
- How to Recover a Database

How to Backup Using dbaascli for User Configured Backup

User configured backups involve the following 4 steps:

1. Set up network and security rules to allow database hosts to store backup to the desired storage destination. For Exadata Cloud Service, refer to *Network Setup for Exadata Cloud Service Instances*.
2. Set up Backup Destination
3. Configure database with a Backup Destination

Copyright © 2008, 2026, Oracle and/or its affiliates.

4. Backup database

- Setting up Backup Destination
  Depending on the Oracle Database Cloud Service, Oracle Exadata Cloud@Customer, or Oracle Exadata Cloud Service, you can choose one of the following backup destinations for storing backups.
- Configuring Database for Backup
  Once the backup destination is set up and available, you can use the dbaascli utility to configure the database with backup destination and associated configuration parameters, for example, Backup Retention Recovery Window, Backup Scheduling, Archivelog Scheduling, and so on.
- Backing Up a Container Database (CDB)
- Backing Up a Pluggable Database (PDB)
- Fetching Backup Job Status

Related Topics

- Network Setup for Exadata Cloud Service Instances

Setting up Backup Destination

Depending on the Oracle Database Cloud Service, Oracle Exadata Cloud@Customer, or Oracle Exadata Cloud Service, you can choose one of the following backup destinations for storing backups.

- **Oracle Cloud Infrastructure Object Storage:** Applies to Exadata Cloud@Customer and Exadata Cloud Service.
- **Zero Data Loss Recovery Appliance (ZDLRA):** Applies to Exadata Cloud@Customer.
- **Network File System (NFS):** Applies to Exadata Cloud@Customer.
- **Fast Recovery Area (FRA):** Applies to Exadata Cloud@Customer.

Related Topics

- Backup Destinations
  A backup destination is a target for sending database backups. The target may be a local filesystem, Object store, NFS mount or ZDLRA.

Configuring Database for Backup

Once the backup destination is set up and available, you can use the dbaascli utility to configure the database with backup destination and associated configuration parameters, for example, Backup Retention Recovery Window, Backup Scheduling, Archivelog Scheduling, and so on.

The configuration parameters are specified through an input configuration file having a list of **parameter=value** pairs. The template for this input file can be generated using `--getConfig` option, which can be modified according to desired destination settings and other preferences and then used as input to `--configure` option.

```
dbaascli database backup --getConfig --dbName <value> --configFile <value>
```

Returns the Backup configuration of the database in a file specified by the user. If the backup is never configured, it returns the default template where the user can fill the configuration parameters with their values and use it as input to --configure command option.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using `dbaascli database backup --configure`, you can set the backup configuration information for a database in the following cases:

1. **User Managed Backup:** Set complete backup configuration based on the destination.
2. **Oracle Managed Backup:** Set additional configuration parameters if a parameter is compatible with Oracle Managed Backup.
   Note that using `dbaascli database backup --configure` to set incompatible parameters with Oracle Managed Backups can lead to backup/recovery issues. Currently, Oracle does not validate whether a parameter is compatible or not. Until Oracle implements validation, it is your responsibility to verify compatibility.

   ```
   dbaascli database backup --configure --dbName <value> --configFile <value>
   ```

   This asynchronous command generates a universally unique identifier (UUID), which you can use to track the status.

   ```
   dbaascli database backup status --uuid <value> --dbname <value>
   ```

   To retrieve the current configuration for the validation, run:

   ```
   dbaascli database --dbaname --getConfig
   ```

> **Note**
>
> When using Oracle Managed Backups, you may want to change some of the backup configuration settings based on your application requirements. You can use `dbaascli` to modify these parameters if they are tagged with `Compatible with console automatic backup`.

Configuration Parameters for Backup

Table 5-18 General Configuration Parameters (valid for all backup destinations except Local Storage (FRA))

Parameter	Description	Compatible with Oracle Managed Backup
`bkup_rman_compression`	Level of compression applied to automatic backups. Valid values are `NONE`, `basic`, `low`, `medium`, and `high`. Default: `low`. `NONE` disables RMAN compression.	Yes
`bkup_set_section_size`	Enables the use of the RMAN multisection backup feature. Valid values are `yes` and `no`.	Yes

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-18 (Cont.) General Configuration Parameters (valid for all backup destinations except Local Storage (FRA))

Parameter	Description	Compatible with Oracle Managed Backup
bkup_section_size	RMAN section size that is used for automatic backups. The default value is 64G. Applicable only when bkup_set_section_size is set to yes.	Yes
bkup_channels_node	The number of RMAN channels per node used for automatic backups. Valid values are between 1 and 32.	Yes
bkup_daily_time	Start time of the automatic daily backup expressed in 24-hour time format as hh:mm.	**Exadata Cloud@Customer:** Yes **Exadata Cloud Service:** No Not compatible with Oracle Managed Backup in Exadata Cloud Service. Scheduling responsibility resides with the Control Plane.
bkup_archlog_frequency	Interval in minutes between automatic backups of archived database log files. Valid values are 15, 20, 30, 60, 120 through 1440 in 1-hour intervals expressed in minutes. Default: 30 for Exadata Cloud@Customer.	Yes
bkup_l0_day	This parameter controls the Level 0 day of the week for both OSS and NFS. Day of the week when a level 0 backup is taken. Valid values are mon, tue, wed, thu, fri, sat, sun. Longer formats, for example, Monday, Tuesday are also supported. Applicable only when bkup_oss is set to yes. Default: sun.	**Exadata Cloud@Customer:** Yes **Exadata Cloud Service:** No Not compatible with Oracle Managed Backup in Exadata Cloud Service. Scheduling responsibility resides with the Control Plane.

Configuration Parameters for Object Storage Service (OSS) Destination

Table 5-19 Configuration Parameters for Object Storage Service (OSS) Destination

Parameter	Description	Compatible with Oracle Managed Backup
bkup_oss	Object storage service will be used as the backup destination. Valid values are yes and no.	No

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-19 (Cont.) Configuration Parameters for Object Storage Service (OSS) Destination

Parameter	Description	Compatible with Oracle Managed Backup
`bkup_oss_recovery_window`	The retention period for the backups is up to 90. Applicable only when `bkup_oss` is set to `yes`. Default: 30.	No
`bkup_oss_url`	Location of the storage container that is used for backup to cloud storage. Applicable only when `bkup_oss` is set to `yes`.	No
`bkup_oss_user`	User name of the Oracle Cloud user having the write privileges on the cloud storage container specified in `bkup_oss_url`. Applicable only when `bkup_oss` is set to `yes`.	No
`bkup_oss_passwd`	The password of the Oracle Cloud user having the write privileges on the cloud storage container specified in `bkup_oss_url`. Applicable only when `bkup_oss` is set to `yes`.	No

> **Note**
>
> Currently, Zero Data Loss Recovery Appliance (ZDLRA) destination is supported only on Exadata Cloud@Customer.

Configuration Parameters for Zero Data Loss Recovery Appliance (ZDLRA) Destination

Table 5-20 Configuration Parameters for Zero Data Loss Recovery Appliance (ZDLRA) Destination

Parameter	Description	Compatible with Oracle Managed Backup
`bkup_zdlra`	Enables backups to a Recovery Appliance. Valid values are `yes` and `no`.	No
`bkup_zdlra_url`	Location of the Recovery Appliance that is being used for backups. Applicable only when `bkup_zdlra` is set to `yes`.	No

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-20 (Cont.) Configuration Parameters for Zero Data Loss Recovery Appliance (ZDLRA) Destination

Parameter	Description	Compatible with Oracle Managed Backup
`bkup_zdlra_user`	The virtual private catalog (VPC) user name for the Recovery Appliance specified in `bkup_zdlra_url`. Applicable only when `bkup_zdlra` is set to `yes`.	No
`bkup_zdlra_passwd`	The password of the Recovery Appliance user specified in `bkup_zdlra_url`. Applicable only when `bkup_zdlra` is set to `yes`.	No

For more information, see *Customizing Real Time Redo Transport (RTRT) Behavior for Recovery Appliance Backups*.

Configuration Parameters for Network File System (NFS) Destination

Table 5-21 Configuration Parameters for Network File System (NFS) Destination

Parameter	Description	Compatible with Oracle Managed Backup
`bkup_nfs`	Enables backups to NFS mounted directory. Valid values are `yes` and `no`.	No
`bkup_nfs_loc`	The NFS mounted location, the directory provided must be a mount point and available on all nodes. Applicable only when `bkup_nfs` is set to `yes`.	No
`bkup_nfs_recovery_window`	The retention period for backups on NFS storage is expressed as a number of days up to 90. Applicable only when `bkup_nfs` is set to `yes`. Default: 30.	No

> **Note**
>
> Currently, Network File System (NFS) destination is supported only on Exadata Cloud@Customer.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Configuration Parameters for Local Storage (FRA) Destination

Table 5-22 Configuration Parameters for Local Storage (FRA) Destination

Parameter	Description	Compatible with Oracle Managed Backup
`bkup_disk`	Enables backups to local Exadata storage. Valid values are `yes` and `no`.	No
`bkup_disk_recovery_window`	The retention period for backups on local Exadata storage is expressed as a number of days up to 14. Applicable only when `bkup_disk` is set to `yes`. Default: 7.	No

> **Note**
>
> Currently, Local Storage (FRA) destination is supported only on Exadata Cloud@Customer.

Related Topics

- Customizing Real Time Redo Transport (RTRT) Behavior for Recovery Appliance Backups
- 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 recover
  To recover a database, use the `dbaascli database recover` command.
- 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 recover
  To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.

Backing Up a Container Database (CDB)

Before performing a backup, you must set up the backup configuration using `dbaascli database backup --configure` command as a prerequisite. For more information, see *dbaascli database backup*.

Related Topics

- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

Backing Up a Pluggable Database (PDB)

Copyright © 2008, 2026, Oracle and/or its affiliates.

You can also perform a backup of a specific PDB. This is useful if one or more PDBs need to be backed up immediately or adhoc need basis, instead of backing up the complete CDB which has its own schedule. For more information, see *dbaascli pdb backup*.

Related Topics

- dbaascli pdb backup
  To backup a pluggable database (PDB), query PDB backups, and delete a PDB backup, use the `dbaascli pdb backup` command.

Fetching Backup Job Status

As the backup operation is being run, its status can be monitored by providing the job ID that is displayed to the user when the backup command is executed. For more information, see *dbaascli database backup* and *dbaascli pdb backup*.

Related Topics

- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.
- dbaascli pdb backup
  To backup a pluggable database (PDB), query PDB backups, and delete a PDB backup, use the `dbaascli pdb backup` command.

How to Recover a Database

To perform recovery of the database, there are various options depending on the type of fault, availability of data, and its' backup and recovery requirement. For more information, see *dbaascli database recover*.

- Recovering a PDB
- Fetching Recover Job Status

Related Topics

- dbaascli database recover
  To recover a database, use the `dbaascli database recover` command.

Recovering a PDB

You can also recover a particular PDB instead of complete CDB whenever that can be avoided. It reduces the impact of the outage as only the particular PDB is being recovered while the rest of the CDB and PDBs are still available for business and applications. For more information, see *dbaascli pdb recover*.

Related Topics

- dbaascli pdb recover
  To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.

Fetching Recover Job Status

Similar to backup, the recovery commands' status can be monitored by providing the job ID that is displayed to the user when the recovery command is executed. To know the status of

Copyright © 2008, 2026, Oracle and/or its affiliates.

recovery operation: For more information, see *dbaascli database recover* and *dbaascli pdb recover*.

Related Topics

- dbaascli database recover
  To recover a database, use the `dbaascli database recover` command.
- dbaascli pdb recover
  To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.

Configuring and Customizing Backups with dbaascli

In addition to the console-based automated backup option, use dbaascli to configure and customize backups.

> **Note**
>
> `bkup_api` will be deprecated in a future release. 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*.

> **Note**
>
> In dbaascli Release 25.1.2.0.0, the backup configuration parameters have been renamed. However, you can still use the old parameter names, as they are retained for backward compatibility.

- Customizing Backup Settings by Using a Generated Configuration File
  You can customize backup settings for a database deployment by generating a file containing the current customizable settings, editing the file, and then using the file to update the backup settings.

Related Topics

- User Configured Backup
  It is the user's responsibility to configure and execute backup operations using `dbaascli` according to their preferences.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Customizing Backup Settings by Using a Generated Configuration File

You can customize backup settings for a database deployment by generating a file containing the current customizable settings, editing the file, and then using the file to update the backup settings.

> **Warning**
>
> In general, `bkup_api` configured backups, as described in this section, **should not** be used in conjunction with console-enabled automatic backups, other than the exceptions noted specifically below. If using parameters, other than those noted as safe below, then do not enable console-based backups; otherwise, conflicting conditions or over-writes of settings can occur, and backups may not execute successfully.

> **Note**
>
> In dbaascli Release 25.1.2.0.0, the backup configuration parameters have been renamed. However, you can still use the old parameter names, as they are retained for backward compatibility.

To generate a configuration file with the current backup settings and use it to update the settings:

1. Connect to a virtual machine as `opc` user.

   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.

2. Start a `root` user command shell:

   ```
   sudo -s
   ```

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 --dbname` *`dbname`*
   Where:

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

   For example:

   - To generate config file:

     ```
     # dbaascli database backup --getconfig --dbname <dbname>
     DBAAS CLI version MAIN
     Executing command database backup --getconfig --dbname <dbname>
     vi Session log: /var/opt/oracle/log/v23db8/database/backup/
     dbaastools_2025-01-28_12-36-04-PM_236081.log
     logfile:/var/opt/oracle/log/dtrs/dcs-dtrs.0.*.log
     File /var/opt/oracle/log/dtrs/staging/tmpl/
     backup-6b4fff1e-370c-429b-9a28-723e9c3ba953-2025128-03611.cfg created
     dbaascli execution completed
     ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

- To customize backup settings using the config file:

```
# vi /var/opt/oracle/log/dtrs/staging/tmpl/
backup-6b4fff1e-370c-429b-9a28-723e9c3ba953-2025128-03611.cfg
# dbaascli database backup --configure --configFile /var/opt/oracle/log/
dtrs/staging/tmpl/
backup-6b4fff1e-370c-429b-9a28-723e9c3ba953-2025128-03611.cfg --dbname
v23db8
```

4. Edit the parameter values in the generated file to change any settings you want to customize in the backup configuration.

   The following parameters can be modified to customize the backup configuration:

> **Note**
>
> `Compatible with Console Automatic Backups=Yes` indicates the parameter is safe to change, even when using console-based automatic backups. If using parameters with `Compatible with Console Automatic Backups=No`, then do not enable backups through the console.

Table 5-23 Backup Configuration Parameters - Cron Parameters Specific to bkup_api

Parameter	Description	Compatible with Console Automatic Backups*
Old name: `bkup_cron_entry` New name: `scheduleBackups`	Enables 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 © 2008, 2026, Oracle and/or its affiliates.

Table 5-24 Backup Configuration Parameters - General RMAN Configuration Parameters (valid for all backup destinations except Local Storage (FRA))

Parameter	Description	Compatible with Console Automatic Backups*
Old name: `bkup_rman_compression` New name: `compressionLevel`	Level of compression applied to automatic backups. Valid values are `NONE`, `basic`, `low`, `medium`, and `high`. Default value is `low`. A value of `NONE` disables RMAN compression. If RMAN compression is enabled, then any TDE encrypted datafile will be decrypted, compressed, and RMAN encrypted.	Yes
Old name: `bkup_section_size` New name: `sectionSize`	RMAN section size that is used for automatic backups. Default value is 64G.	Yes
Old name: `bkup_channels_node` New name: `channelsPerNode`	Number of RMAN channels per node used for automatic backups. Valid values are between 1 and 32. Default value is 4.	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 1 hour intervals expressed in minutes. Default value is 30 for Exadata Cloud@Customer.	Yes
Old name: `bkup_type` New name: `backupDestination`	The type of the location where the backup resides. Any one of the following values. OSS or ZDLRA or DISK or NFS	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 © 2008, 2026, Oracle and/or its affiliates.

Table 5-24 (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 NFS, while it is disabled for ZDLRA and DISK. This setting can be modified only for NFS.	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. By default, the value is set to false for ZDLRA, and this setting cannot be changed.	Yes
Old name: `bkup_archlog_fra_retention` New name: `archivelogRetentionDays`	Specifies the number of days archive log to be retained in FRA.	Yes

Table 5-25 Backup Configuration Parameters - Local Storage (FRA) Parameters

Parameter	Description	Compatible with Console Automatic Backups*
`backupDestination=disk`	Enables backups to local Exadata storage. Valid values are `yes` and `no`.	No
Old name: `bkup_disk_recovery_window` New name: `diskRecoveryWindow`	Retention period for backups on local Exadata storage, expressed as a number of days up to 14. Applicable only when `bkup_disk` is set to `yes` or `backupdestination` is set to `DISK` Default value is 7.	No

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-26 Backup Configuration Parameters - Network File System (NFS) Parameters

Parameter	Description	Compatible with Console Automatic Backups*
`backupDestination=nfs`	Enables backups to NFS mounted directory. Valid values are `yes` and `no`.	No
Old name: `bkup_nfs_loc` New name: `nasLocation`	The NFS mounted location, the directory provided must be a mountpoint and available on all nodes.	No
Old name: `bkup_nfs_recovery_window` New name: `nasRecoveryWindow`	Retention period for backups on NFS storage, expressed as a number of days up to 90. Applicable only when `bkup_nfs` is set to `yes` or `backupdestination` is set to `NFS`. Default value is 30.	No

Table 5-27 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
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

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-27 (Cont.) Backup Configuration Parameters - Object Storage Service (OSS) Parameters

Parameter	Description	Compatible with Console Automatic Backups*
Old name: `bkup_oss_passwd` New name: `ossAuthToken`	Password of the Oracle Cloud user having write privileges on the cloud storage container specified in `ossURL`. Applicable only when `bkup_oss` is set to `yes` or `backupdestination` is set to `OSS`.	No

> **Note**
>
> Zero Data Loss Recovery Appliance (ZDLRA) parameters are only valid for Exadata Cloud@Customer installations.

Table 5-28 Backup Configuration Parameters - Zero Data Loss Recovery Appliance (ZDLRA) Parameters

Parameter	Description	Compatible with Console Automatic Backups*
`backupDestination=zdlra`	Enables backups to a Recovery Appliance. Valid values are `yes` and `no`.	No
Old name: `bkup_zdlra_url` New name: `zdlraUrl`	Location of the Recovery Appliance that is being used for backups. Applicable only when `bkup_zdlra` is set to `yes` or `backupdestination` is set to `ZDLRA`.	No
Old name: `bkup_zdlra_user` New name: `zdlraVpcUserName`	The virtual private catalog (VPC) user name for the Recovery Appliance specified in `bkup_zdlra_url`. Applicable only when `bkup_zdlra` is set to `yes` or `backupdestination` is set to `ZDLRA`.	No
Old name: `bkup_zdlra_passwd` New name: `zdlraVpcPassword`	Password of the Recovery Appliance user specified in `bkup_zdlra_url`. Applicable only when `bkup_zdlra` is set to `yes` or `backupdestination` is set to `ZDLRA`.	No

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 5-29 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 (only for NFS and OSS backups)
Old name: `bkup_rcat_user` New name: `catalogUserName`	Recovery catalog user name. Applicable only when `bkup_use_rcat` is set to `yes`.	Yes (only for NFS and OSS backups)
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 (only for NFS and OSS backups)
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 (only for NFS and OSS backups)

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.

5. Use the `dbaascli database backup --configure` command to update the backup settings using the file containing your updated backup settings:

```
# dbaascli database backup --configure --configFile <file_name> --dbname
    <database_name>
```

6. Use the `dbaascli database backup --status` to check the status of the configuration update:

```
# dbaascli database backup --status --uuid <uuid> --dbname
    <database_name>
```

7. Exit the `root` user command shell:

```
exit
```

> **Note**
>
> Any changes you make by using the `bkup_api` command are not reflected in the Oracle Database Exadata Cloud@Customer console.

Related Topics

- Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management
- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Creating an On-Demand Backup by Using dbaascli

You can use the `dbaascli` to create an on-demand backup of a complete database or an individual pluggable database (PDB):

> **Note**
>
> `bkup_api` will be deprecated in a future release. 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 Backups and Recovery*.

> **Note**
>
> Using this method for a manual backup is safe to use in conjunction with automatic backups managed through the Console. Manual backups done this way will appear in the console after some time due to synchronization.

To change which database configuration files are backed up:

1. Connect as the `oracle` user to a compute node.

   For detailed instructions, see *Connecting to a Compute Node with SSH*.
2. Start a root-user command shell:

   ```
   # sudo -s
   #
   ```
3. Enter the `dbaascli` command:
   - To create a backup that follows the current retention policy:

     ```
     # dbaascli database backup --start --dbname <database_name>
     ```

     Where `dbname` is the database name for the database that you want to back up.
   - To create an on-demand backup of a specific PDB, use the following `bkup_api` command:

     ```
     # /var/opt/oracle/bkup_api/bkup_api bkup_start --dbname=dbname --
     pdb=pdbname
     ```
   - To create a long-term backup:

     ```
     # dbaascli database backup --start --archival --dbname --tag
     <archival_tag>
     ```
   - To create an on-demand RMAN level 0 backup:

     ```
     # dbaascli database backup --dbname <value> --start --level0
     ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

You can use this option to manually perform an RMAN level 0 (full) backup if the scheduled weekly level 0 backup fails or following a major structural change in the database, such as adding a new data file or tablespace. This option is only valid for backup configurations that use cloud storage only.

4. After you start an on-demand backup, the backup process runs in the background. To check the progress of the backup process, run the following `dbaascli` command on the same compute node where the backup is running:

   ```
   # dbaascli database backup --dbname <value> --status --uuid <backup_uuid>
   ```

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

   ```
   # exit
   # exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management

Backups, configured in the Exadata Cloud Infrastructure console, API or `dbaascli` work for a variety of backup and recovery use cases. If you require use cases not supported by the cloud-managed backups, then you can manage database backup and recovery manually, using the Oracle Recovery Manager (RMAN) utility. For information about using RMAN, see the *Oracle Database Backup and Recovery User's Guide for Release 19*.

Managing backup and recovery, using RMAN, on Exadata Cloud Infrastructure requires taking full ownership of both database and archive log backups, and the cloud-managed backups should no longer be used. Before manual backups are started, the cloud-managed backup functionality should be disabled. This is needed so the cloud backup jobs do not purge archive logs before they are manually backed up and do not conflict with the manual backups.

You can use the `dbaascli` utility to disable cloud-managed backups, including disabling the automatic archive log purge job.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- Oracle Database Backup and Recovery User's Guide for Release 19

Customizing Real Time Redo Transport (RTRT) Behavior for Recovery Appliance Backups

Real Time Redo Transport can be enabled or disabled using `dbaascli`. This is available from `dbaastools` 21.4.1 release.

Copyright © 2008, 2026, Oracle and/or its affiliates.

For example, to enable RTRT for the database *myTestDB*:

```
dbaascli database backup --dbName myTestDB --configure –enableRTRT
```

For more information, see *dbaascli database backup*.

> **Note**
>
> When Real-Time Redo Transport (RTRT) is enabled with Zero Data Loss Recovery Appliance (ZDLRA) as the backup destination, the tooling automatically provisions a VPC user on the protected database. This user has the same identifier as the VPC user on the recovery appliance and is granted `SYSOPER` privileges to enable redo transport and backup operations.

> **Note**
>
> - Enabling or disabling RTRT works only for Oracle Database version 12.2 and higher.
> - If you have enabled Data Guard association between databases, then after enabling RTRT on the Primary database, you must manually copy the password file from the Primary to the Standby database. To learn more about copying the password file from the primary to the standby database, refer to *12c: Data Guard Physical Standby - Managing password files in a RAC Physical Standby (Doc ID 1984091.1)*.

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.
- 12c: Data Guard Physical Standby - Managing password files in a RAC Physical Standby (Doc ID 1984091.1)
- About Real-Time Redo Transport

Alternative Backup Methods

Learn about alternative backup methods that are available in addition to the OCI Console.

Backup for databases on Exadata Database Service on Cloud@Customer 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).

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **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*.

RMAN is the backup tool included with the Oracle Database. For information about using RMAN, see the Oracle Database Backup and Recovery User's Guide for Release 19. Using RMAN to back up databases on Exadata Cloud Infrastructure provides the most flexibility in terms of backup options, but also the most complexity.

> **Note**
>
> While using RMAN for restoring databases backed up through any method described herein is considered safe, RMAN should NEVER be used to set up backups in conjunction with either console (and OCI API / CLI), nor in conjunction with `dbaascli`. If you choose to orchestrate backups manually leveraging RMAN, you should not use either console automated backups, nor should you use `dbaascli`. You must first completely disable console based automated backups. For more information, see *Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management*.

The `dbaascli` method offers a middle ground between RMAN and console automated backups in terms of flexibility and simplicity. Use `dbaascli` if needed functionality is not supported with console automated backups, but when you wish to avoid complexity of using RMAN directly. In certain cases, `dbaascli` can be used to modify the console automated backup configuration, but this is not generally the case. Generally, `dbaascli` must be used instead of enabling backups in the console.

- Recovering a Database Using Oracle Recovery Manager (RMAN)

Related Topics

- User Configured Backup
  It is the user's responsibility to configure and execute backup operations using `dbaascli` according to their preferences.
- Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management

Recovering a Database Using Oracle Recovery Manager (RMAN)

If you backed up your database using dbaascli, then you can manually restore that database backup by using the Oracle Recovery Manager (RMAN) utility. For information about using RMAN, see the *Oracle Database Backup and Recovery User's Guide for Release 19*.

> **Note**
>
> While recovering using RMAN is safe, you must not use RMAN to initiate backups or edit backup setting in conjunction with either `dbaascli` usage or in conjunction with automated console backups. Doing so could result in conflicting conditions or over-writes of settings, and backups may not execute successfully.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Oracle Database Backup and Recovery User's Guide for Release 19

Standalone Backup

ExaDB-C@C now supports standalone (ad-hoc) backups for databases in the following configurations:

- Standalone databases (no Data Guard)
- Primary databases in a Data Guard configuration
- Standby databases in a Data Guard configuration

Backup Destinations

The backup destination for standalone backups is determined by the automatic backup configuration:

- Automatic backups disabled (NONE): destination is Object Storage
- Automatic backups enabled with Object Storage: destination is Object Storage
- Automatic backups enabled with Network File System (NFS): destination is the same NFS as configured for automatic backups
- Automatic backups enabled with Oracle Zero Data Loss Recovery Appliance (ZDLRA): destination is the same Recovery Appliance as configured for automatic backups

Destination-Specific Behavior

- Oracle Zero Data Loss Recovery Appliance (ZDLRA):
  - Standalone backups are retained for up to the recovery window configured for automatic backups
  - Automatic backups must already be configured to ZDLRA
  - Customers are responsible for managing ZDLRA wallets
- Network File System (NFS):
  - No retention window; customers must delete backups manually
- Object Storage:
  - No retention window; customers must delete backups manually
- Using the Console to Create a Standalone Backup
  To take a standalone backup, complete this procedure.

Using the Console to Create a Standalone Backup

To take a standalone backup, complete this procedure.

1. Go to Database Details page.
2. Click **Backups**.
3. Click **Create backup**.
4. On the resulting Create backup page, enter a descriptive **Display name**.
5. Click **Create**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Patch and Update an Oracle Exadata Database Service on Cloud@Customer System

Learn to update and patch the Oracle Exadata Database Service on Cloud@Customer System

- Perform User Managed Maintenance Updates
- Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to perform patching operations on Exadata database virtual machines and Database Homes by using the Console, API, or the CLI.
- Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System Manually
  This topic describes the procedures for patching and updating various components in Oracle Exadata Database Service on Cloud@Customer outside of the cloud automation. For information related to patching and updating with `dbaascli`, refer to "Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli."
- Resolving Dependency Issues Associated with Additional Non-Exadata Software Packages for Guest VM Upgrade
  If you've installed non-Exadata software packages beyond those provided by Oracle, and the precheck fails during a Guest VM upgrade due to conflicts between and Oracle-installed RPMs, you can use the following procedure to resolve the conflicts and proceed with the upgrade.

Perform User Managed Maintenance Updates

Maintaining a secure Oracle Exadata Database Service on Cloud@Customer system in the best working order requires you to perform the following tasks regularly:

- Patching the Oracle Grid Infrastructure and Oracle Database software on the VM Cluster virtual machines. For information and instructions, see *Patching and Updating an Exadata Cloud@Customer System*.
- Updating the operating system on the VM Cluster virtual machines. For information and instructions, see *Updating Guest VM Operating System* and *Oracle Clusterware Configuration and Administration*.

Related Topics

- Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to perform patching operations on Exadata database virtual machines and Database Homes by using the Console, API, or the CLI.
- Updating Guest VM Operating System
  Learn to update the operating system image on Oracle Exadata Database Service on Cloud@Customer VM cluster nodes in an automated manner from the OCI console and APIs.
- Administering Oracle Clusterware

Copyright © 2008, 2026, Oracle and/or its affiliates.

Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System

Learn how to perform patching operations on Exadata database virtual machines and Database Homes by using the Console, API, or the CLI.

For information and instructions on patching the system by using the `dbaascli` utility, see *Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System Manually*.

For more information and examples for applying database quarterly patches on Exadata Cloud@Customer refer to My Oracle Support note: *How to Apply Database Quarterly Patch on Exadata Cloud Service and Exadata Cloud at Customer Gen 2 (Doc ID 2701789.1)*.

For more guidance on achieving continuous service during patching operations, see the *Application Checklist for Continuous Service for MAA Solutions* white paper.

- Patching and Updating VM Clusters and Database Homes
  Learn how to perform patching operations on VM Cluster Grid Infrastructure (GI) and Database Homes using the Console or API
- Updating Guest VM Operating System
  Learn to update the operating system image on Oracle Exadata Database Service on Cloud@Customer VM cluster nodes in an automated manner from the OCI console and APIs.
- Upgrading Oracle Grid Infrastructure on an Oracle Exadata Database Service on Cloud@Customer VM Cluster
  Learn to upgrade Oracle Grid Infrastructure on an Oracle Exadata Database Service on Cloud@Customer VM cluster using the Oracle Cloud Infrastructure Console or API.
- Upgrading Oracle Databases
  Learn to upgrade an Exadata database instance to Oracle Database 19c and Oracle AI Database 26ai by using the Console and the API.

Related Topics

- Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System Manually
  This topic describes the procedures for patching and updating various components in Oracle Exadata Database Service on Cloud@Customer outside of the cloud automation. For information related to patching and updating with `dbaascli`, refer to "Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli."
- https://support.oracle.com/epmos/faces/DocContentDisplay?id=2701789.1
- Application Checklist for Continuous Service for MAA Solutions

Patching and Updating VM Clusters and Database Homes

Learn how to perform patching operations on VM Cluster Grid Infrastructure (GI) and Database Homes using the Console or API

- About Patching and Updating VM Cluster's GI and Database Homes
- Prerequisites for Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System
  Check and apply the latest Cloud patches that are dowloaded and made available by Oracle on the CPS host.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the Console for Patching and Updating VM Cluster's GI and Database Homes
  Learn how to use the console to view the history of patch operations on VM cluster and Database Homes, apply patches, and monitor the status of patch operations.
- Using the API for Patching and Updating VM Cluster and Database Homes
  Use various API features to help manage patching an Oracle Exadata Database Service on Cloud@Customer system.

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

Patching a VM cluster updates components on each of the VM guests in the VM cluster. VM cluster patching updates the grid infrastructure (GI) and Database Home patching updates the Oracle Database software shared by the databases in that home.

For more information on available patches, see My Oracle Support note https://support.oracle.com/epmos/faces/DocContentDisplay?id=2333222.1.

Consider the following best practices:

- Because patching a system requires a reboot, plan to run the operations at a time when they will have minimal impact on users.
- Oracle recommends that you back up your databases before you apply any patches. For information about backing up the databases, see *Managing Database Backup and Recovery on Exadata Database Service on Cloud@Customer*.
- Oracle recommends your Oracle Grid Infrastructure RU version be no more than 6 months behind your latest database RU version. When updating the database version, you should also update the GI version if possible.
- 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. See *Using the Console to Move a Database to Another Home*.

Related Topics

- Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer
  Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Cloud@Customer.
- Using the Console to Move a Database to Another Home
  You can update the version of a VM cluster database by moving it to a Database Home that is running the version of Oracle Database you are interested in.

Prerequisites for Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System

Check and apply the latest Cloud patches that are dowloaded and made available by Oracle on the CPS host.

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 © 2008, 2026, Oracle and/or its affiliates.

Using the Console for Patching and Updating VM Cluster's GI and Database Homes

Learn how to use the console to view the history of patch operations on VM cluster and Database Homes, apply patches, and monitor the status of patch operations.

Oracle recommends that you use the precheck action to ensure your VM cluster or Database Home has met the requirements for the patch you want to apply.

- Using the Console to Perform Grid Infrastructure Updates
  Learn to apply Grid Infrastructure updates on a VM cluster.
- Using the Console to Perform Update Operation on a Database Home
  Learn to apply patches on a Database Home.
- Using the Console to View Update History
  Each update 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.
- Using the Console to Move a Database to Another Home
  You can update the version of a VM cluster database by moving it to a Database Home that is running the version of Oracle Database you are interested in.

Using the Console to Perform Grid Infrastructure Updates

Learn to apply Grid Infrastructure updates on a VM cluster.

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

   VM Clusters is selected by default.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the VM cluster on which you want to perform a patch operation.
4. On the resulting VM Cluster Details page, click **Updates**.
5. Review the list of available patches for the VM cluster.

   The Oracle Grid Infrastructure software images are generally available Oracle Grid Infrastructure software images that you can use to patch your cluster. Oracle images that can be used for patching have the update Type of "Patch".

   The Custom Grid Infrastructure software images are the Grid Infrastructure software image that you have created in advance.

   Use the **Select a Compartment** selector to specify the compartment that contains the database software image.

   Use the **Region** filter to access the software images created in a different region.
6. Click the Actions icon (three dots) for the patch you are interested in, and then click one of the following actions:
   - **Precheck:** Check for any prerequisites to make sure that the patch can be successfully applied. Oracle highly recommends that you run this operation before you apply a patch. Precheck does not cause any availability impact to the cluster, everything remains operational.
   - **Apply Grid Infrastructure update:** Applies the selected patch.

Copyright © 2008, 2026, Oracle and/or its affiliates.

7. Confirm when prompted.

The patch list displays the status of the operation. While the precheck is running, the patch's status shows `Checking`. While a patch is being applied, the patch's status shows `Applying` and the VM cluster's status shows `Updating`. During patching, lifecycle operations on the VM cluster and its resources are temporarily unavailable. If patching completes successfully, the patch's status changes to `Applied` and the VM cluster's status changes to `Available`. You can view more details about an individual patch operation by clicking **Update History**. Grid Infrastructure patching is done in a rolling fashion, node by node, and the cluster resources will be stopped and restarted on each node.

Using the Console to Perform Update Operation on a Database Home

Learn to apply patches on a Database Home.

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

   VM Clusters is selected by default.

2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.

3. In the list of VM clusters, click the VM cluster where the Database Home is located.
4. Click **Database Homes**.
5. In the list of Database Homes, click the Database Home on which you want to perform a patch operation.
6. Click **Updates**.
7. Review the list of available patches for the Database Home.

   The Oracle Database Software Images are generally available Oracle Database software images that you can use to update your database. Oracle images that can be used for patching have the update Type of "Patch".

   The Custom Database Software Images are the database software image that you have created in advance.

   Use the **Select a Compartment** selector to specify the compartment that contains the database software image.

   Use the **Type** filter to select **Oracle-published images** or **Custom images**.

8. Click the Actions icon (three dots) for the Release Update (RU) you are interested in, and then click **Apply**.

   On the Apply Database Home Update page, you can choose to run a precheck or apply the updates.

   - **Precheck:** Check for any prerequisites to make sure that the update can be successfully applied. Oracle highly recommends that you run this operation before you apply an update. The Precheck does not cause any availability impact to the cluster, everything remains operational.
   - **Apply:** Applies the selected Release Update.
   - Click **Advanced options**, choose one of the following **Datapatch options**.
     - **Run datapatch on the databases**: This is the default option, which applies the datapatch operation to all databases running from the Database Home being updated.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **Skip datapatch on the databases**: Select this option to skip the datapatch execution on all databases running from the Database Home being updated.
- **Skip datapatch for closed PDBs only**: Select this option to skip datapatch on only those PDBs that are in a closed state during the update.

The patch list displays the status of the operation. While the precheck is running, the patch's status shows `Checking`. While a patch is being applied, the patch's status shows `Applying`, the status of the Database Home and the databases in it display as `Updating`, and lifecycle operations on the VM cluster and its resources are temporarily unavailable. Patches are applied to the Database Home in a rolling fashion, node by node, and each database in the home is stopped and then restarted. This may result in temporary service disruption. If patching completes successfully, the patch's status changes to `Applied` and the Database Home's status changes to `Available`. You can view more details about an individual patch operation by clicking **Update History**.

Using the Console to View Update History

Each update 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.

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

- Using the Console to View the Update History of a VM Cluster
  Learn how to view the history of patches applied on a VM cluster.
- Using the Console to View the Update History of a Database Home
  Learn how to view the history of patches applied on a Database Home.

Using the Console to View the Update History of a VM Cluster

Learn how to view the history of patches applied on a VM cluster.

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

   VM Clusters is selected by default.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the VM cluster you are interested in.
4. Click **Update History**.

The history of patch operations for that VM cluster is displayed, along with the history of patch operations on its Database Homes.

Using the Console to View the Update History of a Database Home

Learn how to view the history of patches applied on a Database Home.

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

   VM Clusters is selected by default.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the VM cluster where the Database Home is located.
4. Click **Database Homes**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

A list of Database Homes is displayed.

5. In the list of Database Homes, click the Database Home you are interested in.
6. Click **Update History**.

The history of patch operations for that Database Home is displayed, along with the history of patch operations on the VM cluster to which it belongs.

Using the Console to Move a Database to Another Home

You can update the version of a VM cluster database by moving it to a Database Home that is running the version of Oracle Database you are interested in.

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

   VM Clusters is selected by default.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the VM cluster where the database you want to move is located.
4. Click **Database Homes**.
5. In the list of Database Homes, click the Database Home you are interested in.
6. Click **Databases**.
7. In the list of databases, click the database you are interested in.
8. Click Actions and then select **Move Database**.
9. Select the target Database Home.
10. Click **Move**.

    The database will be stopped in the current home and then restarted in the destination home.
11. 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 and Database statuses display as `Updating`. The Database Home location, shown under **Database Version**, displays as `Moving Database`. When the operation completes, Database Home is updated with the current home. Datapatch is executed automatically, as part of the database move, to complete post-patch SQL actions for all patches, including one-offs, on the new Database Home. If the database move operation is unsuccessful, then the status of the database displays as `Failed`, and the **Database Home** field provides information about the reason for the failure.

Using the API for Patching and Updating VM Cluster and Database Homes

Use various API features to help manage patching an Oracle Exadata Database Service on Cloud@Customer system.

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 patching VM clusters, Database Homes and Databases.

Copyright © 2008, 2026, Oracle and/or its affiliates.

VM cluster:

- `UpdateVmCluster`

Database Homes:

- `CreateDbHome`
- `UpdateDbHome`
- `DeleteDbHome`

Database:

- `CreateDatabase`
- `UpdateDatabase`
- `DeleteDatabase`

Use `UpdateVMCluster` to patch the Oracle Grid Infrastructure on the VM Cluster. Use `UpdateDbHome` to patch the Database Software of the Database Home. Use `UpdateDatabase` to move a database to a different Database Home, thereby updating the database to the same version as the target Database Home.

For the complete list of APIs for the Database service, see "Database Service API".

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- UpdateVmCluster
- CreateDbHome
- UpdateDbHome
- DeleteDbHome
- CreateDatabase
- UpdateDatabase
- DeleteDatabase
- Database Service API

Updating Guest VM Operating System

Learn to update the operating system image on Oracle Exadata Database Service on Cloud@Customer 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 that your Exadata Cloud@Customer VM cluster, database system, or Database Home meets the requirements for that patch. If the precheck is not successful, then 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Supported Software Versions and Update Restrictions
  Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):
- Using the Console to Update Guest VM Operating System
  To update the guest VM operating system with the Console, be prepared to provide values for the fields required.
- Using the Console to Rollback or Retry Failed Guest VM Operating System Update
  To update the guest VM operating system with the Console, be prepared to provide values for the fields required.
- Using the API to Update Guest VM Operating System
  Review the list of API calls to update guest VM operating system.

Supported Software Versions and Update Restrictions

Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):

> **Note**
>
> These are just the minimum requirements. If you want to update Grid Infrastructure and/or Oracle Database to meet the Exadata 23.1 requirements, then the recommendation is to update to the latest available versions of Grid Infrastructure and Oracle Database, and not to the minimum.

- **Exadata Image (Guest OS):** Exadata image release 22.1.0 (May 2022) or 21.2.10 (March 2022). Systems running versions older than 21.2.10 will first need to upgrade to at least 22.1.0 (May 2022) or 21.2.10 (March 2022) before updating to 23.1.0.0.0. This applies to both storage and database servers.
  - In addition to performing minor version updates to the Exadata VM Cluster images, you can update to a new major version if the currently installed version is 19.2 or higher. For example, if the VM cluster is on version 20, then you can update it to version 21.
  - The latest 4 (N to N-3) or more minor versions of each major version of the VM Cluster images are available through the console to apply.
- **Oracle Grid Infrastructure:** Exadata image release 23.1.0.0.0 supports the following minimum or newer Oracle Grid Infrastructure versions.
  - Release 19c: Version 19.15, April 2022 Release Update (RU) and newer (Default)
  - Release 21c: Version 21.6, April 2022 Release Update (RU) and newer
- **Oracle Database:** Exadata System Software 23.1 supports the following minimum versions or newer for new database installations.
  - Release 19c: Version 19.15, April 2022 Release Update (RU) and newer (Default)
  - Additional supported database releases under Market Driven Support or Quarterly Updates exception approval:
    * Release 12.2.0.1, Release Update (RU) 12.2.0.1.220118 (Jan 2022)
    * Release 12.1.0.2, Bundle Patch 12.1.0.2.220719 (Jul 2022) - requires patch 30159782

Copyright © 2008, 2026, Oracle and/or its affiliates.

    * Release 11.2.0.4, Bundle Patch 11.2.0.4.210119 (Jan 2021) - requires patch 30159782, patch 33991024

- If you have an Exadata infrastructure maintenance operation scheduled to start within the next 24 hours, then the Exadata Image update feature is not available.
- Once the VM cluster is upgraded to Exadata Database Service Guest VM OS 23.1, you will be able to add a new VM or a new database server to this VM cluster if Exadata Cloud@Customer Infrastructure is running an Exadata System Software version 22.1.16 and later.

> **Note**
>
> Upgrade to Exadata System Software 23.1 for Exadata Cloud@Customer Infrastructure will be available with February 2024 update cycle.

Using the Console to Update Guest VM Operating System

To update the guest VM operating system with the Console, be prepared to provide values for the fields required.

> **Note**
>
> Once the VM cluster is upgraded to Exadata Database Service Guest VM OS 23.1, you will be able to add a new VM or a new database server to this VM cluster if Exadata Cloud@Customer Infrastructure is running an Exadata System Software version 22.1.16 and later.
> Upgrade to Exadata System Software 23.1 for Exadata Cloud@Customer Infrastructure will be available with February 2024 update cycle.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
4. Click **Updates (OS)**.
5. Review the list of available software updates and locate the operating system 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 any time, and the precheck you run just before running a patch may find errors that the previous precheck did not find.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> If the precheck fails, the system displays a message in the **Apply Exadata OS Image Update** dialog that the last precheck has failed. Oracle recommends that you run the precheck again. Click the Actions icon (three dots)at the end of the row listing the OS patch to view the dialog.

- **Apply Exadata OS Image Update.** This link displays the Apply Exadata Image Update dialog that you use to apply the patch. The dialog shows the name of the database system you are patching, the current version of the database, and the new version of the database after the patch is applied. To start the process, click **Apply Exadata OS Image Update**.
- **Copy OCID.** This copies the Oracle Cloud ID. This can be used when troubleshooting a patch or to give to Support when contacting them.

> **Note**
>
> While the patch is running:
>
> – Run Precheck and Apply OS Image Update are not available. When the patch has completed, these actions are available again.
>
> – If the Exadata infrastructure containing this VM cluster is scheduled for maintenance that conflicts with the patching operation, the patch fails and the system displays a message explaining why. After the infrastructure maintenance is complete, run the patch operation again.

7. Confirm when prompted.

Using the Console to Rollback or Retry Failed Guest VM Operating System Update

To update the guest VM operating system with the Console, be prepared to provide values for the fields required.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
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 cluster you want to patch to display the cluster details.

   If applying update has failed, then on the **VM Cluster Details** page, a banner with the options **Roll Back** and **Retry Apply** is displayed.

   Choose an appropriate option.

   a. Click **Retry Apply**.

      Apply Exadata OS Image Update dialog is displayed with the options **Apply Exadata Image Update** and **Run Precheck**.

      Choose an appropriate option.

      (or)

   b. Click **Roll Back**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Confirm Rollback Operation dialog is displayed.

Click **Roll Back to confirm**.

4. You can also Apply Exadata Image Update from the **Updates (OS)** page.

Using the API to Update Guest VM Operating System

Review the list of API calls to update guest VM operating system.

For information about using the API and signing requests, see *REST APIs* and *Security Credentials*. For information about SDKs, see *Software Development Kits and Command Line Interface*.

Use these API operations to upgrade the Oracle Grid Infrastructure in a VM cluster and view the cluster's update history:

- `ListVmClusterUpdates`
- `GetVmClusterUpdate`
- `ListVmClusterUpdateHistoryEntries`
- `GetVmClusterUpdateHistoryEntry`
- `UpdateVmCluster`

For the complete list of APIs, see *Database Service API*.

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- ListVmClusterUpdates
- GetVmClusterUpdate
- ListVmClusterUpdateHistoryEntries
- GetVmClusterUpdateHistoryEntry
- UpdateVmCluster
- Database Service API

Upgrading Oracle Grid Infrastructure on an Oracle Exadata Database Service on Cloud@Customer VM Cluster

Learn to upgrade Oracle Grid Infrastructure on an Oracle Exadata Database Service on Cloud@Customer VM cluster using the Oracle Cloud Infrastructure Console or API.

Upgrading enables you to provision Oracle Database Homes and databases that use the most current Oracle Database software.

- About Upgrading Oracle Grid Infrastructure
  Upgrading the Oracle Grid Infrastructure (GI) on a VM cluster involves upgrading all the compute nodes in the instance. The upgrade is performed in a rolling fashion, with only one node being upgraded at a time.
- Using the Console to Manage Oracle Grid Infrastructure Upgrade
  You can use the Console to perform a precheck prior to upgrading your Oracle Grid Infrastructure (GI) and to perform the GI upgrade operation.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Using the API to Manage Oracle Grid Infrastructure Upgrade
  Review the list of API calls to manage Oracle Grid Infrastructure upgrade.

About Upgrading Oracle Grid Infrastructure

Upgrading the Oracle Grid Infrastructure (GI) on a VM cluster involves upgrading all the compute nodes in the instance. The upgrade is performed in a rolling fashion, with only one node being upgraded at a time.

- Oracle recommends running an upgrade precheck to identify and resolve any issues that would prevent a successful upgrade.
- You can monitor the progress of the upgrade operation by viewing the associated *work requests*.
- If you have an Exadata infrastructure maintenance operation scheduled to start within the next 24 hours, then the GI upgrade feature is not available.
- During the upgrade, you cannot perform other management operations such as starting, stopping, or rebooting nodes, scaling CPU, provisioning or managing Database Homes or databases, restoring a database, or editing IORM settings. The following Data Guard operations are not allowed on the VM cluster undergoing a GI upgrade:
  - Enable Data Guard
  - Switchover
  - Failover to the database using the VM cluster (a failover operation to standby on another VM cluster is possible)

Related Topics

- Work Requests Integration

Using the Console to Manage Oracle Grid Infrastructure Upgrade

You can use the Console to perform a precheck prior to upgrading your Oracle Grid Infrastructure (GI) and to perform the GI upgrade operation.

- Using the Console to Precheck Your VM Cluster Prior to Upgrading
- Using the Console to Upgrade Oracle Grid Infrastructure of a VM Cluster

Using the Console to Precheck Your VM Cluster Prior to Upgrading

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
4. Click **Updates**.
5. Click the Actions icon (three dots) at the end of the row listing the Oracle Grid Infrastructure (GI) upgrade, then click **Run Precheck**.
6. In the **Confirm** dialog, confirm you want to upgrade to begin the precheck operation.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Upgrade Oracle Grid Infrastructure of a VM Cluster

> **Note**
>
> - When planning to upgrade your Grid Infrastructure to 26ai, make sure that for each ASM diskgroup, `compatible.rdbms` has a value set to 19.0.0.0 and later.
> - Minimum requirements for upgrading Grid Infrastructure from 19c to 26ai:
>   - Exadata Guest VM running Exadata System Software 23.1.8
>   - Exadata Infrastructure running Exadata System Software 23.1.x

> **Note**
>
> Currently, Grid Infrastructure upgrade from 19c to 26ai is not supported for single node VM clusters.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
4. Click **Updates**.
5. Click the Actions icon (three dots) at the end of the row listing the Oracle Grid Infrastructure (GI) upgrade, then click **Upgrade Grid Infrastructure**.
6. In the Upgrade Grid Infrastructure dialog, confirm you want to upgrade the GI by clicking **Upgrade Grid Infrastructure**.

   If you haven't run a precheck, you have the option of clicking **Run Precheck** in this dialog to precheck your cloud VM cluster prior to the upgrade.

Using the API to Manage Oracle Grid Infrastructure Upgrade

Review the list of API calls to manage Oracle Grid Infrastructure upgrade.

For information about using the API and signing requests, see *REST APIs* and *Security Credentials*. For information about SDKs, see *Software Development Kits and Command Line Interface*.

Use these API operations to upgrade the Oracle Grid Infrastructure in a VM cluster and view the cluster's update history:

- `ListVmClusterUpdates`
- `GetVmClusterUpdate`
- `ListVmClusterUpdateHistoryEntries`
- `GetVmClusterUpdateHistoryEntry`

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `UpdateVmCluster`

For the complete list of APIs, see *Database Service API*.

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- ListVmClusterUpdates
- GetVmClusterUpdate
- ListVmClusterUpdateHistoryEntries
- GetVmClusterUpdateHistoryEntry
- UpdateVmCluster
- Database Service API

Upgrading Oracle Databases

Learn to upgrade an Exadata database instance to Oracle Database 19c and Oracle AI Database 26ai by using the Console and the API.

The upgrade is accomplished by moving the Exadata database to a Database Home that uses the target software version. For Oracle Database release and software support timelines, see *Release Schedule of Current Database Releases (Doc ID 742060.1)*.

- Prerequisites to Upgrade Oracle Databases
  Review the list of prerequisites to upgrade an Exadata Cloud@Customer Oracle Database instance.
- About Upgrading an Oracle Database
  Before you upgrade the database, become familiar with the following procedures to prepare your database for upgrade.
- How the Upgrade Operation Is Performed by the Database Service
  Familiarize yourself with what the Database service does during the upgrade process.
- Rolling Back an Oracle Database Unsuccessful Upgrade
  If your upgrade does not complete successfully, then you have the option of performing a rollback.
- After Upgrading an Oracle Database
  After a successful upgrade, note the following:
- Using the Console to Manage Oracle Database Upgrade
  Oracle recommends that you use the precheck action to ensure that your database has met the requirements for the upgrade operation.
- Using the API to Upgrade Oracle Databases
  Review the list of API calls to upgrade Oracle Databases.

Related Topics

- Release Schedule of Current Database Releases (Doc ID 742060.1)

Prerequisites to Upgrade Oracle Databases

Review the list of prerequisites to upgrade an Exadata Cloud@Customer Oracle Database instance.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- To upgrade to 19c, Oracle Linux 7 is the minimum requirement, and to upgrade to 26ai, Oracle Linux 8 is the minimum requirement. See *How to update the Exadata System Software (DomU) to 19 from 18 on the Exadata Cloud Service in OCI* for instructions on manually updating the operating system.
- The Oracle Grid Infrastructure can be version 19c or 26ai for Oracle Database 19c. However, the Oracle Grid Infrastructure must be version 26ai for Oracle AI Database 26ai. See Upgrading Exadata Grid Infrastructure for instructions on using the Oracle Cloud Infrastructure Console or API to upgrade Grid Infrastructure. If patches are available for your Grid Infrastructure, Oracle recommends applying them prior to performing a database upgrade.
- You must have an available Oracle Database Home that uses the four most recent versions of Oracle Database 19c or Oracle AI Database 26ai available in Oracle Cloud Infrastructure. See *Using the Console to Create Oracle Database Home on Exadata Cloud@Customer* for information on creating a Database Home. You can use Oracle-published software images or a *custom database software image* based on your patching requirements to create Database Homes.
- You must ensure that all pluggable databases in the container database that is being upgraded can be opened. Pluggable databases that cannot be opened by the system during the upgrade can cause an upgrade failure.

Your Oracle database must be configured with the following settings in order to upgrade:

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

See the *Oracle Database documentation* for your database's release version to learn more about these settings.

Related Topics

- How to update the Exadata System Software (DomU) to 19 from 18 on the Exadata Cloud Service in OCI (Doc ID 2521053.1)
- Using the Console to Create Oracle Database Home on Oracle Exadata Database Service on Cloud@Customer
  To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- Manage Software Images
- Oracle Database Documentation

About Upgrading an Oracle Database

Before you upgrade the database, become familiar with the following procedures to prepare your database for upgrade.

- Database upgrades involve database downtime. Keep this in mind when scheduling your upgrade.
- Oracle recommends that you back up your database and test the new software version on a test system or a cloned version of your database before you upgrade a production database. See *Creating an On-Demand Backup by Using the bkup_api Utility* for information on creating an on-demand manual backup.
- Oracle recommends running an upgrade precheck operation for your database prior to attempting an upgrade so that you can discover any issues that need mitigation prior to the time you plan to perform the upgrade. The precheck operation does not affect database availability and can be performed at any time that is convenient for you.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- If your databases uses Data Guard, you can upgrade either the primary or the standby first.
- An upgrade operation cannot take place while an automatic backup operation is underway. Before upgrading, Oracle recommends disabling automatic backups and performing a manual backup. See *Creating an On-Demand Backup by Using the bkup_api Utility* and *Customizing the Automatic Backup Configuration* for more information.
- After upgrading, you cannot use automatic backups taken prior to the upgrade to restore the database to an earlier point in time.
- If you are upgrading a database that uses version 11.2 software, the resulting version 19c database will be a non-container database (non-CDB).

Related Topics

- Configuring and Customizing Backups with dbaascli
  In addition to the console-based automated backup option, use dbaascli to configure and customize backups.
- Creating an On-Demand Backup by Using dbaascli
  You can use the `dbaascli` to create an on-demand backup of a complete database or an individual pluggable database (PDB):

How the Upgrade Operation Is Performed by the Database Service

Familiarize yourself with what the Database service does during the upgrade process.

- Executes an automatic precheck. This allows the system to identify issues needing mitigation and to stop the upgrade operation.
- Sets a guaranteed restore point, enabling it to perform a flashback in the event of an upgrade failure.
- Moves the database to a user-specified Oracle Database Home that uses the desired target software version.
- Runs the Database Upgrade Assistant (DBUA) software to perform the upgrade.

Rolling Back an Oracle Database Unsuccessful Upgrade

If your upgrade does not complete successfully, then you have the option of performing a rollback.

Details about the failure are displayed on the **Database Details** page in the Console, allowing you to analyze and resolve the issues causing the failure.

A rollback resets your database to the state prior to the upgrade. All changes to the database made during and after the upgrade will be lost. The rollback option is provided in a banner message displayed on the database details page of a database following an unsuccessful upgrade operation. See *Using the Console to Roll Back a Failed Database Upgrade* for more information.

Related Topics

- Using the Console to Roll Back a Failed Database Upgrade

After Upgrading an Oracle Database

After a successful upgrade, note the following:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Check that automatic backups are enabled for the database if you disabled them prior to upgrading. See *Customizing the Automatic Backup Configuration* for more information.
- Edit the Oracle Database

  `COMPATIBLE`

  parameter to reflect the new Oracle Database software version. See *What Is Oracle Database Compatibility?* for more information.
- If your database uses a *`database_name`*`.env` file, ensure that the variables in the file have been updated to point to the 19c Database Home. These variables should be automatically updated during the upgrade process.
- If you are upgrading a non-container database to Oracle Database version 19c, you can convert the database to a pluggable database after converting. See *How to Convert Non-CDB to PDB (Doc ID 2288024.1)* for instructions on converting your database to a pluggable database.
- If your old Database Home is empty and will not be reused, you can remove it. See *Using the Console to Delete an Oracle Database Home* for more information.

Related Topics

- Configuring and Customizing Backups with dbaascli
  In addition to the console-based automated backup option, use dbaascli to configure and customize backups.
- What Is Oracle Database Compatibility?
- How to Convert Non-CDB to PDB - Step by Step Example (Doc ID 2288024.1)
- Using the Console to Delete an Oracle Database Home
  To delete an Oracle Database home with the Console, use this procedure.

Using the Console to Manage Oracle Database Upgrade

Oracle recommends that you use the precheck action to ensure that your database has met the requirements for the upgrade operation.

- Using the Console to Run Oracle Database Upgrade Precheck or Perform Upgrade
- Using the Console to Roll Back a Failed Database Upgrade
- Using the Console to View the Upgrade History of a Database

Using the Console to Run Oracle Database Upgrade Precheck or Perform Upgrade

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
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 you want to upgrade.
4. In the list of databases on the **VM Cluster Details** page, click the name of the database you want to upgrade to view the **Database Details** page.
5. Click the Actions menu and then select **Upgrade**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

6. In the Upgrade Database dialog, select the following:
   - **Oracle Database version:** The drop-down selector lists only Oracle Database versions that are compatible with an upgrade from the current software version the database is using. The target software version must be higher than the database's current version.
   - **Target Database Home:** Select a Database Home for your database. The list of Database Homes is limited to those homes using the most recent versions of Oracle Database 19c software. Moving the database to the new Database Home results in the database being upgraded to the major release version and patching level of the new Database Home.
7. Click one of the following:
   - **Run Precheck:** This option starts an upgrade precheck to identify any issues with your database that need mitigation before you perform an upgrade.
   - **Upgrade Database:** This option starts upgrade operation. Oracle recommends performing an upgrade only after you have performed a successful precheck on the database.

Using the Console to Roll Back a Failed Database Upgrade

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the name of the VM cluster that contains the database with the failed upgrade.
4. Find the database that was unsuccessfully upgraded, and click its name to display details about it.
5. The database must display a banner at the top of the details page that includes a **Rollback** button and details about what issues caused the upgrade failure.
6. Click **Rollback**.
7. In the **Confirm rollback** dialog, confirm that you want to initiate a rollback to the previous Oracle Database version.

Using the Console to View the Upgrade History of a Database

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
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 you want to upgrade
4. In the list of databases on the VM Cluster Details page, click the name of the database for which you want to view the upgrade history.
5. On the Database Details page, click **Update History**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the API to Upgrade Oracle Databases

Review the list of API calls to upgrade Oracle 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 the following APIs to manage database upgrades::

- `getDatabaseUpgradeHistoryEntry`
- `ListDatabaseUpgradeHistoryEntries`
- `UpgradeDatabase`

For the complete list of APIs, see *Database Service API*.

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- getDatabaseUpgradeHistoryEntry
- ListDatabaseUpgradeHistoryEntries
- UpgradeDatabase
- Database Service API

Patching and Updating an Oracle Exadata Database Service on Cloud@Customer System Manually

This topic describes the procedures for patching and updating various components in Oracle Exadata Database Service on Cloud@Customer outside of the cloud automation. For information related to patching and updating with `dbaascli`, refer to "Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli."

> **Note**
>
> For more guidance on achieving continuous service during patching operations, see the *Application Checklist for Continuous Service for MAA Solutions* white paper.

- Updating Software Manually
  For daylight savings time and some routine or one-off patches it can be necessary for you to patch software manually.
- Updating the Guest VM Operating System Manually
  Learn about standard Exadata tools and techniques you can use to update the operating system components on the Oracle Exadata Database Service on Cloud@Customer Guest VMs.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli
  Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud@Customer system.
- Application Checklist for Continuous Service for MAA Solutions

Updating Software Manually

For daylight savings time and some routine or one-off patches it can be necessary for you to patch software manually.

To perform routine patching of Oracle Database and Oracle Grid Infrastructure software, Oracle recommends that you use the facilities provided by Oracle Exadata Database Service on Cloud@Customer. However, under some circumstances, it can be necessary for you to patch the Oracle Database or Oracle Grid Infrastructure software manually:

- **Daylight Savings Time (DST) Patching:** Because they cannot be applied in a rolling fashion, patches for the Oracle Database DST definitions are not included in the routine patch sets for Exadata Database Service on Cloud@Customer. If you need to apply patches to the Oracle Database DST definitions, you must do so manually. See My Oracle Support Doc ID 412160.1.
- **Non-routine or One-off Patching:** If you encounter a problem that requires a patch which is not included in any routine patch set, then work with Oracle Support Services to identify and apply the appropriate patch.

For general information about patching Oracle Database, refer to information about patch set updates and requirements in *Oracle Database Upgrade Guide* for your release.

Related Topics

- https://support.oracle.com/epmos/faces/DocumentDisplay?cmd=show&type=NOT&id=412160.1

Updating the Guest VM Operating System Manually

Learn about standard Exadata tools and techniques you can use to update the operating system components on the Oracle Exadata Database Service on Cloud@Customer Guest VMs.

You are responsible for managing patches and updates to the operating system environment on the Database Server Virtual Machines (VMs). For more information, read about updating Exadata Database Machine servers in *Oracle Exadata Database Machine Maintenance Guide*.

- Preparing for an Operating System Update
  To prepare for an operating system update for Oracle Exadata Database Service on Cloud@Customer, review this checklist of tasks.
- Updating the Operating System on All Virtual Machines of an Oracle Exadata Database Service on Cloud@Customer System
  To update the operating system on the Database Server Virtual Machines (VMs), use the `patchmgr` tool.
- Installing Additional Operating System Packages
  Review these guidelines before you install additional operating system packages for Oracle Exadata Database Service on Cloud@Customer.

Related Topics

- Updating Oracle Exadata Database Machine Database Servers

Copyright © 2008, 2026, Oracle and/or its affiliates.

Preparing for an Operating System Update

To prepare for an operating system update for Oracle Exadata Database Service on Cloud@Customer, review this checklist of tasks.

Before you update your operating system, do each of these preparation tasks:

Determine the latest software update. Before you begin an update, to determine the latest software to use, review Exadata Cloud Service Software Versions in My Oracle Support note 2333222.1.

Related Topics

- https://support.oracle.com/epmos/faces/DocumentDisplay?cmd=show&type=NOT&id=2333222.1

Updating the Operating System on All Virtual Machines of an Oracle Exadata Database Service on Cloud@Customer System

To update the operating system on the Database Server Virtual Machines (VMs), use the `patchmgr` tool.

> **Note**
>
> Customers who do not have My Oracle Support patch download privilege may obtain the Exadata patchmgr update utility and recent Exadata System Software releases using the Exadata Cloud@Customer Gen 2 utility `exadata_updates.sh`. For more information, see My Oracle Support Doc 2730739.1.

The `patchmgr` utility manages the entire update of one or more virtual machines remotely, including the pre-restart, restart, and post-restart steps of an Oracle Exadata Database Service on Cloud@Customer system.

You can run the utility either from one of your Oracle Exadata Database Service on Cloud@Customer virtual machines, or from another server running Oracle Linux. The server on which you run the utility is known as the **driving system**. You cannot use the driving system to update itself. Therefore, if the driving system is one of the virtual machines in a VM cluster that you are updating, then you must run the `patchmgr` utility more than once. The following scenarios describe typical ways of performing the updates:

- Non-Exadata Driving System
  The simplest way to run the update the system is to use a separate Oracle Linux server to update all virtual machines in one operation.
- Exadata Virtual Machine Driving System
  You can use one virtual machine to drive the updates for the rest of the virtual machines in the VM cluster. Then, you can use one of the updated nodes to drive the update on the original driving system. For example, consider updating a half rack system with four virtual machines; `node1`, `node2`, `node3`, and `node4`. You could first use `node1` to drive the updates of `node2`, `node3`, and `node4`. Then, you could use `node2` to drive the update of `node1`.

The driving system requires `root` user SSH access to each virtual machine being updated.

The following procedure is based on an example that assumes the following:

- The system has two virtual machines, `node1` and `node2`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- The target Exadata software version is `18.1.4.0.0.180125.3`.
- Each node is used as the driving system to update the other node.

1. Gather the environment details.
   a. Using SSH, connect to *node1* as the `opc` user.
   For detailed instructions, see *Connecting to a Compute Node with SSH*.
   b. Start a `root` user command shell.

      ```
      sudo su -
      ```

   c. Run the following command to determine the current Exadata software version.

      ```
      imageinfo -ver
      ```

      For example:

      ```
      # imageinfo -ver 19.2.13.0.0.200428
      ```

   d. Switch to the `grid` user, and identify all nodes in the cluster.

      ```
      su - grid
      ```

      ```
      olsnodes
      ```

      For example:

      ```
      olsnodes
      node1
      node2
      ```

2. Configure the driving system.
   a. Switch back to the `root` user on `node1` and check whether an SSH key pair (`id_rsa` and `id_rsa.pub`) exists. If not, then generate it.

      ```
      ls /root/.ssh/id_rsa*
      ls: cannot access /root/.ssh/id_rsa*: No such file or directory

      ssh-keygen -t rsa
      Generating public/private rsa key pair.
      Enter file in which to save the key (/root/.ssh/id_rsa):
      Enter passphrase (empty for no passphrase):
      Enter same passphrase again:
      Your identification has been saved in /root/.ssh/id_rsa.
      Your public key has been saved in /root/.ssh/id_rsa.pub.
      The key fingerprint is:
      93:47:b0:83:75:f2:3e:e6:23:b3:0a:06:ed:00:20:a5 root@node1.example.com
      The key's randomart image is:
      +--[ RSA 2048]----+
      |o..     + .      |
      |o.     o *       |
      |E     . o o      |
      ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
| . .     =       |
|  o .   S =      |
|   +     = .     |
|    +   o o      |
|   . .   + .     |
|      ...        |
+-----------------+
```

b. Distribute the public key to the target nodes, and verify this step. In the example, the only target node is `node2`.

```
scp -i ~root/.ssh/id_rsa.pub opc@node2:/tmp/id_rsa.node1.pub


ls -al /tmp/id_rsa.node1.pub
-rw-r--r-- 1 opc opc 442 Feb 28 03:33 /tmp/id_rsa.node1.pub


date
Wed Feb 28 03:33:45 UTC 2018
```

c. On the target node (`node2` in the example), add the root public key of `node1` to the root `authorized_keys` file.

```
cat /tmp/id_rsa.node1.pub >> ~root/.ssh/authorized_keys
```

d. Download `patchmgr` into `/root/patch` on the driving system (`node1` in this example). You can download the patchmgr bundle from Oracle Support by using My Oracle Support Patch ID 21634633. Always obtain the latest available Exadata patchmgr update utility to install any Exadata System Software release.

For further information, see also `dbnodeupdate.sh` and `dbserver.patch.zip`: Updating Exadata Database Server Software using the `DBNodeUpdate` Utility and `patchmgr`: My Oracle Support Doc ID 1553103.1.

e. Unzip the `patchmgr` bundle.
Depending on the version that you downloaded, the name of your ZIP file can differ.

```
cd /root/patch/18.1.4.0.0.180125.3
unzip dbserver.patch.zip
Archive:  p21634633_181400_Linux-x86-64.zip   creating:
dbserver_patch_5.180228.2/
creating: dbserver_patch_5.180228.2/ibdiagtools/
inflating: dbserver_patch_5.180228.2/ibdiagtools/cable_check.pl
inflating: dbserver_patch_5.180228.2/ibdiagtools/setup-ssh
inflating: dbserver_patch_5.180228.2/ibdiagtools/VERSION_FILE
extracting: dbserver_patch_5.180228.2/ibdiagtools/xmonib.sh
inflating: dbserver_patch_5.180228.2/ibdiagtools/monitord
inflating: dbserver_patch_5.180228.2/ibdiagtools/checkbadlinks.pl
creating: dbserver_patch_5.180228.2/ibdiagtools/topologies/
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/
VerifyTopologyUtility.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/verifylib.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Node.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Rack.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Group.pm
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Switch.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topology-zfs
inflating: dbserver_patch_5.180228.2/ibdiagtools/dcli
creating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
remoteScriptGenerator.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/CommonUtils.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
SolarisAdapter.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
LinuxAdapter.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
remoteLauncher.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
remoteConfig.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/spawnProc.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
runDiagnostics.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/OSAdapter.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/SampleOutputs.txt
inflating: dbserver_patch_5.180228.2/ibdiagtools/infinicheck
inflating: dbserver_patch_5.180228.2/ibdiagtools/ibping_test
inflating: dbserver_patch_5.180228.2/ibdiagtools/tar_ibdiagtools
inflating: dbserver_patch_5.180228.2/ibdiagtools/verify-topology
inflating: dbserver_patch_5.180228.2/installfw_exadata_ssh
creating: dbserver_patch_5.180228.2/linux.db.rpms/
inflating: dbserver_patch_5.180228.2/md5sum_files.lst
inflating: dbserver_patch_5.180228.2/patchmgr
inflating: dbserver_patch_5.180228.2/xcp
inflating: dbserver_patch_5.180228.2/ExadataSendNotification.pm
inflating: dbserver_patch_5.180228.2/ExadataImageNotification.pl
inflating: dbserver_patch_5.180228.2/kernelupgrade_oldbios.sh
inflating: dbserver_patch_5.180228.2/cellboot_usb_pci_path
inflating: dbserver_patch_5.180228.2/exadata.img.env
inflating: dbserver_patch_5.180228.2/README.txt
inflating: dbserver_patch_5.180228.2/exadataLogger.pm
inflating: dbserver_patch_5.180228.2/patch_bug_26678971
inflating: dbserver_patch_5.180228.2/dcli
inflating: dbserver_patch_5.180228.2/patchReport.py
extracting: dbserver_patch_5.180228.2/dbnodeupdate.zip
creating: dbserver_patch_5.180228.2/plugins/
inflating: dbserver_patch_5.180228.2/plugins/010-check_17854520.sh
inflating: dbserver_patch_5.180228.2/plugins/020-check_22468216.sh
inflating: dbserver_patch_5.180228.2/plugins/040-check_22896791.sh
inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_bash
inflating: dbserver_patch_5.180228.2/plugins/050-check_22651315.sh
inflating: dbserver_patch_5.180228.2/plugins/005-check_22909764.sh
inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_perl
inflating: dbserver_patch_5.180228.2/plugins/030-check_24625612.sh
inflating: dbserver_patch_5.180228.2/patchmgr_functions
inflating: dbserver_patch_5.180228.2/exadata.img.hw
inflating: dbserver_patch_5.180228.2/libxcp.so.1
inflating: dbserver_patch_5.180228.2/imageLogger
inflating: dbserver_patch_5.180228.2/ExaXMLNode.pm
inflating: dbserver_patch_5.180228.2/fwverify
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

f. In the directory that contains the `patchmgr` utility, create the `dbs_group` file, which contains the list of virtual machines to update. Include the nodes listed after running the `olsnodes` command in step 1, except for the driving system. In this example, `dbs_group` only contains `node2`.

```
cd /root/patch/18.1.4.0.0.180125.3/dbserver_patch_5.180228
cat dbs_group
node2
```

3. Run a patching precheck operation.

```
./patchmgr -dbnodes dbs_group -precheck -iso_repo zipped_iso_file -
target_version target-version -nomodify_at_prereq
```

> **Note**
>
> Run the precheck operation with the `-nomodify_at_prereq` option to prevent any changes to the system that could impact the backup you take in the next step. Otherwise, the backup might not be able to roll the system back to its original state, should it be necessary.

The output should look similar to the following example:

```
./patchmgr -dbnodes dbs_group -precheck -iso_repo /root/patch/
18.1.4.0.0.180125.3/exadata_ol6_18.1.4.0.0.180125.3_Linux-x86-64.zip -
target_version 18.1.4.0.0.180125.3 -nomodify_at_prereq

********************************************************************************
********************************
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the
latest release of dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.
********************************************************************************
********************************
2018-02-28 21:22:45 +0000        :Working: DO: Initiate precheck on 1
node(s)
2018-02-28 21:24:57 +0000        :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:26:15 +0000        :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:26:47 +0000        :Working: DO: dbnodeupdate.sh running a
precheck on node(s).
2018-02-28 21:28:23 +0000        :SUCCESS: DONE: Initiate precheck on
node(s).
```

4. Back up the current system.

```
./patchmgr -dbnodes dbs_group -backup -iso_repo zipped_iso_file -
target_version target-version -allow_active_network_mounts
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Ensure that you take the backup at this point, before any modifications are made to the system.

The output should look similar to the following example:

```
./patchmgr -dbnodes dbs_group -backup -iso_repo /root/patch/
18.1.4.0.0.180125.3/exadata_ol6_18.1.4.0.0.180125.3_Linux-x86-64.zip -
target_version 18.1.4.0.0.180125.3 -allow_active_network_mounts

********************************************************************************
********************************
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the
latest release of dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.
********************************************************************************
********************************
2018-02-28 21:29:00 +0000        :Working: DO: Initiate backup on 1
node(s).
2018-02-28 21:29:00 +0000        :Working: DO: Initiate backup on node(s)
2018-02-28 21:29:01 +0000        :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:30:18 +0000        :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:30:51 +0000        :Working: DO: dbnodeupdate.sh running a
backup on node(s).
2018-02-28 21:35:50 +0000        :SUCCESS: DONE: Initiate backup on
node(s).
2018-02-28 21:35:50 +0000        :SUCCESS: DONE: Initiate backup on 1
node(s).
```

5. Remove all custom RPMs from the target virtual machines. Custom RPMs are reported in precheck results. They include RPMs that were manually installed after the system was provisioned.
   - If you are updating the system from version 12.1.2.3.4.170111, and the precheck results include `krb5-workstation-1.10.3-57.el6.x86_64`, then remove it. This item is considered a custom RPM for this version.
   - Do not remove `exadata-sun-vm-computenode-exact` or `oracle-ofed-release-guest`. These two RPMs are handled automatically during the update process.
6. Perform the update. To ensure that the update process in not interrupted, use the command `nohup`. For example:

```
nohup ./patchmgr -dbnodes dbs_group -upgrade -nobackup -iso_repo
zipped_iso_file -target_version target-version -
allow_active_network_mounts &
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

The output should look similar to the following example:

```
nohup ./patchmgr -dbnodes dbs_group -upgrade -nobackup -iso_repo /root/
patch/18.1.4.0.0.180125.3/exadata_ol6_18.1.4.0.0.180125.3_Linux-x86-64.zip
-target_version 18.1.4.0.0.180125.3 -allow_active_network_mounts &

********************************************************************************
********************************
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the
latest release of dbserver.patch.zip)
NOTE
NOTE    Database nodes will reboot during the update process.
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.
********************************************************************************
******************************
2018-02-28 21:36:26 +0000        :Working: DO: Initiate prepare steps on
node(s).
2018-02-28 21:36:26 +0000        :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:37:44 +0000        :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:38:43 +0000        :SUCCESS: DONE: Initiate prepare steps on
node(s).
2018-02-28 21:38:43 +0000        :Working: DO: Initiate update on 1
node(s).
2018-02-28 21:38:43 +0000        :Working: DO: Initiate update on node(s)
2018-02-28 21:38:49 +0000        :Working: DO: Get information about any
required OS upgrades from node(s).
2018-02-28 21:38:59 +0000        :SUCCESS: DONE: Get information about any
required OS upgrades from node(s).
2018-02-28 21:38:59 +0000        :Working: DO: dbnodeupdate.sh running an
update step on all nodes.
2018-02-28 21:48:41 +0000        :INFO   : node2 is ready to reboot.
2018-02-28 21:48:41 +0000        :SUCCESS: DONE: dbnodeupdate.sh running
an update step on all nodes.
2018-02-28 21:48:41 +0000        :Working: DO: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000        :SUCCESS: DONE: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000        :Working: DO: Waiting to ensure node2 is
down before reboot.
2018-02-28 21:56:18 +0000        :Working: DO: Initiate prepare steps on
node(s).
2018-02-28 21:56:19 +0000        :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:57:37 +0000        :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:57:42 +0000        :SEEMS ALREADY UP TO DATE: node2
2018-02-28 21:57:43 +0000        :SUCCESS: DONE: Initiate update on node(s)
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

7. After the update operation completes, verify the version of the Exadata software on the virtual machine that was updated.

   ```
   imageinfo -ver
   18.1.4.0.0.180125.3
   ```

8. Repeat steps 2 through 7 of this procedure using the updated virtual machine as the driving system to update the remaining virtual machine. In this example update, you would now use `node2` to update `node1`.
9. As `root` On each virtual machine, run the `uptrack-install` command to install the available `ksplice` updates.

   ```
   uptrack-install --all -y
   ```

   ```
   uptrack-install --all -y
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- https://support.oracle.com/epmos/faces/DocumentDisplay?cmd=show&type=NOT&id=2730739.1
- https://support.oracle.com/epmos/faces/DocumentDisplay?cmd=show&type=NOT&id=1553103.1
- https://support.oracle.com/epmos/faces/ui/patch/PatchDetail.jspx?patchId=21634633

Installing Additional Operating System Packages

Review these guidelines before you install additional operating system packages for Oracle Exadata Database Service on Cloud@Customer.

You are permitted to install and update operating system packages on Oracle Exadata Database Service on Cloud@Customer as long as you do not modify the kernel or InfiniBand-specific packages. However, Oracle technical support, including installation, testing, certification and error resolution, does not apply to any non-Oracle software that you install.

Also be aware that if you add or update packages separate from an Oracle Exadata software update, then these package additions or updates can introduce problems when you apply an Oracle Exadata software update. Problems can occur because additional software packages add new dependencies that can interrupt an Oracle Exadata update. For this reason, Oracle recommends that you minimize customization.

If you install additional packages, then Oracle recommends that you have scripts to automate the removal and reinstallation of those packages. After an Oracle Exadata update, if you install additional packages, then verify that the additional packages are still compatible, and that you still need these packages.

For more information, refer to *Oracle Exadata Database Machine Maintenance Guide*.

Related Topics

- Installing, Updating and Managing Non-Oracle Software

Copyright © 2008, 2026, Oracle and/or its affiliates.

Resolving Dependency Issues Associated with Additional Non-Exadata Software Packages for Guest VM Upgrade

If you've installed non-Exadata software packages beyond those provided by Oracle, and the precheck fails during a Guest VM upgrade due to conflicts between and Oracle-installed RPMs, you can use the following procedure to resolve the conflicts and proceed with the upgrade.

For updates that do not change the major Oracle Linux version, this integrated capability allows you to update additional non-Exadata software packages as part of the Exadata database server update. It simplifies the handling of package dependency issues that may arise when such non-Exadata software packages are present on the system.

You can iteratively run precheck to identify and resolve dependency issues associated with the additional non-Exadata software packages. Once the required updates are understood, you can confidently perform the Exadata database server update and update the additional packages in a single, coordinated operation.

Ensure that the configuration file exists on the target server to trigger the setup of a temporary YUM repository for non-Exadata software packages.

- **File Location:** `/etc/exadata/additional-packages.txt`
- **Ownership and Permissions:** This file must be owned and modifiable by the `root` user only.

If the file exists, it is used to collect information about the required non-Exadata software packages and to set up and enable a temporary YUM repository. If the file is not present, no repository is configured.

You may also create a symbolic link at `/etc/exadata/additional-packages.txt` that points to a configuration file located elsewhere—typically on a shared mount.

File Format
The file must contain a list of non-Exadata software packages, with each entry on a new line. Supported formats include:

- `http(s)://path/to/package.rpm`: Full URL to the RPM file
- `/full/path/to/package.rpm`: Absolute path to a local RPM file
- `repo:package.rpm`: Reference to a package in an existing YUM repository

> **Note**
>
> - If using the `repo:` format, ensure the referenced repository is defined in the target server's YUM configuration.
> - Local files can reside in standard local directories, NFS mounts, or ACFS mounts.

Example: `additional-packages.txt`

```
/u01/elfutils-debuginfod-client-0.190-2.el8.x86_64.rpm
/u01/elfutils-libelf-devel-0.190-2.el8.x86_64.rpm
/u01/keyutils-libs-devel-1.5.10-9.0.1.el8.x86_64.rpm
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
https://example.com/packages/krb5-devel-1.18.2-28.0.1.el8_10.x86_64.rpm
https://example.com/packages/memstrack-0.2.5-2.el8.x86_64.rpm
/u01/pigz-2.4-4.el8.x86_64.rpm
/u01/sssd-nfs-idmap-2.9.4-3.0.1.el8_10.x86_64.rpm
https://example.com/packages/timedatex-0.5-3.el8.x86_64.rpm
https://example.com/packages/zlib-devel-1.2.11-25.el8.x86_64.rpm
```

Interim Software Updates

For authorized environments, learn how to download interim software updates.

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

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

- Create an Interim Software Update
- Download an Interim Software Update
- Delete an Interim Software Update
- Move an Interim Software Update Resource to Another Compartment
- Using the API to Manage Interim Software Updates

Create an Interim Software Update

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

   VM Clusters is selected by default.
2. Under **Resources**, click **Interim software updates**.

   Manual 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**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Download an Interim Software Update

The patch download path is valid for four days. Download the patch within the specified timeframe.

1. On the Interim Software Updates page, click the Actions menu (three dots) of the Interim Software Update you want to download and then select **Download**.

   The system starts downloading the patch.

2. You can also download a patch from the Interim Software Updates details page.

   - Click the Actions button (three dots) for the patch you're interested in, and select **Download**.

> **Note**
>
> You can only download the patches that are in **Available** state.

Interim Software Updates Lifecycle States:

- **Available**: Patch has been created successfully and the time-to-live (TTL) has not expired.
- **Creating**: The patch creation process is in progress.
- **Expired**: The lifetime of the patch download link has expired, which means you cannot download it.
- **Failed**: The patch create failed due to some error.
- **Terminating**: The patch deletion process is in progress.
- **Terminated**: The patch has been deleted.

Delete an Interim Software Update

Be discrete in deleting interim software updates. However, you can delete the interim software updates that have expired to free up space in the Object Store.

1. On the Interim Software Updates page, click the Actions menu (three dots) of the Interim Software Update you want to delete and then select **Delete**.
2. In the resulting dialog, enter the name of the patch to confirm and then click **Delete**.
3. You can also delete a patch from the Interim Software Updates details page.

   - Click the Actions button (three dots) for the patch you're interested in, and select **Delete**.

Move an Interim Software Update Resource to Another Compartment

1. On the Interim Software Updates page, click the Actions menu (three dots) of the Interim Software Update you want to move and then select **Move resource**.
2. In the resulting dialog, choose a new compartment, and click **Move Resource**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

3. You can also move a patch from the Interim Software Updates details page.
   - Click the Actions button (three dots) for the patch you're interested in, and select **Move Resource**.

Using the API to Manage Interim Software Updates

ExaDB-C@C and ExaDB-D use the same API to manage interim software updates.

For information about using the API and signing requests, see *REST APIs* and *Security Credentials*. For information about SDKs, see *Software Development Kits and Command Line Interface*.

Use these API operations to manage interim software updates:

- `CreateOneoffPatch`
- `DeleteOneoffPatch`
- `DownloadOneoffPatch`
- `UpdateOneoffPatch`
- `ListOneoffPatches`
- `GetOneoffPatch`
- `ChangeOneoffPatchCompartment`

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- OneoffPatch Reference

Use Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer

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

- About Using Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer
  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 Cloud@Customer
  Review the list of prerequisites for using Data Guard with Oracle Exadata Database Service on Cloud@Customer.
- Working with Data Guard
  Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.
- Support for Concurrent Data Guard Operations in Multiple Standby Environments

Copyright © 2008, 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.
- Using the Console to Manage an Oracle Data Guard Group
  Learn how to enable a Data Guard Group between databases, change the role of a database in a Data Guard Group using either a switchover or a failover operation, and reinstate a failed database.
- Using the API to Manage Data Guard Associations on an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to use the API to manage Data Guard associations on an Oracle Exadata Database Service on Cloud@Customer system.
- Using the API to Manage Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer System
  Learn how to use the API to manage Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer system.

About Using Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer

Oracle Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruptions.

Oracle Data Guard maintains these standby databases as copies of the production database. Then, if the production database becomes unavailable because of a planned or an unplanned outage, Oracle Data Guard can switch any standby database to the production role, minimizing the downtime associated with the outage. Oracle Data Guard can be used with traditional backup, restoration, and cluster techniques to provide a high level of data protection and data availability. Oracle Data Guard transport services are also used by other Oracle features such as Oracle Streams and Oracle GoldenGate for efficient and reliable transmission of redo from a source database to one or more remote destinations.

For complete information on Oracle Data Guard, see the Oracle Data Guard Concepts and Administration documentation and Oracle Data Guard Broker Concepts on the Oracle Database Documentation portal.

This topic explains how to use the Console or the API to 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 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Prerequisites for Using Oracle Data Guard with Oracle Exadata Database Service on Cloud@Customer

Review the list of prerequisites for using Data Guard with Oracle Exadata Database Service on Cloud@Customer.

- VM Clusters
  A VM cluster Data Guard implementation requires two Exadata database VM Clusters, one containing the primary database and one containing the standby database.
- Password
  To change the SYS password or rotate TDE keys, use OCI API.
- Adding a Node to a VM Cluster
  When adding a node to a VM cluster, an instance of the Data Guard database is automatically created on the new node.
- Removing a Node from a VM Cluster
  When removing a node from a VM cluster, the instance and it's metadata on the removing node is deleted automatically.

VM Clusters

A VM cluster Data Guard implementation requires two Exadata database VM Clusters, one containing the primary database and one containing the standby database.

> **Note**
>
> Oracle strongly recommends the primary and standby databases for any production workloads be on different Exadata Cloud Infrastructures for better fault isolation and disaster protection.

Password

To change the SYS password or rotate TDE keys, use OCI API.

Adding a Node to a VM Cluster

When adding a node to a VM cluster, an instance of the Data Guard database is automatically created on the new node.

Related Topics

- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.

Removing a Node from a VM Cluster

When removing a node from a VM cluster, the instance and it's metadata on the removing node is deleted automatically.

Related Topics

- Using the Console to Scale Infrastructure Compute and Storage
  To scale infrastructure compute and storage, complete this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Working with Data Guard

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

The primary and standby databases constitute a Data Guard Group. Most of your applications access the primary database. A standby database is a transactionally consistent copy of the primary database.

Data Guard maintains the standby database by transmitting and applying redo data from the primary database. If the primary database becomes unavailable, you can use Data Guard to switchover or failover the standby database to the primary role. This is true even if you have more than one standby database.

- Switchover
  A switchover reverses the primary and standby database roles.
- Failover
  A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable.
- Reinstate
  Reinstates a database into the standby role in a Data Guard Group.

Switchover

A switchover reverses the primary and standby database roles.

Each database continues to be part of the Data Guard Group in its new role. A switchover ensures no data loss. You can use a switchover before you perform planned maintenance on the primary database. Performing planned maintenance on a Exadata database compute node with a Data Guard Group is typically done by switching the primary to the standby role, performing maintenance on the standby, and then switching it back to the primary role.

Failover

A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable.

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

Reinstate

Reinstates a database into the standby role in a Data Guard Group.

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> You cannot terminate a primary database that is part of a Data Guard Group that contains one or more standby databases. You will have to terminate the standby databases first. Alternatively, you can switch over the primary database to the standby role, and then terminate the former primary.
>
> You cannot terminate a VM cluster that includes Data Guard enabled databases. You must first terminate the standby databases that are part of the Data Guard Group.

Support for Concurrent Data Guard Operations in Multiple Standby Environments

In addition to the Support for Concurrent Data Guard, Container Database (CDB), and Pluggable Database (PDB) Operations, these enhancements enable you to perform concurrent operations on CDBs and PDBs alongside Data Guard migration operations in environments with multiple standby databases.

You can now perform the following operations in parallel within the same Oracle Home:

- Perform Add Standby operations concurrently on different databases.
- Create or delete a CDB while an Add Standby operation is running on another database, and vice versa.
- Create or delete a PDB while an Add Standby operation is running on another database, and vice versa.
- Performing Data Guard actions (switchover, failover, reinstate, convert to snapshot and convert to physical) while an Add Standby operation is running on another database, and vice versa.

Similarly, you can perform the following Migrate Data Guard operations in parallel:

- Perform migrate Data Guard operations concurrently on different databases.
- Perform migrate Data Guard operations while an Add Standby operation is running on another database, and vice versa.
- Create or delete a CDB while a migrate Data Guard operations is running on another database, and vice versa.
- Create or delete a PDB while a migrate Data Guard operations is running on another database, and vice versa.
- Performing Data Guard actions (switchover, failover, reinstate, convert to snapshot and convert to physical) while a migrate Data Guard operations is running on another database, and vice versa.

Enhanced Data Guard Health Status Reporting

The enhanced Data Guard health status reporting provides comprehensive insights into protection mode, switchover and failover readiness, and data loss exposure across primary and standby databases.

With clear visual indicators (green, yellow, red, gray), you can quickly assess the readiness of your databases for role transitions and failover events. Additionally, detailed redo transport lag

Copyright © 2008, 2026, Oracle and/or its affiliates.

metrics help you evaluate potential data loss scenarios, enabling proactive disaster recovery planning and improved operational reliability.

Protection Mode

Data Guard associations provide two protection modes:

- Maximum Availability – Enables zero data loss failover and uses synchronous redo transport to at least one standby database.
- Maximum Performance – Enables near-zero data loss failover with no performance impact by using asynchronous transport to all standby targets.

Switchover Readiness

- Primary Database: status indicators – Green, Yellow, Red, Gray.
  - Green (Healthy): All standby databases have passed all switchover readiness checks.
  - Yellow (Warning): Applicable only if the primary database has multiple standby databases. A subset of standby databases is switchover-ready. Address the failed checks.
  - Red (Critical): None of the standby databases are switchover-ready. Address failed checks to avoid extended downtime.
  - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Standby Database: status indicators – Green, Red, Gray.
  - Green (Healthy): The standby database has passed all switchover readiness checks.
  - Red (Critical): This standby database is not switchover-ready. Address failed checks to avoid extended downtime during switchover attempts.
  - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Disabled Standby Database and Snapshot Standby Database: status indicators – Gray
  - Gray (Unknown): Disabled standby databases always appear as Unknown because their health status cannot be determined.

Failover Readiness

- Primary Database: status indicators – Green, Yellow, Red, Gray.
  - Green (Healthy): All standby databases have passed all failover readiness checks.
  - Yellow (Warning): Applicable only if the primary database has multiple standby databases. A subset of standby databases is failover-ready. Address the failed checks.
  - Red (Critical): None of the standby databases are failover-ready. Address failed checks to avoid extended downtime.
  - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Standby Database and Snapshot Standby: status indicators – Green, Red, Gray.
  - Green (Healthy): The standby database has passed all failover readiness checks.
  - Red (Critical): This standby database is not failover-ready. Address failed checks to avoid extended downtime during failover.
  - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Disabled Standby Database: status indicators – Gray
  - Gray (Unknown): Disabled standby databases always appear as Unknown because their health status cannot be determined.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> When the health status cannot be determined for any database in a Data Guard Group at a given point in time, the status appears as null in the SDK/CLI, and Terraform, and as UNKNOWN (gray icon) in the Console.

Data Loss Exposure

- Primary Database: Data loss exposure is defined as the redo transport lag between the primary and standby databases. "Last Computed" if displayed represents the last time the metric could be retrieved and calculated.
- Standby Database: Data loss exposure is defined as the redo transport lag between the primary and standby databases. "Last Computed" if displayed represents the last time the metric could be retrieved and calculated.

Refresh Data Guard health status

Run an explicit refresh to get the latest status.

Using the Console to Manage an Oracle Data Guard Group

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

When you enable Data Guard, a separate Data Guard Group is created between the primary and the standby databases.

- Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System
  Learn to set up a Data Guard Group between databases.
- Using the Console to View Data Guard Group of Databases in an Exadata VM Cluster
  To view the role of each database in a Data Guard Group in an Exadata VM Cluster, follow this procedure.
- Using the Console To View and Edit Data Guard Group Details
  You can switch between Data Guard types based on the Oracle Database software license type you have deployed.
- Using the Console To Perform a Database Switchover
  You can initiate a switchover operation on a standby database that is a member of the Data Guard Group.
- Using the Console To Perform a Database Failover
  You can initiate a failover operation on a standby database that is a member of the Data Guard Group.
- Using the Console 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.
- Using the Console To Terminate a Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer System
  On a VM cluster, you remove a Data Guard Group by terminating all the standby databases.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Using the Console to Enable Data Guard on an Oracle Exadata Database Service on Cloud@Customer System

Learn to set up a Data Guard Group between databases.

> **Note**
>
> Data Guard relies on a reliable network with sufficient throughput between the primary and standby clusters. Since Oracle does not own the network, some evaluation should be done prior to implementing Data Guard to ensure the required network bandwidth is available. It is recommended that Assessing and Optimizing Network Performance be followed to understand the achievable throughput between the clusters and evaluate whether the requirements of the database are met. By default, the maximum socket buffer size is set to a higher value for cross-region ExaDB-C@C Data Guard configurations.

> **Note**
>
> 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.
>
> As part of the latest release we are introducing an enhanced user experience and new APIs to improve performance and provide additional Data Guard capabilities including support for multiple standby databases via cloud automation.
>
> - With the new API, your new Data Guard configuration will be created as a Data Guard Group resource.
> - If you have an existing Data Guard setup, you can continue to use current capabilities with no impact. However, if you wish to create multiple standby databases, you will need to migrate to the new API model, which can be done at any time.
> - If you currently have automation that manages Data Guard operations using the existing Data Guard Association API, you will need to update your applications to use the new API to take advantage of these new capabilities
> Oracle currently supports both the existing Data Guard Association API and the new Data Guard Group API and the associated user interfaces.

> **Note**
>
> A parallel operation on the Standby, if it fails, should be retried after a 5-minute interval.

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

   VM Clusters is selected by default.
2. Select your **Compartment**.

Copyright © 2008, 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 VM cluster that contains the database for which you want to assume the primary role, and then click the name of that database.
4. Click **Data Guard associations**.
5. Click **Add standby**.
6. On the **Add standby** page, configure your Data Guard Group.
   - To explicitly run a Data Guard precheck, click **Run Precheck**. When you click Run Precheck, the following actions are performed:
     - The system closes the Add Standby page.
     - A message appears on the Primary Database Details page indicating that the Data Guard precheck is in progress, and the **Add Standby** button is disabled.

     Click the **View Details** link to navigate to the Work Requests page for the status:
     - Review the precheck status: **Succeeded** or **Failed**.
     - If the precheck has failed, then resolve the reported errors and retry the Data Guard setup.

     > **Note**
     >
     > While the precheck is running, the primary database and the VM Cluster remain in the **AVAILABLE** state.

   - – **Choose the Data Guard experience:**
     - **Use the new Data Guard Group Resource** With this option, your new Data Guard configuration will be created as a Data Guard Group resource. This option with new APIs supports adding multiple standby databases and provides other enhancements. If you currently have automation that manages Data Guard operations using the existing Data Guard Association API, you can update your applications to use the new API to take advantage of these new capabilities.
     - **Use the existing Data Guard Association Resource** Choose this option if your automation for managing Data Guard operations relies on the existing Data Guard Association API. However, you will not be able to add multiple standby databases and will not get the enhancements provided by the new API.
   - **Data Guard Group details:**
     - Select a Data Guard type, **Active Data Guard or Data Guard**, based on the Oracle Database software license type you have deployed. If you have deployed Oracle Database Enterprise Edition Extreme Performance (License Included), then you may choose either Data Guard or Active Data Guard. If you have deployed Bring Your Own License (BYOL) Oracle Database Enterprise Edition without the Active Data Guard option, then you would select Data Guard, which is the default.
       - **Active Data Guard:** Active Data Guard is a licensed option to the Oracle Database Enterprise Edition and enables advanced capabilities that extend the basic Data Guard functionality. These capabilities include Real-Time Query and DML Offload, Automatic Block Repair of physical data corruptions, Standby Block Change Tracking, Global Data Services, and Application Continuity.

Copyright © 2008, 2026, Oracle and/or its affiliates.

    - **Data Guard:** Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data. 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. Data Guard maintains these standby databases as transactionally consistent copies of the production database.

  - **Protection mode:** The protection mode used for this Data Guard Group.

    Maximum Performance provides the highest level of data protection that is possible without affecting the performance of a primary database.

    Maximum Availability provides the highest level of protection of data with zero data loss synchronous transport without compromising the availability of the database.

  - **Transport type:** The redo transport type used for the Data Guard association between the primary and this standby database.

    - **Async:** Asynchronous transport mode used with Maximum Performance protection mode.
    - **Sync:** Synchronous transport mode used with Maximum Performance and Maximum Availability protection mode.

    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.
      - The default settings are:
        - **Protection Mode:** Max Performance
        - **Transport Type:** Async
    - **Creating the second to Nth standby:** You cannot modify the Protection Mode or Transport Type for any subsequent standby databases.
      - The Protection Mode is inherited from the first standby.
      - The default Transport Type is set to Async.

- **Select peer VM cluster:** Specify the following values for the standby:
  - **Peer Region:** The primary and standby databases could be running on two different VM clusters on a shared ExaDB-C@C system or on two geographically separated ExaDB-C@C systems managed from the same or different Oracle Cloud Infrastructure regions.
  - **Peer Exadata Cloud@Customer infrastructure**: Select the Exadata Database Service on Cloud@Customer infrastructure where the standby database is located. Click the **CHANGE COMPARTMENT** hyperlink to choose a compartment.
  - **Peer VM Cluster:** Select the Exadata database compute node that contains the standby database. Click the **CHANGE COMPARTMENT** hyperlink to choose a compartment.
- **Choose Database Home:** Select an existing Database Home or create one as applicable.
  - **Select an existing Database Home:** If one or more Database Homes already exist for the database version you have selected, then this option is selected by default. And, you will be presented with a list of Database Homes. Select a Database Home from the list.

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> Although only Database homes of the same version and RU are listed, the homes displayed may have different one-off patches than the primary. Though acceptable to have different one-offs, the best practice is to have identical database homes between primary and standby.

- **Create a new Database Home:** If no Database Homes exist for the database version you have selected, then this option is selected by default. You can create the new Database Home with the same DSI as the primary database, or choose a different image. Note that DSIs are not available across regions. A separate DSI must be created in the peer region using the same RU as the primary.
  a. Click **Change Database Image** to select a database software image for the new Database Home.
  b. In the resulting **Browse Database Images**, do the following:
     i. Select the compartment containing the database software image you want to use to create the new Database Home.
     ii. 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.
     iii. Click **Select**.

> **Note**
>
> If you are using the new Data Guard Group resource, you must first create the database home before adding the standby database.

- **Configure standby database:**
  - **Provide a unique name for the database:**

> **Note**
>
> You cannot modify the `db_name`, `db_unique_name`, and SID prefix after creating the database.

Optionally, specify a unique name for the database. This attribute defines the value of the `db_unique_name` database parameter. The value is case insensitive. The `db_unique_name` must contain only the permitted characters.

Review the following guidelines when selecting a database name:

* maximum of 30 characters
* can contain alphanumeric and underscore (_) characters
* begin with an alphabetic character
* unique across the fleet/tenancy

If a unique name is not entered, then the `db_unique_name` defaults to the following format <*db_name*>_<*3 char unique string*>_<*region-name*>.

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - **Database password:** Enter the database admin password of the primary database in the Database password field. This same database admin password will be used for the standby.
    The admin password and the TDE password must be the same. If they are not, follow the instructions in *Changing the Database Passwords* to align them.
  - **TDE wallet password:** Enter the TDE wallet password.
- (Optional) Select **Show Advanced Options**.
  - **Provide the Oracle SID prefix:** Optionally, specify the Oracle SID prefix for the database. The instance number is automatically appended to the SID prefix to become the `instance_name` database parameter. If not provided, then the SID prefix defaults to the first 12 characters of the `db_unique_name`.
    Review the following guidelines when selecting a database name:
    * maximum of 12 characters
    * contain only alphanumeric characters
    * begin with an alphabetic character
    * unique in the VM cluster
  - **Tags:** (Optional) You can choose to apply tags. If you have permissions to create a resource, 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 if you should apply tags, then skip this option (you can apply tags later), or ask your administrator.

7. Click **Add**.

A work request is issued to configure the Data Guard association. The progress of the request and the stages of provisioning can be viewed on the Work Requests page of the respective Standby database.

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

Related Topics

- Network Requirements for Oracle Exadata Database Service on Cloud@Customer
  To provide secure and reliable network connectivity for different application and management functions, Exadata Database Service on Cloud@Customer uses different networks.
- Changing the Database Passwords
  To change the SYS password, or to change the TDE wallet password, use this procedure.

Using the Console to View Data Guard Group of Databases in an Exadata VM Cluster

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

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click on the VM Cluster containing the databases you wish to view their roles in Data Guard associations.

Copyright © 2008, 2026, Oracle and/or its affiliates.

4. Click **Databases**.

   The role of each database in this VM Cluster is indicated in the **Data Guard role** column.

Using the Console To View and Edit Data Guard Group Details

You can switch between Data Guard types based on the Oracle Database software license type you have deployed.

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

   VM Clusters is selected by default.

2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.

3. In the list of VM clusters, click the VM cluster that contains the primary database you want to switch Data Guard type.

4. Click the name of the primary database.

5. Click **Data Guard Group**.

   A list of databases that are members of the Data Guard Group is displayed with the **Data Guard type** you have chosen for each Data Guard Group member.

6. To edit Data Guard Group details, click the Actions icon (three dots), and then click **Edit**.

   Edit Data Guard Group screen is displayed.

7. Do the following on the Edit Data Guard Group screen.

   - Select an applicable **Data Guard type**.
   - Select the **Protection mode**.

     > **Note**
     >
     > Select the **Transport type**.

   - Set the **Database password**.
   - Click **Edit Data Guard** to save the changes.

Using the Console To Perform a Database Switchover

You can initiate a switchover operation on a standby database that is a member of the Data Guard Group.

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

   VM Clusters is selected by default.

2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.

3. In the list of VM clusters, click the VM cluster that contains the primary database you want to switch over.

4. Click the name of the primary database.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. Click **Data Guard Group**.
6. Select the standby database in the Data Guard Group on which you want to perform a switchover. Click the Actions icon (three dots), and then click **Switchover**.
7. In the Switchover database dialog box, enter the database admin password, and then click **Switchover**.

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

Using the Console To Perform a Database Failover

You can initiate a failover operation on a standby database that is a member of the Data Guard Group.

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

   VM Clusters is selected by default.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the VM cluster that contains the primary database's peer standby you want to fail over to.
4. Click the name of the standby database.
5. Click **Data Guard Group**.
6. Select the standby database in the Data Guard Group on which you want to perform a failover. Click the Actions icon (three dots), and then click **Failover**.
7. In the Failover database dialog box, enter the database admin password, and then click **Failover**.

> **Note**
>
> You can initiate a failover even if the primary database is in a healthy state; however, exercise caution when performing a failover.

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

Using the Console 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.

Before you can reinstate a version 12.2 or later database, you must perform some steps on the database host to stop the database or start it in `MOUNT` mode.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Set your `ORACLE_UNQNAME` environment variable to the value of the Database Unique Name, and then run these commands:

```
srvctl stop database -d db-unique-name -o abort


srvctl start database -d db-unique-name -o mount
```

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

   VM Clusters is selected by default.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the VM cluster that contains the primary database.
4. Click the name of the primary database.
5. Click **Data Guard Group**.

   You will see the database you want to reinstate listed.
6. Click the Actions icon (three dots), and then click **Reinstate**.
7. In the Reinstate database dialog box, enter the database admin password, and then click **Reinstate**.

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

Using the Console To Terminate a Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer System

On a VM cluster, you remove a Data Guard Group by terminating all the standby databases.

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

   VM Clusters is selected by default.
2. Choose your **Compartment**.

   A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the VM cluster that contains the standby database you want to terminate.
4. Click the name of the standby database.
5. For the standby database you want to terminate, click the Actions icon (three dots), and then click **Terminate**.
6. In the Terminate Database dialog box, enter the name of the database, and then click **Terminate**.

Using the API to Manage Data Guard Associations on an Oracle Exadata Database Service on Cloud@Customer System

Learn how to use the API to manage Data Guard associations on an Oracle Exadata Database Service on Cloud@Customer system.

Copyright © 2008, 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.

The following table lists the REST API endpoints to manage Data Guard associations.

Operation	REST API Endpoint
Create a Data Guard association.	CreateDataGuardAssociation
View details of the specified Data Guard association's configuration information.	GetDataGuardAssociation
View the list of all Data Guard associations for the specified database.	ListDataGuardAssociations
Perform a switchover to transition a primary database of a Data Guard association into standby role.	SwitchoverDataGuardAssociation
Perform a failover to transition a standby database identified by the `databaseId` parameter into the specified Data Guard association's primary role after the existing primary database fails or becomes unreachable.	FailoverDataGuardAssociation
Reinstate a database identified by the `databaseId` parameter into standby role in a Data Guard association.	ReinstateDataGuardAssociation For more information, see Using the Console To Reinstate a Database.
Delete a standby database.	DeleteDatabase

For the complete list of APIs, see Database Service API.

Using the API to Manage Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer System

Learn how to use the API to manage Data Guard Group on an Oracle Exadata Database Service on Cloud@Customer system.

> **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.

Copyright © 2008, 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.

The following table lists the REST API endpoints to manage Data Guard Group.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
Create / Add Standby Database	CreateDatabase	• The same API is being used for creating a first standby and adding m

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
		orestandbydatabases. It uses the existing created database A

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment

Plwithsourceas"DATAGUARD".

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
Data Guard Action - Failover	FailoverDataGuard	Failover should be triggered on respective standby that to be c

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
		omeprimary.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
Data Guard Action - Reinstate	ReinstateDataGuard	Reinstate should be triggered on respective standby to be reinstated

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
		stated.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
Data Guard Action - Switchover	SwitchOverDataGuard	Switchover should be triggered on respective standby that to b

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
		ecomeprimary.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
Delete Standby	DeleteDatabase	• Delete standby remains same as existing. - DeleteDatabase • Cal

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
		Ignore respective standby to be deleted.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
Migrate Data Guard Association to multiple Standby	MigrateDataGuardAssociationToMultiDataGuards	• Migrate the existing Data Guard association to Data Guard Gro

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
		upmodel. Newstandbycanbeaddedonlyaftermigrationiscomp

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
Update Data Guard Group Configuration	UpdateDataGuard	It takes either standby or primary database OCID to update the c

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint	Comment
		onfiguration.

For the complete list of APIs, see Database Service API.

Manage Keys Using an External Keystore

Review the use cases and implementation details of an external keystore.

- Configuring an External Keystore
  This configuration process is crucial for using external keystores to manage and secure encryption keys for your databases on Exadata systems. Ensure proper installation, password configuration, and communication setup for seamless operation.
- Storing the Keys in an External Keystore
  You can now encrypt Oracle Databases on Oracle Exadata Database Service on Cloud@Customer by storing the keys in an external keystore.
- Restrictions on Adding a Virtual Machine to a VM Cluster Configured with an External Keystore
  When a database is protected by an external keystore, adding a new virtual machine (VM) to the cluster is restricted.

Configuring an External Keystore

This configuration process is crucial for using external keystores to manage and secure encryption keys for your databases on Exadata systems. Ensure proper installation, password configuration, and communication setup for seamless operation.

Work with your external keystore vendor for detailed configuration steps. The details outlined below provides a high-level overview of the generic steps required to configure an external keystore.

External Keystore Server Installation: You are responsible for installing and configuring the external keystore server using the documentation provided by the vendor.

External Keystore Password Format: The format of the external keystore password varies depending on the provider.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Network Configuration: Ensure that a connection is established between the Guest VM and the external keystore server by:

- Setting up the required network.
- Opening the necessary ports.
- Enabling the protocol specified by the external keystore vendor.

PKCS#11 Library Installation: Install the PKCS#11-related software and configure the PKCS#11 library on the VMs according to the external keystore vendor's documentation.

Limitations:

- Only one vendor PKCS#11 library can be present on a Guest VM at a time.
- External keystore interfaces cannot be used to associate keys with Oracle Databases on Oracle Exadata Database Service on Cloud@Customer.
- While the external keystore interface allows you to view the keys associated with the databases, it may not support performing key management operations directly from the interface.

Communication Validation: Verify that the PKCS#11 library can successfully communicate with the external keystore. Note that cloud automation does not perform prechecks to validate this connection. If the key is inaccessible, the database will return an error with the relevant details.

For more information, see https://support.oracle.com/support/?kmExternalId=FAQ2403.

Storing the Keys in an External Keystore

You can now encrypt Oracle Databases on Oracle Exadata Database Service on Cloud@Customer by storing the keys in an external keystore.

Applicable Database Versions: 26ai and 19c

When provisioning a database, you have the option to choose from different key management solutions: Oracle Software Keystore, Oracle Key Vault (OKV), or an External Keystore.

- The selected key management solution applies to the entire Container Database (CDB) and all Pluggable Databases (PDBs) contained within it. If a CDB is configured to use an external keystore, all associated PDBs will also use the external keystore. You cannot select different key management solutions at the PDB level.
- While the key management solution must be consistent across the CDB and its PDBs, different PDBs within the same CDB can use distinct encryption keys, providing flexibility in key usage across the PDBs.

This functionality ensures that sensitive encryption keys are securely stored in an external keystore, offering an added layer of security for your databases.

Restrictions on Adding a Virtual Machine to a VM Cluster Configured with an External Keystore

When a database is protected by an external keystore, adding a new virtual machine (VM) to the cluster is restricted.

If one or more databases on a VM cluster are configured with an external keystore, the following message is displayed:

Copyright © 2008, 2026, Oracle and/or its affiliates.

```
"While you should be able to add the virtual machine to the existing VM cluster, the
database instance will not be extended to the newly created VM. This is because one or
more databases on this VM cluster are configured with an external keystore. You must
configure the external keystore on the newly created VM, then run the dbaascli command
to extend the database instances to the new VM."
```

Managing Encryption Keys on External Devices

Learn how to store and manage database encryption keys.

There are two options to store and manage database encryption keys for your databases on Oracle Exadata Database Service on Cloud@Customer:

1. In an auto-login wallet file stored in an Oracle Advanced Cluster File System (Oracle ACFS) accessible by the customer VM operating system.
2. Oracle Key Vault.

- Customer-Managed Keys in Oracle Exadata Database Service on Cloud@Customer
  Customer-managed keys for Oracle Exadata Database Service on Cloud@Customer is a feature that enables you to migrate the Oracle Database TDE Master Encryption Key for an Oracle Database from the password-protected wallet file stored on the Oracle Exadata Database Service on Cloud@Customer equipment to an OKV server that you control.
- About Oracle Key Vault
  Oracle Key Vault is a full-stack, security-hardened software appliance built to centralize the management of keys and security objects within the enterprise.
- Overview of Key Store
  Integrate your on-premises Oracle Key Vault (OKV) with customer-managed database cloud services to secure your critical data on-premises.
- Required IAM Policy for Managing OKV on Oracle Exadata Database Service on Cloud@Customer
  Review the identity access management (IAM) policy for managing OKV on Oracle Exadata Database Service on Cloud@Customer Systems.
- Tagging Resources
  You can apply tags to your resources to help you organize them according to your business needs.
- Moving Resources to a Different Compartment
  You can move OKV Vault, Secret, and Keystore resources from one compartment to another.
- Setting Up Your Oracle Exadata Database Service on Cloud@Customer to Work With Oracle Key Vault
- Managing Your Key Store
- Administer Transparent Data Encryption (TDE) Keys
  Use this procedure to change the key management configuration.
- How to Manually Clone a Pluggable Database (PDB) from a Remote Container Database (CDB) When Data is Encrypted with Master Encryption Key (MEK) in Oracle Key Vault (OKV)
- How to Upgrade Oracle Key Vault (OKV) Home in ExaDB-C@C

Copyright © 2008, 2026, Oracle and/or its affiliates.

Customer-Managed Keys in Oracle Exadata Database Service on Cloud@Customer

Customer-managed keys for Oracle Exadata Database Service on Cloud@Customer is a feature that enables you to migrate the Oracle Database TDE Master Encryption Key for an Oracle Database from the password-protected wallet file stored on the Oracle Exadata Database Service on Cloud@Customer equipment to an OKV server that you control.

The Oracle Key Vault (OKV) provides fault-tolerant, highly available and scalable key and secrets management for your encrypted ExaDB-C@C databases. Use customer-managed keys when you need security governance, regulatory compliance, and homogenous encryption of data, while centrally managing, storing, and monitoring the life cycle of the keys you use to protect your data.

You can:

- Switch from Oracle-managed keys to customer-managed keys on databases that are not enabled with Oracle Data Guard.
- Rotate your keys to maintain security compliance.
- Rotating the PDB key is also supported. Rotate CDB and PDB key operations are allowed only if the database is customer-managed.

Requirements

- To enable the management of customer-managed encryption keys, you must create a policy in the tenancy that allows a particular dynamic group to do so. For more information, see Setting Up Your Oracle Exadata Database Service on Cloud@Customer to Work With Oracle Key Vault.
- Pluggable databases must be configured in United Mode. For more information about United Mode, see Managing Keystores and TDE Master Encryption Keys in United Mode. Isolated Mode is not supported. For more information about Isolated Mode, see Managing Keystores and TDE Master Encryption Keys in Isolated Mode
- If an Exadata Database Service was configured for Oracle Key Vault using the procedures published at Migration of File based TDE to OKV for Exadata Database Service on Cloud at Customer Gen2 (Doc ID 2823650.1), then you should open a My Oracle Support (MOS) Service Request to have Oracle cloud operations update the control plane configuration to reflect the Oracle Key Vault information for the the specific Exadata Database service

About Oracle Key Vault

Oracle Key Vault is a full-stack, security-hardened software appliance built to centralize the management of keys and security objects within the enterprise.

> **Note**
>
> The Oracle Key Vault is a customer-provisioned and managed system and it is not part of Oracle Cloud Infrastructure managed services.

Related Topics

- Oracle Key Vault

Copyright © 2008, 2026, Oracle and/or its affiliates.

Overview of Key Store

Integrate your on-premises Oracle Key Vault (OKV) with customer-managed database cloud services to secure your critical data on-premises.

Oracle Key Vault integration enables you to take complete control of your encryption keys and store them securely on an external, centralized key management device.

OKV is optimized for Oracle wallets, Java keystores, and Oracle Advanced Security Transparent Data Encryption (TDE) master keys. Oracle Key Vault supports the OASIS KMIP standard. The full-stack, security-hardened software appliance uses Oracle Linux and Oracle Database technology for security, availability, and scalability, and can be deployed on your choice of compatible hardware.

OKV also provides a REST interface for clients to auto-enroll endpoints and setup wallets and keys. For Oracle Autonomous AI Database on Exadata Cloud@Customer to connect to OKV REST interface, create a key store in your tenancy to store the IP address and administrator credentials of your OKV. Exadata Cloud@Customer temporarily stores the OKV REST user administrator password required to connect to the OKV appliance in a password-protected wallet file so that the software running in the customer VM can connect to the OKV server. Following the migration of the TDE keys to OKV, the cloud automation software will remove the password from the wallet file. Ensure that you create a secret with Oracle's Vault Service, which will store the password required for autonomous databases to connect to OKV for key management.

For more information, see "Oracle Key Vault".

Related Topics

- Oracle Key Vault

Required IAM Policy for Managing OKV on Oracle Exadata Database Service on Cloud@Customer

Review the identity access management (IAM) policy for managing OKV on Oracle Exadata Database Service on Cloud@Customer Systems.

A **policy** is an IAM document that specifies who has what type of access to your resources. It is used in different ways: to mean an individual statement written in the policy language; to mean a collection of statements in a single, named "policy" document (which has an Oracle Cloud ID (OCID) assigned to it), and to mean the overall body of policies your organization uses to control access to resources.

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

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

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

Related Topics

- Let database admins manage DB systems
- Getting Started with Policies
- Common Policies
- Details for the Database Service

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs.

You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see "Resource Tags".

Related Topics

- Resource Tags

Moving Resources to a Different Compartment

You can move OKV Vault, Secret, and Keystore resources from one compartment to another.

After you move an OCI resource to a new compartment, inherent policies apply immediately and affect access to the resource. Moving an OKV Vault resource doesn't affect access to any OKV Vault Keys or OKV Vault Secrets that the OKV Vault contains. You can move an OKV Vault Keys or OKV Vault Secrets from one compartment to another independently of moving the OKV Vault it's associated with. For more information, see *Managing Compartments*.

Related Topics

- Managing Compartments

Setting Up Your Oracle Exadata Database Service on Cloud@Customer to Work With Oracle Key Vault

Prerequisites

1. Ensure that OKV is set up and the network is accessible from the Exadata client network. Open ports 443, 5695, and 5696 for egress on the client network for the OKV client software and Oracle database instance to access the OKV server.
2. Ensure that the REST interface is enabled from the OKV user interface.
3. Create "OKV REST Administrator" user.
   You can use any qualified username of your choice, for example, `okv_rest_user`. The databases can be key-managed in the same or in different on-premises OKV clusters. ExaDB-C@C requires a REST user with the `Create Endpoint` privilege.
4. Gather OKV administrator credentials and IP address, which is required to connect to OKV.

For more information, see *Network Port Requirements*, *Managing Oracle Key Vault Users*, and *Managing Administrative Roles and User Privileges*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Step 1: Create a Vault in OCI Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password
- Step 2: Create a Policy Statement for Group to Access Secret in OCI Vault Service
- Step 3: Create Key Store

Related Topics

- Network Port Requirements
- Managing Oracle Key Vault Users
- Managing Administrative Roles and User Privileges

Step 1: Create a Vault in OCI Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password

Your Exadata Cloud@Customer infrastructure communicates with OKV over REST each time an Oracle Database is provisioned to register the Oracle Database and request a wallet on OKV. Therefore, Exadata infrastructure needs access to the REST admin credentials to register with the OKV server.

These credentials are stored securely in the Oracle Vault Service in OCI as a Secret and accessed by your Exadata Cloud@Customer infrastructure only when needed. When needed, the credentials are stored in a password-protected wallet file.

To store the OKV administrator password in the OCI Vault service, create a vault by following the instructions outlined in *Managing Vaults* and create a Secret in that vault by following the instructions outlined in *Managing Secrets*.

Related Topics

- Managing Vaults
- Managing Secrets

Step 2: Create a Policy Statement for Group to Access Secret in OCI Vault Service

To enable a user group to access the Secret in the OCI Vault that stores the OKV Administrator password, define an IAM policy that grants the group permission to read secret bundles within the specified compartment.

Policy Statement:

```
Allow group <group> to read secret-bundles in compartment <compartmentName>
where target.secret.id='<secret ocid>'
```

For more information, see:

- Getting a Secret's Contents
- Details for Vault, Key Management, and Secrets

Step 3: Create Key Store

Follow these steps to create a Key Store to connect to an on-premises encryption key appliance such as Oracle Key Vault (OKV).

Copyright © 2008, 2026, Oracle and/or its affiliates.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click **Key Stores**.

   Key Stores page displays the list of name of key stores, the number of databases associated with each database, and the date on which each key store was created.
4. Click **Create Key Store**.
5. In the **Create Key Store** dialog, enter the following general information:
   - **Name your key store:** A user-friendly description or other information that helps you easily identify the Key Store resource. Avoid entering confidential information.
   - **Oracle Key Vault connection settings**
     - **Connection IP addresses:** Enter at least one OKV cluster node IP address; multiple comma-separated IP addresses (of the same OKV cluster) are possible, for example, 193.10.20.1, 193.10.20.2.
     - **Administrator username:** Enter the user name of the *okv_rest_user*.
     - **Administrator Password Secret:** The administrator password is stored with the secret management service within OCI. Select the OCI Vault Secret in your tenancy that contains *okv_rest_user* password stored as Secret.
   - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator. Avoid entering confidential information.
6. Click **Create**.
7. Ensure that you use the same Key Store resource while provisioning a database or while migrating TDE encryption method for a database to OKV.

   For more information, see *Managing Vaults*, *Managing Keys*, and *Managing Secrets*.

Related Topics

- Managing Vaults
- Managing Keys
- Managing Secrets

Managing Your Key Store

- View Key Store Details
  Follow these steps to view Key Store details that include Oracle Key Vault (OKV) connection details and the list of associated databases.
- Edit Key Store Details
  You can edit a Key Store only if it is not associated with any CDBs.
- Move a Key Store to Another Compartment
  Follow these steps to move a Key Store on an Oracle Exadata Database Service on Cloud@Customer system from one compartment to another compartment.
- Delete a Key Store
  You can delete a Key Store only if it is not associated with any CDBs.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- View Key Store Associated Container Database Details
  Follow these steps to view details of the container database associated with a Key Store.
- Using the API to Manage Key Store
  Learn how to use the API to manage key store.

View Key Store Details

Follow these steps to view Key Store details that include Oracle Key Vault (OKV) connection details and the list of associated databases.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click **Key Stores**.

   Key Stores page displays the list name of Key Stores, the number of databases associated with each database, and the date on which each Key Store was created.
4. Click the name of the Key Store or click the Actions icon (three dots), and then click **View Details**.
5. Click the link in the **Administrator Password Secret** field to view secret details.

   The **Associated Databases** section displays the list of CDBs associated with this Key Store.

Edit Key Store Details

You can edit a Key Store only if it is not associated with any CDBs.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click **Key Stores**.
4. Click the name of the Key Store or click the Actions icon (three dots), and then click **View Details**.
5. On the Key Store Details page, click **Edit**.
6. On the Edit Key Store page, make changes as needed, and then click **Save Changes**.

Move a Key Store to Another Compartment

Follow these steps to move a Key Store on an Oracle Exadata Database Service on Cloud@Customer system from one compartment to another compartment.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click **Key Stores**.
4. Click the name of the Key Store, click **Actions**, and then select **Move resource**.

   (or)

   Click the Actions icon (three dots) for the key store you want to move, and then click **Move Resource**.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. On the Move Resource to a Different Compartment page, select the new compartment.
6. Click **Move Resource**.

Delete a Key Store

You can delete a Key Store only if it is not associated with any CDBs.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click **Key Stores**.
4. Click the name of the Key Store, click **Actions**, and then select **Delete**.

   (or)

   Click the Actions icon (three dots) for the key store you want to delete, and then click **Delete**.
5. On the Delete Key Store dialog, click **Delete**.

View Key Store Associated Container Database Details

Follow these steps to view details of the container database associated with a Key Store.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Database Service on Cloud@Customer**.
2. Choose your **Compartment**.
3. Click **Key Stores**.
4. In the resulting Key Stores page, click the name of the Key Store or click the Actions icon (three dots), and then click **View Details**.
5. Click **Associated Databases**.
6. Click the name of the associated database or click the Actions icon (three dots), and then click **View Details**.

Using the API to Manage Key Store

Learn how to use the API to manage key store.

For information about using the API and signing requests, see "REST APIs" and "Security Credentials". For information about SDKs, see "Software Development Kits and Command Line Interface".

The following table lists the REST API endpoints to manage key store.

Operation	REST API Endpoint
Create OKV Key Store	`CreateKeyStore`
View OKV Key Store	`GetKeyStore`
Update OKV Key Store	`UpdateKeyStore`
Delete OKV Key Store	`DeleteKeyStore`
Change Key store compartment	`ChangeKeyStoreCompartment`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Operation	REST API Endpoint
Choose between customer-managed and Oracle-managed encryption	`CreateDatabase`
Get the Key Store (OKV or Oracle-managed) and OKV wallet name	`GetDatabase`
Change Key store type	`changeKeyStoreType`
Rotate OKV and Oracle-managed key	`RotateVaultKey`

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface

Administer Transparent Data Encryption (TDE) Keys

Use this procedure to change the key management configuration.

After provisioning a database on an ExaDB-C@C system, you can change the key management and perform operations such as rotating the TDE keys.

> **Note**
>
> - You can change the key management from Oracle Wallet to other available options.
> - When you change the key management to OKV, the database will experience a shutdown abort operation followed by a restart. Plan to perform the migration in a planned maintenance window.
> - You must rotate TDE keys only via OCI interfaces (Console, API).
> - You cannot rotate an encryption key:
>   - when a database restore is in progress in a given Oracle Home.
>   - when a database patching or database home patching is in progress.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Exadata Database Service on Cloud@Customer**.
2. Choose your compartment from the **Compartment** drop-down.
3. Navigate to the VM Cluster that contains the database for which you want to change encryption management or rotate a key.
   a. Under Exadata Database Service on Cloud@Customer Infrastructure, click **Exadata VM Clusters**.
   b. In the list of VM clusters, locate the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. In the Databases section, click the name of the database for which you want to change encryption management or rotate a key to display its details page.
5. Go to the Encryption section on the database details page.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- **To change key management:**
  If you have configured **Oracle Wallet** as the key management, then the system displays the **Change** option to change the key management to **Oracle Key Vault**.
  a. Click **Change**.
  b. Select the key management as **Oracle Key Vault**.
  You must have a valid encryption key in the Oracle Key Vault service and provide the information in the subsequent steps. For more information, see Key and Secret Management Concepts.
  c. Choose a region.
  d. Choose a compartment.
  You can change the compartment by clicking the **Change Compartment** link.
  e. Click **Save Changes**.
- **To rotate an encryption key:**
  If you have configured **Oracle Key Vault** as the key management, then the system displays the **Rotate** option to rotate the encryption key.
  a. Click **Rotate**.
  b. On the resulting Rotate key dialog, click **Rotate**.

> **Note**
>
> - Migration of TDE keys to Oracle Key Vault (OKV) requires 10 minutes of downtime. During the migration, the database state will be UPDATING and connections may fail due to multiple database restarts to enable OKV. Applications can resume operation after the migration completes and when the database returns to its original ACTIVE state.
> - The OKV keystore password will be set to the TDE wallet password.

> **Caution**
>
> After changing key management, deleting the key from the OKV will cause the database to become unavailable.
>
> On the database details page for this database, the Encryption section displays the encryption key name and the encryption key OCID.

Related Topics

- Key and Secret Management

How to Manually Clone a Pluggable Database (PDB) from a Remote Container Database (CDB) When Data is Encrypted with Master Encryption Key (MEK) in Oracle Key Vault (OKV)

Copyright © 2008, 2026, Oracle and/or its affiliates.

The dbaascli tool lets you clone PDBs when the source CDB and target CDB are the same (local clone) or if they are different (remote clone). However, you cannot clone a remote PDB if the data is encrypted with a MEK in OKV.

> **Note**
>
> To decrypt / encrypt the data during a remote clone, the container database must have access to MEK. The MEK must be made available to the target CDB when it is stored in the OKV server.

- Source CDB and Target CDB are Encrypted with MEK in the Same OKV Server
- Source CDB and Target CDB are Encrypted with MEK in a Different OKV Server

Related Topics

- 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 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.

Source CDB and Target CDB are Encrypted with MEK in the Same OKV Server

1. Get the OKV object ID of the source PDB.
   a. Get the latest encryption key of the source PDB using SQL*Plus.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<source_db_name>.env
[oracle@testserver oracle]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Mon Jun 12 23:13:12 2023
Version 19.19.0.0.0

Copyright (c) 1982, 2022, Oracle.  All rights reserved.

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

SQL> set heading off;
SQL> alter session set container=<SOURCE_PDB>;

Session altered.

SQL> select key_id,keystore_type,activation_time from v$encryption_keys
order by activation_time;

0648E5D8D5559B4F0EBFB8AA5EE730401A
SOFTWARE KEYSTORE
25-MAR-23 12.01.41.075932 AM +00:00

06AFF5B6E27A954F6EBFFC77296B27C9EC
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
SOFTWARE KEYSTORE
25-MAR-23 11.42.51.336955 AM +00:00

SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver oracle]$
```

b. Get the OKV object ID (uuid) of the newest MEK obtained from the step above.

Enter the OKV Endpoint password when prompted and hit the **Enter** key on your keyboard.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<source_db_name>.env
[oracle@testserver oracle]$ $OKV_HOME/bin/okvutil list | grep
06AFF5B6E27A954F6EBFFC77296B27C9EC
E5344379-8B16-4FE9-BF35-F8ECB057571A    Symmetric Key    TDE Master
Encryption Key: MKID 06AFF5B6E27A954F6EBFFC77296B27C9EC
[oracle@testserver oracle]$
```

2. Install OKV REST wallet in the source database.

a. Create the `okv_rest_cli` directory if it does not exist.

```
[root@testserver newdb1]# su oracle
[oracle@testserver oracle]$ mkdir /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli
```

b. Download and extract `okvrestclipackage.zip`.

Select `ALL` if prompted for replacement.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli
[oracle@scaqar06dv0101 okv_rest_cli]$ curl -O -k https://
<source_okv_server_ip1>:5695/okvrestclipackage.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time
Current
                                 Dload  Upload   Total   Spent    Left
Speed
100 3784k  100 3784k    0     0  19.0M      0 --:--:-- --:--:--
--:--:-- 19.1M
[oracle@testserver okv_rest_cli]$ unzip -q okvrestclipackage.zip
[oracle@testserver okv_rest_cli]$
```

c. Modify the `okvrestcli.ini` and `okvrestcli_logging.properties` files as follows.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
server=<source_okv_server_ip1>
user=<source_okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
client_wallet
log_property=/var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
conf/okvrestcli_logging.properties
okv_client_config=/u02/app/oracle/admin/<source_db_name>/okv_home/conf/
okvclient.ora

[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvForma
tter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.OkvFo
rmatter
[oracle@testserver okv_rest_cli]$
```

d. Create the `client_wallet` directory.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv_rest_cli]$
```

e. Create OKV REST wallet using the OKV REST command-line interface.

Enter the source OKV REST password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ export JAVA_HOME=/usr/java/latest;
export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv admin client-wallet add --client-
wallet /var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
client_wallet --wallet-user <source_okv_rest_user>
Password:
{
  "result" : "Success"
}
[oracle@testserver okv_rest_cli]$ ls -ltr /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet
total 8
-rw------- 1 oracle oinstall    0 Jun 16 01:29 ewallet.p12.lck
-rw------- 1 oracle oinstall    0 Jun 16 01:29 cwallet.sso.lck
-rw------- 1 oracle oinstall  976 Jun 16 01:29 ewallet.p12
-rw------- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv_rest_cli]$
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

3. Create a new OKV wallet to store only the PDB MEK obtained in step #1.

   a. Get the OKV wallet name from the source PDB in the format `EXA_DB_NAME_DBID_PDB_NAME_WL`.

      For example, the wallet name would be `EXA_NEWDB1_37508325141_PDB_NAME_WL`.

```
[root@testserver newdb1]# su oracle
[oracle@testserver newdb1]$ source ~/<source_db_name>.env
[oracle@testserver newdb1]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
Version 19.19.0.0.0

Copyright (c) 1982, 2022, Oracle.  All rights reserved.

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

SQL> select name,db_unique_name,dbid from v$database;

NAME      DB_UNIQUE_NAME                 DBID
--------- ------------------------------ ----------
NEWDB1      newdb1_uniq             3750832514

SQL> select value from v$parameter where name='instance_name';

VALUE
--------------------------------------------------------------------------------
--------
newdb11

SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdb1]$
```

   b. Create a new wallet using the OKV REST command-line interface.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet create --
wallet <SOURCE_PDB_OKV_WALLET> --description "Wallet to clone
<source_pdb_name> pdb from <source_db_name>" --unique FALSE
{
  "result" : "Success",
  "value" : {
    "status" : "PENDING",
    "locatorID" : "BA5FBFE1-DB41-4425-8EE4-D58541A1E41A"
  }
}
[root@testserver oracle]#
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

c. Check the status until it is ACTIVE.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet check-status
--wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success",
  "value" : {
    "status" : "PENDING"
  }
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet check-status
--wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success",
  "value" : {
    "status" : "ACTIVE",
    "wallet" : "<SOURCE_PDB_OKV_WALLET>"
  }
}
[root@testserver oracle]#
```

4. Add **Read and Modify**, and **Manage Wallet** permissions from the source database OKV Endpoints to the OKV wallet created in step #3.

   a. Get the Endpoint names from the source database. One per VM.

   Usually, the structure is in the format, `EXA_`*`DB_UNIQUE_NAME_DBID_SID`*`_EP`.

   For example, the Endpoint name of node 1 would be `EXA_`*`NEWDB1_UNIQ_3750832514_NEWDB11`*`_EP`.

```
[root@testserver newdb1]# su oracle
[oracle@testserver newdb1]$ source ~/<source_db_name>.env
[oracle@testserver newdb1]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
Version 19.19.0.0.0

Copyright (c) 1982, 2022, Oracle.  All rights reserved.

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

SQL> select name,db_unique_name,dbid from v$database;

NAME      DB_UNIQUE_NAME                 DBID
--------- ------------------------------ ----------
NEWDB1      newdb1_uniq             3750832514

SQL> select value from v$parameter where name='instance_name';
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
VALUE
--------------------------------------------------------------------------------
--------
newdb11

SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdb1]$
```

b. Add **Read and Modify**, and **Manage Wallet** permissions using the OKV REST command-line interface.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <SOURCE_OKV_EP1> --access
RM_MW
{
  "result" : "Success"
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <SOURCE_OKV_EP2> --access
RM_MW
{
  "result" : "Success"
}
[root@testserver oracle]#
```

5. Store MEK from the source PDB obtained in step #1 into the OKV wallet created in step #3.

   - Add MEK (uuid obtained in step #1.b) using the OKV REST command-line interface.

     Enter the source OKV Endpoint password when prompted.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv managed-object wallet add-member
--uuid E5344379-8B16-4FE9-BF35-F8ECB057571A --wallet
<SOURCE_PDB_OKV_WALLET>
Password:
{
  "result" : "Success"
}
[root@testserver oracle]#
```

6. Install OKV REST wallet in the target database.

Copyright © 2008, 2026, Oracle and/or its affiliates.

a. Create the `okv_rest_cli` directory if it does not exist.

```
[root@testserver newdb1]# su oracle
[oracle@testserver oracle]$ mkdir /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli
```

b. Download and extract `okvrestclipackage.zip`.

Select `ALL` when prompted for replacement.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli
[oracle@scaqar06dv0101 okv_rest_cli]$ curl -O -k https://
<target_okv_server_ip1>:5695/okvrestclipackage.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time
Current
                                 Dload  Upload   Total   Spent    Left
Speed
100 3784k  100 3784k    0     0  19.0M      0 --:--:-- --:--:--
--:--:-- 19.1M
[oracle@testserver okv_rest_cli]$ unzip -q okvrestclipackage.zip
[oracle@testserver okv_rest_cli]$
```

c. Modify the `okvrestcli.ini` and `okvrestcli_logging.properties` files as follows.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli.ini
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
server=<target_okv_server_ip1>
user=<target_okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
client_wallet
log_property=/var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
conf/okvrestcli_logging.properties
okv_client_config=/u02/app/oracle/admin/<target_db_name>/okv_home/conf/
okvclient.ora

[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvForma
tter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.OkvFo
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
rmatter
[oracle@testserver okv_rest_cli]$
```

d. Create the `client_wallet` directory.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv_rest_cli]$
```

e. Create OKV REST wallet using the OKV REST command-line interface.

Enter the target OKV REST password when prompted.

```
[oracle@testserver okv_rest_cli]$ export JAVA_HOME=/usr/java/latest;
export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/bin/okv admin client-wallet add --client-
wallet /var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
client_wallet --wallet-user <target_okv_rest_user>
Password:
{
  "result" : "Success"
}
[oracle@testserver okv_rest_cli]$ ls -ltr /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet
total 8
-rw------- 1 oracle oinstall    0 Jun 16 01:29 ewallet.p12.lck
-rw------- 1 oracle oinstall    0 Jun 16 01:29 cwallet.sso.lck
-rw------- 1 oracle oinstall  976 Jun 16 01:29 ewallet.p12
-rw------- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv_rest_cli]$
```

7. Add **Read Only** and **Manage Wallet** permissions from the target database OKV Endpoints to the source PDB OKV wallet created in step #3.

a. Get the Endpoint names from the target database. One per VM.

Usually, the structure is in the format, `EXA_DB_UNIQUE_NAME_DBID_SID_EP`.

For example, the Endpoint name of node 1 would be `EXA_NEWDB1_UNIQ_3750832514_NEWDB11_EP`.

```
[root@testserver newdb1]# su oracle
[oracle@testserver newdb1]$ source ~/<target_db_name>.env
[oracle@testserver newdb1]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
Version 19.19.0.0.0

Copyright (c) 1982, 2022, Oracle.  All rights reserved.

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

SQL> select name,db_unique_name,dbid from v$database;
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
NAME      DB_UNIQUE_NAME                DBID
--------- ------------------------------ ----------
NEWDB1      newdb1_uniq             3750832514

SQL> select value from v$parameter where name='instance_name';

VALUE
--------------------------------------------------------------------------------
--------
newdb11

SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdb1]$
```

b. Add permissions **Read Only** and **Manage Wallet** using the OKV REST command-line interface.

```
[root@testserver oracle]#  export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <TARGET_OKV_EP1> --access
RO_MW
{
  "result" : "Success"
}
[root@testserver oracle]#  export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <TARGET_OKV_EP2> --access
RO_MW
{
  "result" : "Success"
}
[root@testserver oracle]#
```

8. Clone the PDB.

   - Run `dbaascli` to clone the PDB.

     Enter the source DB SYS user password when prompted.

     ```
     [root@testserver oracle]# dbaascli pdb remoteClone --pdbName
     <source_pdb_name> --dbName <target_db_name> --sourceDBConnectionString
     <source_db_connection_string> --targetPDBName <target_pdb_name>
     DBAAS CLI version 23.2.1.0.0
     Executing command pdb remoteClone --pdbName <source_pdb_name> --dbName
     <target_pdb_name> --sourceDBConnectionString scaqar06dvclu01-
     scan1.us.oracle.com:1521/<source_db_unique_name>.us.oracle.com --
     targetPDBName <target_pdb_name>
     Job id: 197f30e9-209e-4ec5-9700-a13f7915f8b9
     Session log: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
     dbaastools_2023-06-12_10-32-17-PM_188384.log
     ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
Enter REMOTE_DB_SYS_PASSWORD:

Enter REMOTE_DB_SYS_PASSWORD (reconfirmation):

Loading PILOT...
Session ID of the current execution is: 6848
Log file location: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
pilot_2023-06-12_10-32-35-PM_204184
-----------------
Running Plugin_initialization job
Enter REMOTE_DB_SYS_PASSWORD
***************
Completed Plugin_initialization job
-----------------
Running Validate_input_params job
Completed Validate_input_params job
-----------------
Running Perform_dbca_prechecks job
Completed Perform_dbca_prechecks job
-----------------
Running PDB_creation job
Completed PDB_creation job
-----------------
Running Load_pdb_details job
Completed Load_pdb_details job
-----------------
Running Configure_pdb_service job
Completed Configure_pdb_service job
-----------------
Running Configure_tnsnames_ora job
Completed Configure_tnsnames_ora job
-----------------
Running Set_pdb_admin_user_profile job
Completed Set_pdb_admin_user_profile job
-----------------
Running Lock_pdb_admin_user job
Completed Lock_pdb_admin_user job
-----------------
Running Register_ocids job
Skipping. Job is detected as not applicable.
-----------------
Running Prepare_blob_for_standby_in_primary job
Skipping. Job is detected as not applicable.
-----------------
Running Generate_dbsystem_details job
Completed Generate_dbsystem_details job
dbaascli execution completed
[root@testserver oracle]#
```

9. Delete the source PDB OKV wallet created in step #3 using the OKV REST command-line interface.

```
[root@testserver oracle]#  export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet delete --
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
wallet  <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success"
}
[root@testserver oracle]#
```

10. Delete the OKV REST wallet created in step #2.
    - Delete the wallet files in the `dbaas_acfs` directory.

```
[root@testserver oracle]# rm -f /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet/*
[root@testserver oracle]#
```

11. Delete OKV REST wallet created in step #6.
    - Delete the wallet files in the `dbaas_acfs` directory.

```
[root@testserver oracle]# rm -f /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet/*
[root@testserver oracle]#
```

Source CDB and Target CDB are Encrypted with MEK in a Different OKV Server

1. Get the OKV object ID of the source PDB.
   a. Get the latest encryption key of the source PDB using SQL*Plus.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<source_db_name>.env
[oracle@testserver oracle]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Mon Jun 12 23:13:12 2023
Version 19.19.0.0.0

Copyright (c) 1982, 2022, Oracle.  All rights reserved.

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

SQL> set heading off;
SQL> alter session set container=<SOURCE_PDB>;

Session altered.

SQL> select key_id,keystore_type,activation_time from v$encryption_keys
order by activation_time;

0648E5D8D5559B4F0EBFB8AA5EE730401A
SOFTWARE KEYSTORE
25-MAR-23 12.01.41.075932 AM +00:00

06AFF5B6E27A954F6EBFFC77296B27C9EC
SOFTWARE KEYSTORE
25-MAR-23 11.42.51.336955 AM +00:00
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver oracle]$
```

b. Get the OKV object ID (uuid) of the newest MEK obtained from the step above.

Enter the OKV Endpoint password when prompted and hit the **Enter** key on your keyboard.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<source_db_name>.env
[oracle@testserver oracle]$ $OKV_HOME/bin/okvutil list | grep
06AFF5B6E27A954F6EBFFC77296B27C9EC
E5344379-8B16-4FE9-BF35-F8ECB057571A    Symmetric Key    TDE Master
Encryption Key: MKID 06AFF5B6E27A954F6EBFFC77296B27C9EC
[oracle@testserver oracle]$
```

2. Install OKV REST wallet in the source database.

a. Create the `okv_rest_cli` directory if it does not exist.

```
[root@testserver newdb1]# su oracle
[oracle@testserver oracle]$ mkdir /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli
```

b. Download and extract `okvrestclipackage.zip`.

Select `ALL` if prompted for replacement.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli
[oracle@scaqar06dv0101 okv_rest_cli]$ curl -O -k https://
<source_okv_server_ip1>:5695/okvrestclipackage.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time
Current
                                 Dload  Upload   Total   Spent    Left
Speed
100 3784k  100 3784k    0     0  19.0M      0 --:--:-- --:--:--
--:--:-- 19.1M
[oracle@testserver okv_rest_cli]$ unzip -q okvrestclipackage.zip
[oracle@testserver okv_rest_cli]$
```

c. Modify the `okvrestcli.ini` and `okvrestcli_logging.properties` files as follows.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
server=<source_okv_server_ip1>
user=<source_okv_rest_user>
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
client_wallet=/var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
client_wallet
log_property=/var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
conf/okvrestcli_logging.properties
okv_client_config=/u02/app/oracle/admin/<source_db_name>/okv_home/conf/
okvclient.ora


[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvForma
tter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.OkvFo
rmatter
[oracle@testserver okv_rest_cli]$
```

d. Create the `client_wallet` directory.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv_rest_cli]$
```

e. Create OKV REST wallet using OKV REST command-line interface.

Enter the source OKV REST password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ export JAVA_HOME=/usr/java/latest;
export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv admin client-wallet add --client-
wallet /var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
client_wallet --wallet-user <source_okv_rest_user>
Password:
{
  "result" : "Success"
}
[oracle@testserver okv_rest_cli]$ ls -ltr /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet
total 8
-rw------- 1 oracle oinstall    0 Jun 16 01:29 ewallet.p12.lck
-rw------- 1 oracle oinstall    0 Jun 16 01:29 cwallet.sso.lck
-rw------- 1 oracle oinstall  976 Jun 16 01:29 ewallet.p12
-rw------- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv_rest_cli]$
```

3. Create a new OKV wallet to store only the PDB MEK obtained in step #1.

Copyright © 2008, 2026, Oracle and/or its affiliates.

a. Create a new wallet using the OKV REST command-line interface. Name suggested: EXA_*DB_NAME_DBID_PDB_NAME*_WL

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet create --
wallet <SOURCE_PDB_OKV_WALLET> --description "Wallet to clone
<source_pdb_name> pdb from <source_db_name>" --unique FALSE
{
  "result" : "Success",
  "value" : {
    "status" : "PENDING",
    "locatorID" : "BA5FBFE1-DB41-4425-8EE4-D58541A1E41A"
  }
}
[root@testserver oracle]#
```

b. Check the status until it is ACTIVE.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet check-status
--wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success",
  "value" : {
    "status" : "PENDING"
  }
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet check-status
--wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success",
  "value" : {
    "status" : "ACTIVE",
    "wallet" : "<SOURCE_PDB_OKV_WALLET>"
  }
}
[root@testserver oracle]#
```

4. Add **Read and Modify**, and **Manage Wallet** permissions from the source database OKV Endpoints to the OKV wallet created in step #3.

a. Get the Endpoint names from the source database. One per VM.

Usually, the structure is in the format, EXA_*DB_UNIQUE_NAME_DBID_SID*_EP.

For example, the Endpoint name of node 1 would be EXA_*NEWDB1_UNIQ_3750832514_NEWDB11*_EP.

```
[root@testserver newdb1]# su oracle
[oracle@testserver newdb1]$ source ~/<source_db_name>.env
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
[oracle@testserver newdb1]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
Version 19.19.0.0.0

Copyright (c) 1982, 2022, Oracle.  All rights reserved.

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

SQL> select name,db_unique_name,dbid from v$database;

NAME      DB_UNIQUE_NAME                 DBID
--------- ------------------------------ ----------
NEWDB1      newdb1_uniq              3750832514

SQL> select value from v$parameter where name='instance_name';

VALUE
--------------------------------------------------------------------------------
--------
newdb11

SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdb1]$
```

b. Add permissions **Read and Modify**, and **Manage Wallet** using the OKV REST command-line interface.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <SOURCE_OKV_EP1> --access
RM_MW
{
  "result" : "Success"
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <SOURCE_OKV_EP2> --access
RM_MW
{
  "result" : "Success"
}
[root@testserver oracle]#
```

5. Store MEK from the source PDB obtained in step #1 into OKV wallet created in step #3.

   - Add MEK (uuid obtained in step #1.b) using the OKV REST command-line interface.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Enter the source OKV Endpoint password when prompted.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv managed-object wallet add-member
--uuid E5344379-8B16-4FE9-BF35-F8ECB057571A --wallet
<SOURCE_PDB_OKV_WALLET>
Password:
{
  "result" : "Success"
}
[root@testserver oracle]#
```

6. Download OKV wallet created in step #3 from the OKV server to the local filesystem.
   a. Create a new directory with permissions for `oracle` user.

      This directory will store the wallet that will contain only the MEK of the source PDB.

      ```
      [root@testserver oracle]# su oracle
      [oracle@testserver oracle]$ mkdir /home/oracle/<source_pdb_wallet_dir>
      [oracle@testserver oracle]$
      ```

   b. Download the OKV wallet created in step #3 to the directory created in step #6.a using `okvutil`.

      It will prompt twice for a password to encrypt the local wallet. Use the same password as the source Endpoint password. Also, enter the source Endpoint password when prompted.

      ```
      [root@testserver oracle]# su oracle
      [oracle@testserver oracle]$ source ~/nfsa.env
      [oracle@testserver oracle]$ $OKV_HOME/bin/okvutil download -l /home/
      oracle/<source_pdb_wallet_dir> -t wallet -g <SOURCE_PDB_OKV_WALLET>
      Enter new wallet password (<enter> for auto-login):
      Confirm new wallet password:
      Enter Oracle Key Vault endpoint password:
      Download succeeded
      [oracle@testserver oracle]$
      ```

   c. Zip the wallet directory.

      ```
      [root@testserver oracle]# su oracle
      [oracle@testserver oracle]$ cd /home/oracle
      [oracle@testserver oracle]$ zip -r <source_pdb_wallet_dir>.zip
      <source_pdb_wallet_dir>
        adding: <source_pdb_wallet_dir>/ (stored 0%)
        adding: <source_pdb_wallet_dir>/ewallet.p12 (stored 0%)
      [oracle@testserver oracle]$
      ```

7. Delete the source PDB OKV wallet created in step #3.

   ```
   [root@testserver oracle]#  export JAVA_HOME=/usr/java/latest; export
   OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
   okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
   <source_db_name>/okv_rest_cli/bin/okv manage-access wallet delete --
   ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
wallet  <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success"
}
[root@testserver oracle]#
```

8. Delete the OKV REST wallet created in step #1.

   - Delete the wallet files in the `dbaas_acfs` directory.

     ```
     [root@testserver oracle]# rm -f /var/opt/oracle/dbaas_acfs/
     <source_db_name>/okv_rest_cli/client_wallet/*
     [root@testserver oracle]#
     ```

9. Copy the source PDB wallet downloaded to the focal filesystem in step #6 to the target Cluster VM.
10. Delete the source PDB wallet from the source local filesystem created in step #6.

    a. Delete the wallet directory.

       ```
       [root@testserver oracle]# su oracle
       [oracle@testserver oracle]$ rm -rf /home/oracle/<source_pdb_wallet_dir>
       [oracle@testserver oracle]$
       ```

    b. Delete the wallet zip file.

       ```
       [root@testserver oracle]# su oracle
       [oracle@testserver oracle]$ rm -f /home/oracle/
       <source_pdb_wallet_dir>.zip
       [oracle@testserver oracle]$
       ```

11. Install OKV REST wallet in the target database.

    a. Create the `okv_rest_cli` directory if it does not exist.

       ```
       [root@testserver newdb1]# su oracle
       [oracle@testserver oracle]$ mkdir /var/opt/oracle/dbaas_acfs/
       <target_db_name>/okv_rest_cli
       ```

    b. Download and extract `okvrestclipackage.zip`.

       Select `ALL` when prompted for replacement.

       ```
       [root@testserver oracle]# su oracle
       [oracle@testserver oracle]$ cd /var/opt/oracle/dbaas_acfs/
       <target_db_name>/okv_rest_cli
       [oracle@testserver okv_rest_cli]$ curl -O -k https://
       <target_okv_server_ip>:5695/okvrestclipackage.zip
         % Total    % Received % Xferd  Average Speed   Time    Time     Time
       Current
                                       Dload  Upload   Total   Spent    Left
       Speed
       100 3784k  100 3784k    0     0  19.0M      0 --:--:-- --:--:--
       --:--:-- 19.1M
       [oracle@testserver okv_rest_cli]$ unzip -q okvrestclipackage.zip
       [oracle@testserver okv_rest_cli]$
       ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

c. Modify the `okvrestcli.ini` and `okvrestcli_logging.properties` files as follows.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli.ini
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
server=<target_okv_server_ip1>
user=<target_okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
client_wallet
log_property=/var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
conf/okvrestcli_logging.properties
okv_client_config=/u02/app/oracle/admin/<target_db_name>/okv_home/conf/
okvclient.ora

[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvForma
tter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.OkvFo
rmatter
[oracle@testserver okv_rest_cli]$
```

d. Create the `client_wallet` directory.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv_rest_cli]$
```

e. Create OKV REST wallet using the OKV REST command-line interface.

Enter the target OKV REST password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ export JAVA_HOME=/usr/java/latest;
export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/bin/okv admin client-wallet add --client-
wallet /var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
client_wallet --wallet-user <target_okv_rest_user>
/var/opt/oracle/dbaas_acfs/newdb1/okv_rest_cli/logs/okvrest.log.lck
Password:
{
  "result" : "Success"
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
}
[oracle@testserver okv_rest_cli]$ ls -ltr /var/opt/oracle/dbaas_acfs/
newdb1/okv_rest_cli/client_wallet
total 8
-rw------- 1 oracle oinstall    0 Jun 16 01:29 ewallet.p12.lck
-rw------- 1 oracle oinstall    0 Jun 16 01:29 cwallet.sso.lck
-rw------- 1 oracle oinstall  976 Jun 16 01:29 ewallet.p12
-rw------- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv_rest_cli]$
```

12. Upload the source PDB wallet created in step #6 and copied to the target Cluster VM in step #9.

    a. Unzip the source PDB wallet.

    ```
    [root@testserver oracle]# su oracle
    [oracle@testserver oracle]$ cd /home/oracle/
    [oracle@testserver ~]$ unzip <source_pdb_wallet_dir>.zip
    Archive:  nfsa_1672104454_NFSPDB_wallet.zip
       creating: <source_pdb_wallet_dir>/
     extracting: <source_pdb_wallet_dir>/ewallet.p12
    [oracle@testserver ~]$
    ```

    b. Get OKV wallet name from the target database in the format `EXA_DB_NAME_DBID_WL`.

    For example, the wallet name would be `EXA_NEWDB1_37508325141_WL`.

    ```
    [root@testserver newdb1]# su oracle
    [oracle@testserver newdb1]$ source ~/<target_db_name>.env
    [oracle@testserver newdb1]$ sqlplus / as sysdba

    SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
    Version 19.19.0.0.0

    Copyright (c) 1982, 2022, Oracle.  All rights reserved.

    Connected to:
    Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
    Version 19.19.0.0.0

    SQL> select name,db_unique_name,dbid from v$database;

    NAME      DB_UNIQUE_NAME               DBID
    --------- ------------------------------ ----------
    NEWDB1      newdb1_uniq             3750832514

    SQL> select value from v$parameter where name='instance_name';

    VALUE
    ------------------------------------------------------------------------
    --------
    newdb11

    SQL> exit
    Disconnected from Oracle Database 19c EE Extreme Perf Release
    19.0.0.0.0 - Production
    ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
Version 19.19.0.0.0
[oracle@testserver newdb1]$
```

c. Upload the source PDB wallet to the target OKV wallet using `okvutil`.

Enter the source PDB wallet password when prompted. Use the same password as the source Endpoint password.

Also, enter the target Endpoint password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<target_db_name>.env
[oracle@testserver oracle]$ $OKV_HOME/bin/okvutil upload -t WALLET -l /
home/oracle/<source_pdb_wallet_dir> -g <TARGET_OKV_WALLET>
Enter source wallet password:
Enter Oracle Key Vault endpoint password:
WARNING: Object ORACLE.SECURITY.ID.ENCRYPTION. already exists; use -o
to overwrite
Upload succeeded
[oracle@testserver oracle]$
```

13. Clone the PDB.

- Run `dbaascli` to clone the PDB.

```
[root@testserver oracle]# dbaascli pdb remoteClone --pdbName
<source_pdb_name> --dbName <target_db_name> --sourceDBConnectionString
<source_db_connection_string> --targetPDBName <target_pdb_name>
DBAAS CLI version 23.2.1.0.0
Executing command pdb remoteClone --pdbName <source_pdb_name> --dbName
<target_db_name> --sourceDBConnectionString scaqar06dvclu01-
scan1.us.oracle.com:1521/<source_db_unique_name>.us.oracle.com --
targetPDBName <target_pdb_name>
Job id: 7d4f638a-1f3a-4219-a05a-0215588dcae8
Session log: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
dbaastools_2023-06-13_01-29-09-AM_179996.log
Enter REMOTE_DB_SYS_PASSWORD:

Enter REMOTE_DB_SYS_PASSWORD (reconfirmation):

Loading PILOT...
Session ID of the current execution is: 6857
Log file location: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
pilot_2023-06-13_01-29-21-AM_196991
-----------------
Running Plugin_initialization job
Enter REMOTE_DB_SYS_PASSWORD
*************
Completed Plugin_initialization job
-----------------
Running Validate_input_params job
Completed Validate_input_params job
-----------------
Running Perform_dbca_prechecks job
Completed Perform_dbca_prechecks job
-----------------
Running PDB_creation job
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
Completed PDB_creation job
-----------------
Running Load_pdb_details job
Completed Load_pdb_details job
-----------------
Running Configure_pdb_service job
Completed Configure_pdb_service job
-----------------
Running Configure_tnsnames_ora job
Completed Configure_tnsnames_ora job
-----------------
Running Set_pdb_admin_user_profile job
Completed Set_pdb_admin_user_profile job
-----------------
Running Lock_pdb_admin_user job
Completed Lock_pdb_admin_user job
-----------------
Running Register_ocids job
Skipping. Job is detected as not applicable.
-----------------
Running Prepare_blob_for_standby_in_primary job
Skipping. Job is detected as not applicable.
-----------------
Running Generate_dbsystem_details job
Completed Generate_dbsystem_details job
dbaascli execution completed
[root@testserver oracle]#
```

14. Delete the OKV REST wallet created in step #1.

    - Delete the wallet files in the `dbaas_acfs` directory.

      ```
      [root@testserver oracle]# rm -f /var/opt/oracle/dbaas_acfs/
      <target_db_name>/okv_rest_cli/client_wallet/*
      [root@testserver oracle]#
      ```

How to Upgrade Oracle Key Vault (OKV) Home in ExaDB-C@C

After the encryption type is migrated from Oracle Managed Keys to Customer Managed Keys (Oracle Key Vault), the OKV home in the Guest VMs remains with the same version used for the migration.

In case the OKV Server is upgraded the functionality would keep working because of backward compatibility. However, the customer might want to get the new features for the client tools. In that case, upgrade the OKV home and the `PKCS#11` Library.

1. Validate current OKV Home version is minor to OKV Server version.

   a. Get the OKV Home version by running `okvutil`. In this case the value is 21.6.0.0.0.

      ```
      # su oracle
      $ /u02/app/oracle/admin/<dbname>/okv_home/okv/bin/okvutil
      okvutil version 21.6.0.0.0
      Usage: okvutil <command> [-v <verbosity>] [<command args>]
        <command> := list | upload | download | sign | sign-verify |
      changepwd | diagnostics
      ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
Options:
  -v, --verbose <verbosity>
    Print extra information to standard out.
    Possible verbosity values are 0, 1 and 2 (more detailed information
with higher verbosity level).
For help on a particular command, use [okvutil <command> -h].
You have new mail in /var/spool/mail/root
```

b. Get the OKV Server version by logging in to the OKV Server console through the browser. In this case the Version is 21.7.0.0.0.

2. Install OKV REST wallet in the source database. This step must be done in only one node.

a. If not existing, create the `okv_rest_cli` directory.

```
# su oracle
$ mkdir /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli
```

b. Download and extract `okvrestclipackage.zip`. If prompted for replacement, select `ALL`.

```
$ cd /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli
$ curl -O -k https://100.75.59.249:5695/okvrestclipackage.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time
Current
                                 Dload  Upload   Total   Spent    Left
Speed
100 3865k  100 3865k    0     0  5102k      0 --:--:-- --:--:--
--:--:-- 5106k
$ unzip -q okvrestclipackage.zip
```

c. Modify `okvrestcli.ini` with the next info.

```
$ vi /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/
okvrestcli.ini

$ cat !$
cat /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/okvrestcli.ini
#Provide absolute path for log_property, okv_client_config properties
[Default]
#log_property=./conf/okvrestcli_logging.properties
#server=[OKV IP Address]
#okv_client_config=./conf/okvclient.ora
#user=[OKV username]
#password=[user password]

#[Profile1]
#server=
#okv_client_config=
#user=

#[Profile2]
#server=
#okv_client_config=
#user=

server=<okv_server_ip>
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
user=<okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/
client_wallet
```

d. Create the `client_wallet` directory.

```
$ mkdir /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/client_wallet
```

e. Create OKV REST Wallet using OKV REST CLI. It will prompt for the source OKV REST password.

```
$ export JAVA_HOME=/usr/java/latest; export OKV_RESTCLI_CONFIG=/var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/okvrestcli.ini; /var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/bin/okv admin client-wallet add
--client-wallet /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/
client_wallet --wallet-user <okv_rest_user>
Password:
{
  "result" : "Success"
}
```

3. Prepare the OKV Home directories in Guest VM 1.

a. Rename OKV Home directory as the current OKV Home version.

```
$ mv /u02/app/oracle/admin/<dbname>/okv_home/okv /u02/app/oracle/admin/
<dbname>/okv_home/okv<current_okv_home_version>
```

b. Create a new OKV Home directory as the OKV Server version.

```
$ mkdir /u02/app/oracle/admin/<dbname>/okv_home/okv<okv_server_version>
```

c. Create a symlink of regular OKV Home name to the directory created in step 3.b.

```
$ ln -s /u02/app/oracle/admin/<dbname>/okv_home/
okv<okv_server_version> /u02/app/oracle/admin/<dbname>/okv_home/okv
```

4. Upgrade OKV Home in Guest VM 1.

a. Get OKV Endpoint name via `okvutil`. It will prompt for OKV Endpoint Password (TDE Password) The entry is the one called "Template".

Note that the hostname should be the same as the current Guest VM Hostname. (Usually, this is named as `EXA_<DBNAME>_<resourceID>_<CURRENT_Guest VM_HOST_NAME>_EP`. *<resourceID>* can be get by listing the DB via `dbaascli` system `getDatabases`).

```
$ /u02/app/oracle/admin/<dbname>/okv_home/okv/bin/okvutil list
Enter Oracle Key Vault endpoint password:
Unique ID                                Type            Identifier
DC690343-5694-4FC8-BFE4-6C7F1A550F67    Opaque Object   TDE Wallet
Metadata
9E317DDB-0542-553B-A47D-FCC31AB6DD7C    Symmetric Key   TDE Master
Encryption Key: MKID AaTAGyAWyk/fv7pnl8qx4s0AAAAAAAAAAA
D9D840AF-A60E-5850-AA86-8C9F216F5501    Symmetric Key   TDE Master
Encryption Key: MKID AUP0Tq+un08Mv1+onNhT4RUAAAAAAAAAAA
364EFC2F-1909-4F34-BF1B-90D3D03DA7EB    Private Key Private Key
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
A9D0134F-C895-4F33-BF85-351B754E9FF9    Opaque Object   TDE Wallet
Metadata
E1AC8D2F-90E9-4F88-BFEE-2883FCBB7271    Opaque Object   TDE Wallet
Metadata
25B7DE14-3849-4F67-BFBE-1934BFE3559B    Opaque Object   TDE Wallet
Metadata
4ED713ED-FE2B-4F35-BF7D-BCBEA8327A0B    Symmetric Key   TDE Master
Encryption Key: MKID 06EA813441C26B4F53BFD58E55C4BE90F4
6162E200-EF0A-4F89-BF25-A8596B3AD7B0    Opaque Object   Certificate
Request
85A55486-28E5-4FFB-BF1C-B93C4C0BAD74    Secret Data Oracle Secret Data:
ID HSM_PASSWORD
67E74D97-56F6-407A-A035-009D953F907A    Template    Default template
for EXA_DB1902_7274B2A2-6F71-4516-B2BB-6D67CC3824FC_SCAQAE08DV0308_EP
E621EA72-5DD1-4F4F-BFD4-451E5B7DB8A9    Symmetric Key   TDE Master
Encryption Key: MKID 0625BA455B03CD4F57BFA5D2290FD379A1
```

b. Re-enroll the Endpoint in Guest VM 1.

```
$ export JAVA_HOME=/usr/java/latest; export OKV_RESTCLI_CONFIG=/var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/okvrestcli.ini; /var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/bin/okv admin endpoint re-
enroll --endpoint <endpoint_name>
{
  "result" : "Success"
}
```

c. Provision Endpoint in Guest VM 1. It will prompt for OKV Endpoint password (TDE password).

```
$ export JAVA_HOME=/usr/java/latest; export OKV_RESTCLI_CONFIG=/var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/okvrestcli.ini; /var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/bin/okv admin endpoint
provision --endpoint <endpoint_name> --location /u02/app/oracle/admin/
<dbname>/okv_home/okv --auto-login FALSE
Enter Oracle Key Vault endpoint password:
{
  "result" : "Success"
}
```

5. Validate the OKV Home Upgrade was successfully done.

a. Validate the OKV Endpoint lists the entries in OKV Wallet. It will prompt for OKV Endpoint password (TDE password).

```
$ /u02/app/oracle/admin/db1902/okv_home/okv/bin/okvutil list
Enter Oracle Key Vault endpoint password:
Unique ID                               Type            Identifier
DC690343-5694-4FC8-BFE4-6C7F1A550F67    Opaque Object   TDE Wallet
Metadata
9E317DDB-0542-553B-A47D-FCC31AB6DD7C    Symmetric Key   TDE Master
Encryption Key: MKID AaTAGyAWyk/fv7pnl8qx4s0AAAAAAAAAAA
D9D840AF-A60E-5850-AA86-8C9F216F5501    Symmetric Key   TDE Master
Encryption Key: MKID AUP0Tq+un08Mv1+onNhT4RUAAAAAAAAAAA
364EFC2F-1909-4F34-BF1B-90D3D03DA7EB    Private Key Private Key
A9D0134F-C895-4F33-BF85-351B754E9FF9    Opaque Object   TDE Wallet
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
Metadata
E1AC8D2F-90E9-4F88-BFEE-2883FCBB7271    Opaque Object   TDE Wallet
Metadata
25B7DE14-3849-4F67-BFBE-1934BFE3559B    Opaque Object   TDE Wallet
Metadata
4ED713ED-FE2B-4F35-BF7D-BCBEA8327A0B    Symmetric Key   TDE Master
Encryption Key: MKID 06EA813441C26B4F53BFD58E55C4BE90F4
6162E200-EF0A-4F89-BF25-A8596B3AD7B0    Opaque Object   Certificate
Request
85A55486-28E5-4FFB-BF1C-B93C4C0BAD74    Secret Data Oracle Secret Data:
ID HSM_PASSWORD
67E74D97-56F6-407A-A035-009D953F907A    Template    Default template
for EXA_DB1902_7274B2A2-6F71-4516-B2BB-6D67CC3824FC_SCAQAE08DV0308_EP
E621EA72-5DD1-4F4F-BFD4-451E5B7DB8A9    Symmetric Key   TDE Master
Encryption Key: MKID 0625BA455B03CD4F57BFA5D2290FD379A1
You have new mail in /var/spool/mail/root
```

b. Get the OKV Home version by running `okvutil`. The version should be the same as OKV Server version. In this case the value should be 21.7.0.0.0.

```
# su oracle

$ /u02/app/oracle/admin/<dbname>/okv_home/okv/bin/okvutil
okvutil version 21.7.0.0.0
Usage: okvutil <command> [-v <verbosity>] [<command args>]
  <command> := list | upload | download | sign | sign-verify |
changepwd | diagnostics
Options:
  -v, --verbose <verbosity>
    Print extra information to standard out.
    Possible verbosity values are 0, 1 and 2 (more detailed information
with higher verbosity level).
For help on a particular command, use [okvutil <command> -h].
You have new mail in /var/spool/mail/root
```

6. Repeat the Steps 3 - 5 in the rest of the Guest VMs.
7. Repeat the Steps 1 - 6 for any other DB that needs to upgrade its OKV Home.
8. Stop the Guest VM 1 instances of all Databases with OKV-based TDE. This can be done via the Console, `srvctl` command or SQL* Plus.
9. Run `root.sh` from the selected OKV Home. It usually should be the one with the newer OKV version. It will prompt for replace `PKCS11` Library, `YES` should be selected.

```
# /u02/app/oracle/admin/<dbname>/okv_home/okv/bin/root.sh
```

10. Start the Guest VM 1 instances of all databases with OKV-based TDE. This can be done via the Console, `srvctl` command or SQL* Plus.
11. Repeat the Steps 8 - 10 in the rest of the Guest VMs.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Migrate to Oracle Exadata Database Service on Cloud@Customer

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

A recommended approach for migrating to Oracle Exadata Database Service on Cloud@Customer is using Zero Downtime Migration

- Moving to Oracle Cloud Using Zero Downtime Migration

Related Topics

- Migrating Databases to the Cloud

Moving to Oracle Cloud Using Zero Downtime Migration

Oracle now offers the Zero Downtime Migration service, a quick and easy way to move on-premises databases to Oracle Cloud Infrastructure.

Zero Downtime Migration leverages Oracle Active Data Guard to create a standby instance of your database in an Oracle Cloud Infrastructure system. You switch over only when you are ready, and your source database remains available as a standby. Use the Zero Downtime Migration service to migrate databases individually or at the fleet level. See *Move to Oracle Cloud Using Zero Downtime Migration* for more information.

Related Topics

- Move to Oracle Cloud Using Zero Downtime Migration

Overview of Exadata Cloud@Customer Gen1 to Out-of-Place Cloud Upgrade to Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure

Gen1 is the first generation of Exadata Database Service on Cloud@Customer, which is deployed in conjunction with Gen1 Oracle Cloud At Customer (OCC) as Control Plane deployed in the customer data center. Oracle Exadata Database Service on Cloud@Customer Gen2 is managed from Oracle Cloud Infrastructure (OCI) Control Plane, which runs in OCI public cloud.

Out-of-Place Cloud Upgrade to Exadata Database Service on Cloud@Customer Gen2 Infrastructure: If you are running Exadata Cloud@Customer Gen1 on X6 or X7 infrastructure, with this offering Oracle will replace Gen1 X6 or X7 infrastructure with new Gen2 Exadata Cloud@Customer Infrastructure, and provide instructions to use Oracle Zero Downtime Migration (ZDM) to migrate your databases on the Exadata Cloud@Customer Gen1 platform to Exadata Cloud@Customer Gen2 platform. Replacing your Exadata Cloud@Customer Gen1 X6 or X7 infrastructure and migrating your databases to the Exadata Cloud@Customer Gen2 platform is called an out-of-place cloud upgrade.

- Scope for Exadata Cloud@Customer Gen1 to Gen2 Out-of-Place Cloud Upgrade

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Hardware and Software Required for Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure
  Review this checklist to prepare for the out-of-place cloud upgrade to new Gen2 infrastructure:
- Using Oracle Zero Downtime Migration (ZDM) to Migrate Oracle Databases
  Use ZDM to migrate Oracle databases from Exadata Cloud@Customer Gen1 to Oracle Exadata Database Service on Cloud@Customer Gen2 infrastructure.
- During Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure
- Post Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure
  The upgrade will move your resources into Oracle Exadata Database Service on Cloud@Customer Gen 2 Cloud Control Plane and onto the new generation hardware.
- Best Practices for Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure
  For the purpose of the upgrade, the recommended tool to use is Oracle Zero Downtime Migration (ZDM).

Scope for Exadata Cloud@Customer Gen1 to Gen2 Out-of-Place Cloud Upgrade

- Exadata Cloud@Customer X6 and X7 System Shapes are eligible for the out-of-place upgrade.
- Databases on Exadata Cloud@Customer Gen1 which participate in a Data Guard configuration are supported by migration. In this case, the primary should be migrated to Exadata Cloud@Customer Gen2 using the regular procedure. Once migration is done, the Data Guard configuration should be set up on the Exadata Cloud@Customer Gen2 side using the regular Gen2 procedure.
- Upgrade from Exadata Cloud@Customer Gen1 to Gen2 is done only when the software versions are compatible on the source and target systems.
  - Oracle Database software: The source and target must be at the same major version. For example, both the source and target must be at version 19c. However, the target can have a higher patch level than the source. For example, the patch versions can be 19.3 for the source and 19.8 for the target. The corresponding equivalence for 12.2 is, both source and the target must be at Oracle Database software version 12.2.0.1, but the patch levels can be `2019JulyRU` on the source and `20200ctRU` on the target.
  - Non-Container Database (CDB) to non-CDB.
  - Multitenant deployment (CDB/PDB) to Multitenant deployment (CDB/PDB).
  - Single-instance Oracle Database: After migration, single-instance database source will be converted to Oracle Real Application Clusters (Oracle RAC) database on the target.
- Permitted differences in software versions include:
  - Oracle Grid Infrastructure
  - Exadata software
  - Guest VM operating system
  - DBaaS tools

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Oracle Databases created on Exadata Cloud@Customer Gen1 using backend tools, dbaasapi and dbaascli, are supported as well besides Oracle Databases created using the Gen1 Console.
- All supported versions of Oracle Database on Exadata Cloud@Customer Gen1 are supported and will be migrated to the same major version on the target. The Gen2 environment will be on the latest supported version of Exadata Cloud@Customer Gen2 for the Guest VM operating system and Oracle Grid Infrastructure.

Note that the following are not in scope for the out-of-place cloud upgrade to new Gen2 hardware:

- Exadata Cloud@Customer Gen1 deployment using Exadata Cloud@Customer Gen1 features not yet available on Gen2 are not expected to use the upgrade procedures until the relevant feature or equivalent is available on Gen2.
- Only Exadata Cloud@Customer upgrade is in scope as part of the procedure here. Upgrade or migration of OCC itself is not in the scope.
- You can reverse the upgrade until both the primary and secondary hardware is at your site. Loss of data is possible depending on the application usage and cutover time. Once the Exadata Cloud@Customer Gen1 hardware is shipped back to Oracle, you cannot reverse the upgrade and you cannot revert to Exadata Cloud@Customer Gen1 as well.

Hardware and Software Required for Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure

Review this checklist to prepare for the out-of-place cloud upgrade to new Gen2 infrastructure:

- Setup Exadata Cloud@Customer Gen2 Environment

  A functioning base Exadata Cloud@Customer Gen2 deployment is a pre-requisite for starting any Exadata Cloud@Customer Gen1 to Gen2 out-of-place upgrade.

  For more information to setup your Gen2 Exadata Cloud@Customer, see *Preparing for Exadata Cloud@Customer*.
- Setup Hardware to Migrate Oracle Databases using Oracle Zero Downtime Migration (ZDM). For more information, see *Prepare a Host for Zero Downtime Migration Software Installation*
- Configure network
  - Provide network access path from the Exadata Cloud@Customer Gen1 servers and the Gen2 servers to the ZDM servers used for the upgrade.
  - Provide network access and SSH access from the ZDM server to the respective Exadata Cloud@Customer infrastructure.
  - For any client access to the target databases, ensure that a network path is available from the client host to the new Exadata Cloud@Customer Gen2 deployed databases.
- Software
  - The upgrade will require minimum versions of the software stack so prior to the upgrade, install the appropriate version of Oracle Grid Infrastructure on the target Exadata Cloud@Customer Gen2 infrastructure.
  - Oracle Database versions supported on Exadata Cloud@Customer Gen1 will continue to be supported. On the target Gen2 infrastructure, install appropriate versions of the Oracle Database software and the one-off patches that exist in the source database.

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - Complete all requirements for ZDM servers in terms of installation, configuration, network access, and SSH access.
- Security
  - Exadata Cloud@Customer Gen2 does not use Oracle Advanced Support Gateway Security (OASG) so cannot request OASG logs.
- Ensure that automatic backup is not configured on the Gen2 target prior to migration.

Related Topics

- Preparing for Oracle Exadata Database Service on Cloud@Customer
  Review OCI as well as the site, network and storage requirements to prepare and deploy Oracle Exadata Database Service on Cloud@Customer in your data center.
- Prepare a Host for Zero Downtime Migration Software Installation
- Best Practices for Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure
  For the purpose of the upgrade, the recommended tool to use is Oracle Zero Downtime Migration (ZDM).

Using Oracle Zero Downtime Migration (ZDM) to Migrate Oracle Databases

Use ZDM to migrate Oracle databases from Exadata Cloud@Customer Gen1 to Oracle Exadata Database Service on Cloud@Customer Gen2 infrastructure.

To familiarize yourself with the features of ZDM, see *Setting Up Zero Downtime Migration Software*. As the first step, download, install and configure ZDM on the host identified for the ZDM server.

Zero Downtime Migration supports both online and offline (backup and recovery) migration. For Exadata Cloud@Customer upgrade from Gen1 to Gen2, it is recommended to use ZDM Physical Migration. Specifically, it is recommended to use **Online Migration with Direct Data Transfer** (online physical migration (`MIGRATION_METHOD=ONLINE_PHYSICAL`) using direct data transfer (`DATA_TRANSFER_MEDIUM=DIRECT`). **Online Migration with Direct Data Transfer** is available with Zero Downtime Migration 21.2 and supports direct data transfer for the physical migration methodology. This new feature allows users to avoid using an intermediate store for backups (normally NFS or the OCI Object Storage). ZDM leverages either active database duplication (for 11.2 databases) or restore from service (for 12+ databases). You can use this method to do migrate your Exadata Cloud@Customer Gen1 databases to Exadata Cloud@Customer Gen2. Examples of the command-line and response files are provided below for reference.

For more information, see:

- *Introduction to Zero Downtime Migration*
- *Preparing for Database Migration*
- *Migrating Your Database with Zero Downtime Migration*

Example 5-4 Active Duplicate

```
zdmcli migrate database -sourcedb z19tgt1 -sourcenode scaqae03client01vm06 -
srcauth zdmauth -srcarg1 user:opc -srcarg2 identity_file:/home/giusr/.ssh/
id_gen1vm -srcarg3 sudo_location:/usr/bin/sudo -targetnode tgt1 -rsp /home/
giusr/activeduplicate_zdm_online_19c.rsp -tgtauth zdmauth -tgtarg1 user:opc -
tgtarg2 identity_file:/home/giusr/.ssh/dbaas_sshkey.priv -tgtarg3
sudo_location:/usr/bin/sudo -schedule NOW -tdekeystorepasswd
ZDM_GET_SRC_INFO ............... COMPLETED
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
ZDM_GET_TGT_INFO ..... COMPLETED
ZDM_PRECHECKS_SRC ..... COMPLETED
ZDM_PRECHECKS_TGT ..... COMPLETED
ZDM_SETUP_SRC ..... COMPLETED
ZDM_SETUP_TGT ..... COMPLETED
ZDM_PREUSERACTIONS ..... COMPLETED
ZDM_PREUSERACTIONS_TGT ..... COMPLETED
ZDM_VALIDATE_SRC ..... COMPLETED
ZDM_VALIDATE_TGT ..... COMPLETED
ZDM_DISCOVER_SRC ..... COMPLETED
ZDM_COPYFILES ..... COMPLETED
ZDM_PREPARE_TGT ..... COMPLETED
ZDM_SETUP_TDE_TGT ..... COMPLETED
ZDM_DUPLICATE_TGT ..... COMPLETED
ZDM_FINALIZE_TGT ..... COMPLETED
ZDM_CONFIGURE_DG_SRC ..... COMPLETED
ZDM_SWITCHOVER_SRC ..... COMPLETED
ZDM_SWITCHOVER_TGT ..... COMPLETED
ZDM_POST_DATABASE_OPEN_TGT .... COMPLETED
ZDM_DATAPATCH_TGT ..... COMPLETED
ZDM_MANIFEST_TO_CLOUD ..... COMPLETED
ZDM_POST_MIGRATE_TGT ..... COMPLETED
ZDM_POSTUSERACTIONS ..... COMPLETED
ZDM_POSTUSERACTIONS_TGT ..... COMPLETED
ZDM_CLEANUP_SRC ..... COMPLETED
ZDM_CLEANUP_TGT ..... COMPLETED
```

Example 5-5 Active Duplicate Response File

```
TGT_DB_UNIQUE_NAME=z19tgt1_uniq2
MIGRATION_METHOD=ONLINE_PHYSICAL
DATA_TRANSFER_MEDIUM=DIRECT
PLATFORM_TYPE=EXACC
SRC_HTTP_PROXY_URL=
SRC_HTTP_PROXY_PORT=
SRC_CONFIG_LOCATION=
SRC_BASTION_HOST_IP=
SRC_BASTION_PORT=
SRC_BASTION_USER=
SRC_BASTION_IDENTITY_FILE=
SRC_HOST_IP=
SRC_TIMEZONE=
SRC_OSS_PROXY_HOST=
SRC_OSS_PROXY_PORT=
SRC_SSH_RETRY_TIMEOUT=
SRC_PDB_NAME=
SRC_DB_LISTENER_PORT=
TGT_HTTP_PROXY_URL=
TGT_HTTP_PROXY_PORT=
TGT_CONFIG_LOCATION=
TGT_BASTION_HOST_IP=
TGT_BASTION_PORT=
TGT_BASTION_USER=
TGT_BASTION_IDENTITY_FILE=
TGT_HOST_IP=
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
TGT_SSH_TUNNEL_PORT=
TGT_SSH_RETRY_TIMEOUT=
TGT_OSS_PROXY_HOST=
TGT_OSS_PROXY_PORT=
TGT_DATADG=
TGT_REDODG=
TGT_RECODG=
TGT_DATAACFS=
TGT_REDOACFS=
TGT_RECOACFS=
BACKUP_PATH=
HOST=
OPC_CONTAINER=
SRC_ZDLRA_WALLET_LOC=
TGT_ZDLRA_WALLET_LOC=
ZDLRA_CRED_ALIAS=
NONCDBTOPDB_CONVERSION=FALSE
NONCDBTOPDB_SWITCHOVER=TRUE
SKIP_FALLBACK=TRUE
TGT_RETAIN_DB_UNIQUE_NAME=
TGT_SKIP_DATAPATCH=FALSE
MAX_DATAPATCH_DURATION_MINS=
DATAPATCH_WITH_ONE_INSTANCE_RUNNING=
SHUTDOWN_SRC=
SKIP_SRC_SERVICE_RETENTION=
SRC_RMAN_CHANNELS=6
TGT_RMAN_CHANNELS=16
ZDM_LOG_OSS_PAR_URL=
ZDM_BACKUP_FULL_SRC_MONITORING_INTERVAL=10
ZDM_BACKUP_INCREMENTAL_SRC_MONITORING_INTERVAL=10
ZDM_BACKUP_DIFFERENTIAL_SRC_MONITORING_INTERVAL=10
ZDM_CLONE_TGT_MONITORING_INTERVAL=10
ZDM_OSS_RESTORE_TGT_MONITORING_INTERVAL=10
ZDM_OSS_RECOVER_TGT_MONITORING_INTERVAL=10
ZDM_BACKUP_RETENTION_WINDOW=
ZDM_BACKUP_TAG=
ZDM_USE_EXISTING_BACKUP=
ZDM_OPC_RETRY_WAIT_TIME=
ZDM_OPC_RETRY_COUNT=
ZDM_SRC_TNS_ADMIN=
ZDM_CURL_LOCATION=
ZDM_USE_EXISTING_UNDO_SIZE=
ZDM_SKIP_DG_CONFIG_CLEANUP=
ZDM_RMAN_COMPRESSION_ALGORITHM=LOW
ZDM_SRC_DB_RESTORE_SERVICE_NAME=
ZDM_RMAN_DIRECT_METHOD=ACTIVE_DUPLICATE
```

Example 5-6 Restore from Service

```
zdmcli migrate database -sourcedb z12tgt1s -sourcenode scaqae03client01vm06 -
srcauth zdmauth -srcarg1 user:opc -srcarg2 identity_file:/home/giusr/.ssh/
id_gen1vm -srcarg3 sudo_location:/usr/bin/sudo -targetnode tgt1 -rsp /home/
giusr/dir_zdm_online_121_sidb.rsp -tgtauth zdmauth -tgtarg1 user:opc -tgtarg2
identity_file:/home/giusr/.ssh/dbaas_sshkey.priv-tgtarg3
sudo_location:/usr/bin/sudo -schedule NOW -tdekeystorepasswd"
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
ZDM_GET_SRC_INFO ..... COMPLETED
ZDM_GET_TGT_INFO ..... COMPLETED
ZDM_PRECHECKS_SRC ..... COMPLETED
ZDM_PRECHECKS_TGT ..... COMPLETED
ZDM_SETUP_SRC ..... COMPLETED
ZDM_SETUP_TGT ..... COMPLETED
ZDM_PREUSERACTIONS ..... COMPLETED
ZDM_PREUSERACTIONS_TGT ..... COMPLETED
ZDM_VALIDATE_SRC ..... COMPLETED
ZDM_VALIDATE_TGT ..... COMPLETED
ZDM_DISCOVER_SRC ..... COMPLETED
ZDM_COPYFILES ..... COMPLETED
ZDM_PREPARE_TGT ..... COMPLETED
ZDM_SETUP_TDE_TGT ..... COMPLETED
ZDM_RESTORE_TGT ..... COMPLETED
ZDM_RECOVER_TGT ..... COMPLETED
ZDM_FINALIZE_TGT ..... COMPLETED
ZDM_CONFIGURE_DG_SRC ..... COMPLETED
ZDM_SWITCHOVER_SRC ..... COMPLETED
ZDM_SWITCHOVER_TGT ..... COMPLETED
ZDM_POST_DATABASE_OPEN_TGT .... COMPLETED
ZDM_DATAPATCH_TGT ..... COMPLETED
ZDM_MANIFEST_TO_CLOUD ..... COMPLETED
ZDM_POST_MIGRATE_TGT ..... COMPLETED
ZDM_POSTUSERACTIONS ..... COMPLETED
ZDM_POSTUSERACTIONS_TGT ..... COMPLETED
ZDM_CLEANUP_SRC ..... COMPLETED
ZDM_CLEANUP_TGT ..... COMPLETED
```

Example 5-7 Restore from Service Response File

```
TGT_DB_UNIQUE_NAME=z12tgt1s_uniq
MIGRATION_METHOD=ONLINE_PHYSICAL
DATA_TRANSFER_MEDIUM=DIRECT
PLATFORM_TYPE=EXACC
SRC_HTTP_PROXY_URL=
SRC_HTTP_PROXY_PORT=
SRC_CONFIG_LOCATION=
SRC_BASTION_HOST_IP=
SRC_BASTION_PORT=
SRC_BASTION_USER=
SRC_BASTION_IDENTITY_FILE=
SRC_HOST_IP=
SRC_TIMEZONE=
SRC_OSS_PROXY_HOST=
SRC_OSS_PROXY_PORT=
SRC_SSH_RETRY_TIMEOUT=
SRC_PDB_NAME=
SRC_DB_LISTENER_PORT=
TGT_HTTP_PROXY_URL=
TGT_HTTP_PROXY_PORT=
TGT_CONFIG_LOCATION=
TGT_BASTION_HOST_IP=
TGT_BASTION_PORT=
TGT_BASTION_USER=
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
TGT_BASTION_IDENTITY_FILE=
TGT_HOST_IP=
TGT_SSH_TUNNEL_PORT=
TGT_SSH_RETRY_TIMEOUT=
TGT_OSS_PROXY_HOST=
TGT_OSS_PROXY_PORT=
TGT_DATADG=
TGT_REDODG=
TGT_RECODG=
TGT_DATAACFS=
TGT_REDOACFS=
TGT_RECOACFS=
BACKUP_PATH=
HOST=
OPC_CONTAINER=
SRC_ZDLRA_WALLET_LOC=
TGT_ZDLRA_WALLET_LOC=
ZDLRA_CRED_ALIAS=
NONCDBTOPDB_CONVERSION=FALSE
NONCDBTOPDB_SWITCHOVER=TRUE
SKIP_FALLBACK=TRUE
TGT_RETAIN_DB_UNIQUE_NAME=
TGT_SKIP_DATAPATCH=FALSE
MAX_DATAPATCH_DURATION_MINS=
DATAPATCH_WITH_ONE_INSTANCE_RUNNING=
SHUTDOWN_SRC=
SKIP_SRC_SERVICE_RETENTION=
SRC_RMAN_CHANNELS=6
TGT_RMAN_CHANNELS=16
ZDM_LOG_OSS_PAR_URL=
ZDM_BACKUP_FULL_SRC_MONITORING_INTERVAL=10
ZDM_BACKUP_INCREMENTAL_SRC_MONITORING_INTERVAL=10
ZDM_BACKUP_DIFFERENTIAL_SRC_MONITORING_INTERVAL=10
ZDM_CLONE_TGT_MONITORING_INTERVAL=10
ZDM_OSS_RESTORE_TGT_MONITORING_INTERVAL=10
ZDM_OSS_RECOVER_TGT_MONITORING_INTERVAL=10
ZDM_BACKUP_RETENTION_WINDOW=
ZDM_BACKUP_TAG=
ZDM_USE_EXISTING_BACKUP=
ZDM_OPC_RETRY_WAIT_TIME=
ZDM_OPC_RETRY_COUNT=
ZDM_SRC_TNS_ADMIN=
ZDM_CURL_LOCATION=
ZDM_USE_EXISTING_UNDO_SIZE=
ZDM_SKIP_DG_CONFIG_CLEANUP=
ZDM_RMAN_COMPRESSION_ALGORITHM=LOW
ZDM_SRC_DB_RESTORE_SERVICE_NAME=
ZDM_RMAN_DIRECT_METHOD=
```

Related Topics

- Introduction to Zero Downtime Migration
- Preparing for Database Migration
- Migrating Your Database with Zero Downtime Migration

Copyright © 2008, 2026, Oracle and/or its affiliates.

During Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure

Monitoring: Oracle will monitor the Oracle Exadata Database Service on Cloud@Customer Gen2 installation from the beginning of the installation just like a regular Gen2 install.

Backups: Backups are done from the Exadata Cloud@Customer Gen1 VM cluster and they will continue to work during the upgrade. Post-migration to Oracle Exadata Database Service on Cloud@Customer Gen2, backups to the Gen1 Oracle Cloud At Customer (OCC) Object Storage Service (OSS) is not allowed and you must use supported backup methods for Oracle Exadata Database Service on Cloud@Customer Gen2.

Post Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure

The upgrade will move your resources into Oracle Exadata Database Service on Cloud@Customer Gen 2 Cloud Control Plane and onto the new generation hardware.

Use the Gen2 OCI Console to manage your Oracle Exadata Database Service on Cloud@Customer Gen2 infrastructure, clusters, databases, and users/groups.

The software stack is upgraded to newer versions, for example, as follows:

- **Exadata software:** 19.x or later
- **Oracle Grid Infrastructure:** 19c
- **Guest VM operating system:** Oracle Linux 7
- **DBaaS Tools:** 20.x
- **CSI:** You will have a new CSI for your Cloud account.

> **Note**
>
> The software stack will be upgraded to the latest versions at that point in time when you perform Out-of-Place Cloud Upgrade to New Gen2 Hardware.

Patching: The infrastructure patching process and notification are different in Gen2. For more information, see *Maintaining an Exadata Cloud@Customer System*.

> **Note**
>
> Post-migration to Exadata Cloud@Customer Gen2, Oracle recommends using supported backup methods for Exadata Cloud@Customer Gen2. It's your responsibility to manually manage any backups to the Gen1 Oracle Cloud At Customer (OCC) Object Storage Service (OSS), and Oracle does not offer it through the OCI Console, API, or CLI.

Related Topics

- Perform User Managed Maintenance Updates

Copyright © 2008, 2026, Oracle and/or its affiliates.

- About Oracle Managed Oracle Exadata Database Service on Cloud@Customer Infrastructure Maintenance Updates
  Oracle performs patches and updates to all of the Oracle-managed system components on Oracle Exadata Database Service on Cloud@Customer.

Best Practices for Out-of-Place Cloud Upgrade to New Oracle Exadata Database Service on Cloud@Customer Gen2 Infrastructure

For the purpose of the upgrade, the recommended tool to use is Oracle Zero Downtime Migration (ZDM).

Some recommended best practices in the context of ZDM's usage for the Gen1 to Gen2 upgrade are:

- While all methods of Physical migration are supported, it is recommended to use **Online Migration with Direct Data Transfer**.

> **Note**
>
> It is not recommended to use either Gen1 Oracle Cloud At Customer OSS or OCI Object Storage for this migration.

- Set `ZDM_RMAN_COMPRESSION_ALGORITHM` to `LOW`.
- The target database Oracle Home must be at the same patch level or at a higher patch level as the source.
- The target database Oracle Home must have all the one-off patches as the source Oracle Home.
- Perform a validation run before the actual run.
- For a CDB source database, it is required to have all the PDBs on the source to be online.

Related Topics

- Zero Downtime Migration Process Phases

Copyright © 2008, 2026, Oracle and/or its affiliates.

6

Oracle Autonomous AI Database on Exadata Cloud@Customer

- What's New in Oracle Autonomous AI Database on Exadata Cloud@Customer
- Introduction to Oracle Autonomous AI Database on Exadata Cloud@Customer

What's New in Oracle Autonomous AI Database on Exadata Cloud@Customer

For complete details on key additions and enhancements for Oracle Autonomous AI Database on Exadata Cloud@Customer, see the New Feature Announcements.

Introduction to Oracle Autonomous AI Database on Exadata Cloud@Customer

Autonomous AI Database offers a fully autonomous, easy-to-use database solution that delivers fast query performance, scales elastically, and eliminates the need for manual database administration.

Oracle Autonomous AI Database on Cloud@Customer provides a highly automated, fully managed database environment running on Oracle Cloud Infrastructure (OCI) with dedicated hardware and software resources. This isolation enables organizations to meet stringent security, availability, and performance requirements—while also reducing operational complexity and cost.

> **Note**
>
> To avoid duplication and deliver a more streamlined, consistent experience, content related to Oracle Autonomous AI Database on Exadata Cloud@Customer has been removed from this guide. Going forward, refer to the Oracle Autonomous Database on Dedicated Exadata Infrastructure documentation, which now serves as the authoritative source for the most current guidance.
> Autonomous AI Database features are largely consistent across Oracle Public Cloud and Exadata Cloud@Customer deployments, with a few differences. Where applicable, these differences are explicitly highlighted using badges.

Copyright © 2008, 2026, Oracle and/or its affiliates.

7

Reference Guides for Oracle Exadata Database Service on Cloud@Customer

- Using the dbaascli Utility with Oracle Exadata Database Service on Cloud@Customer
  Learn to use the `dbaascli` utility on Oracle Exadata Database Service on Cloud@Customer.
- Monitoring and Managing Exadata Storage Servers with ExaCLI
  Learn to use the ExaCLI command-line utility to perform monitoring and management functions on Exadata storage servers in the Exadata Cloud Service.
- 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.
- Metrics for Oracle Exadata Database Service on Cloud@Customer in the Monitoring Service
  Review the metrics emitted by the Oracle Exadata Database Service on Cloud@Customer service in the `oci_database_cluster` and `oci_database` namespaces for Oracle Databases.
- Oracle Exadata Database Service on Cloud@Customer Events
  Oracle Exadata Database Service on Cloud@Customer resources emit events, which are structured messages that indicate changes in resources.
- Policy Details for Oracle Exadata Database Service on Cloud@Customer
  Learn to write policies to control access to Oracle Exadata Database Service on Cloud@Customer resources.
- Managing Exadata Resources with Oracle Enterprise Manager Cloud Control
  To manage and monitor Exadata Cloud and Exadata Cloud@Customer resources, use Oracle Enterprise Manager Cloud Control.
- Observability and Management for Oracle Exadata Database Service on Cloud@Customer
- Security Guide for Oracle Exadata Database Service on Cloud@Customer Systems
  This guide describes security for an Oracle Exadata Database Service on Cloud@Customer system. It includes information about the best practices for securing the Oracle Exadata Database Service on Cloud@Customer system.
- EU Data Act Compliance
  Oracle supports the goal of fair and transparent data sharing and access within the European Union.
- Troubleshooting Oracle Exadata Database Service on Cloud@Customer Systems
  These topics cover some common issues you might run into and how to address them.

Using the dbaascli Utility with Oracle Exadata Database Service on Cloud@Customer

Learn to use the `dbaascli` utility on Oracle Exadata Database Service on Cloud@Customer.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- About Using the dbaascli Utility on Oracle Exadata Database Service on Cloud@Customer
  You can use the `dbaascli` utility to perform various database lifecycle and administration operations on Oracle Exadata Database Service on Cloud@Customer such as creating an Oracle Database, patching an Oracle Database, managing pluggable databases (PDBs), scaling the CPU core count in disconnected mode, and more.
- Creating Oracle Database 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 Cloud@Customer Software Images Using the Dbaascli Utility
  You can list and download the Oracle database software images on an Oracle Exadata Database Service on Cloud@Customer instance, which can then be used for provisioning a database home.
- Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli
  Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud@Customer system.
- Updating Cloud Tooling Using dbaascli
  To update the cloud tooling release for Oracle Exadata Database Service on Cloud@Customer, complete this procedure.
- Creating a Duplicate Database
- Release Notes
  Review the changes made in various releases of `dbaascli`.
- dbaascli Command Reference
  You must use `dbaascli` to create databases and integrate them with the cloud automation framework.

About Using the dbaascli Utility on Oracle Exadata Database Service on Cloud@Customer

You can use the `dbaascli` utility to perform various database lifecycle and administration operations on Oracle Exadata Database Service on Cloud@Customer such as creating an Oracle Database, patching an Oracle Database, managing pluggable databases (PDBs), scaling the CPU core count in disconnected mode, and more.

You must use the DBaaS console or command-line interface to scale resources. The capabilities of the `dbaascli` utility are in addition to, and separate from, the Oracle Cloud Infrastructure Console, API, or command-line interface (CLI). Unless specified differently, you need `root` access to `dbaascli` to run all administration commands.

To use the utility, you must be connected to an Exadata Cloud@Customer virtual machine. See *Connecting to a Virtual Machine with SSH*.

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`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Creating Oracle Database 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
  To get a list of available supported versions for creating Oracle Database, 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

To get a list of available supported versions for creating Oracle Database, 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
   ```

   The command output lists the available database software images.
4. Exit the `root` user command shell:

   ```
   exit
   ```

   For more details on advanced supported options, see `dbaascli cswlib showImages`.

Example 7-1 dbaascli cswlib showImages

```
dbaascli cswlib showImages

DBAAS CLI version MAIN Executing command cswlib showImages
INFO : Log file => /var/opt/oracle/log/list/
list_2021-05-10_10:11:00.56966610630.log

############ List of Available DB Images #############
1.IMAGE_TAG=19.8.0.0.0
VERSION=19.8.0.0.0
DESCRIPTION=19c JUL 2020 DB Image
IMAGE_ALIASES=19000-19800,19000-JUL2020

2.IMAGE_TAG=19.8.0.0.0-NC
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
VERSION=19.8.0.0.0
DESCRIPTION=19c JUL 2020 Non CDB Image
IMAGE_ALIASES=19000-NC19800,19000-NCJUL2020

3.IMAGE_TAG=19.9.0.0.0
VERSION=19.9.0.0.0
DESCRIPTION=19c OCT 2020 DB Image
IMAGE_ALIASES=19000-19900,19000-OCT2020

4.IMAGE_TAG=19.9.0.0.0-NC
VERSION=19.9.0.0.0
DESCRIPTION=19c OCT 2020 Non CDB Image
IMAGE_ALIASES=19000-NC19900,19000-NCOCT2020
```

> **Note**
>
> You can specify the target version in `dbaascli dbhome create` command as `--version` value from the `dbaascli cswlib showImages` command output.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer 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:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--version` specifies the Oracle Database version
- `--imageTag` specifies the Image Tag of the image to be used

For example:

```
dbaascli dbhome create --version 19.9.0.0.0
```

> **Note**
>
> Specifying `imageTag` is optional. To view the Image Tags, refer to command `dbaascli cswlib showImages`. Image Tags are typically same as the version of the database. However, it is kept as a provision for cases where multiple images may need to be released for the same version - each catering to a specific customer requirement.

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli dbhome create`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli dbhome create
  To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.

Creating Oracle Database In the Specified Oracle Database Home

To create an Oracle Database in the specified Oracle Database home of desired version, use the `dbaascli database create` command.

You can use the `dbaascli database create` command to:

- Create a Container Database (CDB) or non-Container Database
- Create a CDB with pluggable databases (PDBs)
- Create an Oracle Database with the specified Character Set
- Create Oracle Databases on a subset of cluster nodes

> **Note**
>
> Databases created on a subset of nodes will not be displayed in the OCI console.

- Create Oracle Database version 12.1.0.2 or higher with the release update JAN 2021 or higher. For databases with lower versions, it is recommended to use the OCI Console based API.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home
   Path
   ```

   Where:

   - `--dbName` specifies the name of the database
   - `--oracleHome` specifies Oracle home location

   To create a CDB, run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home
   Path
   ```

   To create a non-CDB, run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home
   Path --createAsCDB false
   ```

   When prompted, enter the `sys` and `tde` passwords.

4. Exit the `root` user command shell:

   ```
   exit
   ```

   For more details on advanced supported options, see `dbaascli database create`.

- Running Prerequisite Checks Prior to Creating Oracle Database
  To run 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 Cloud@Customer 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

For detailed instructions, see *Connecting to a Virtual Machine with SSH*.

2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli database create --dbName database name --oracleHome Oracle Home Path --executePrereqs
   ```

   Where:

   - `--dbName` specifies the name of the database
   - `--oracleHome` specifies the Oracle home location

4. Exit the `root` user command shell:

   ```
   exit
   ```

   For more details on advanced supported options, see `dbaascli database create`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer 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.

Copyright © 2008, 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 Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli database create
  To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Changing the Database Passwords

To change the SYS password, or to change the TDE wallet password, use this procedure.

The password that you specify in the **Database Admin Password** field when you create a new Exadata Database Service on Cloud@Customer instance or database is set as the password for the SYS, SYSTEM, TDE wallet, and PDB administrator credentials. Use the following procedures if you need to change passwords for an existing database.

> **Note**
>
> If you are enabling Data Guard for a database, then the SYS password and the TDE wallet password of the primary and standby databases must all be the same.

> **Note**
>
> Using the `dbaascli` to change the SYS password will ensure the backup/restore automation can parallelize channels across all nodes in the cluster.

To Change the SYS Password for an Exadata Database Service on Cloud@Customer Database

1. Log onto the Exadata Database Service on Cloud@Customer virtual machine as `opc`.
2. Run the following command:

   ```
   sudo dbaascli database changepassword --dbname database_name --user SYS
   ```

To Change Database Passwords in a Data Guard Environment

1. Run the following command on the primary database:

   ```
   dbaascli database changePassword –dbName <dbname> --user SYS --
   prepareStandbyBlob true --blobLocation <location to create the blob file>
   ```

2. Copy the blob file created to all the standby databases and update the file ownership to `oracle` user.
3. Run the following command on all the standby databases:

   ```
   dbaascli database changePassword –dbName <dbname> --user SYS --
   standbyBlobFromPrimary <location of copies the blob file>
   ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

To Change the TDE Wallet Password for an Exadata Database Service on Cloud@Customer Database

1. Log onto the Exadata Database Service on Cloud@Customer virtual machine as `opc`.
2. Run the following command:

   ```
   sudo dbaascli tde changepassword --dbname database_name
   ```

Managing Oracle Exadata Database Service on Cloud@Customer Software Images Using the Dbaascli Utility

You can list and download the Oracle database software images on an Oracle Exadata Database Service on Cloud@Customer 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 Cloud@Customer instances using the Console or API. These images are stored in Object Storage, and can be used to provision a Database Home in your Exadata instance. See Oracle Database Software Images more information.

You can control the version of Oracle binaries that is installed when you provision a new database on an Oracle Exadata Database Service on Cloud@Customer 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 Cloud@Customer 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
   ```


Copyright © 2008, 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 7-2 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
   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 © 2008, 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 Cloud@Customer 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 Cloud@Customer 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 Cloud@Customer environment.
   The optional parameters are:

   - **version:** specifies an Oracle Database software version. For example, 19.14.0.0.0.
   - **imageTag:** specifies the image tag of the image.
4. Exit the `root` user command shell:

   ```
   # exit
   $
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli

Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud@Customer system.

- Patching Databases using dbaascli
  Using `dbaascli`, you can choose to patch a database by patching Oracle home, or by moving the database to an Oracle home with the desired patch level.
- Patching Oracle Grid Infrastructure
  To apply a patch to Oracle Grid Infrastructure, use the `grid patch` command.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Listing Available Software Images and Versions for Database and Grid Infrastructure
  To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.
- Performing a Precheck Before Patching Databases and Grid Infrastructure
  You can perform a prerequisites-checking operation (also called a "precheck") for the commands in this topic using the applicable precheck flag.
- Resuming or Rolling Back a Patching Operation
  You can resume or revert a failed patching operation. Reverting a patch is known as a rollback.

Patching Databases using dbaascli

Using `dbaascli`, you can choose to patch a database by patching Oracle home, or by moving the database to an Oracle home with the desired patch level.

- Patching an Oracle home (in-place patching). This updates all databases located in the Oracle home.
- Moving a database to a different Oracle home that has the desired Oracle Database software version (out-of-place patching).
- Patching a Database Home (In-Place Database Patching)
  To patch an Oracle home, use the `dbaascli dbHome patch` command.
- Moving a Database to a Different Oracle Home (Out-of-Place Patching)
  To patch an Oracle Database by moving it to an Oracle home that is already at the desired patch level, use the `dbaascli database move` command.

Patching a Database Home (In-Place Database Patching)

To patch an Oracle home, use the `dbaascli dbHome patch` command.

This will patch all databases running in the specified home, and the databases will remain in the home after the patching is complete. The following apply to using the `dbHome patch` command for in-place patching operations:

- You can patch all of your database nodes or a subset of nodes.
- Multi-node patching takes place in a rolling fashion.
- Optionally, you can perform a software-only patch operation. Then, when you are ready, you can run `datapatch` to perform post-patch SQL actions.
- You can patch an Oracle home containing one or more databases.

To patch an Oracle Home (dbhome):

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli dbhome patch --oracleHome dbhome_path --targetVersion
   Oracle_Database_version
   ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Where:

- `--oracleHome` identifies the path of the Oracle home to be patched.
- `--targetVersion` specifies the target Oracle Database version to use for patching, specified as five numeric segments separated by periods (e.g. 19.12.0.0.0).

For example:

```
dbaascli dbhome patch --oracleHome /u02/app/oracle/product/19.0.0.0/
dbhome_2 --targetVersion 19.9.0.0.0
```

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli dbHome patch`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

Moving a Database to a Different Oracle Home (Out-of-Place Patching)

To patch an Oracle Database by moving it to an Oracle home that is already at the desired patch level, use the `dbaascli database move` command.

After the database move operation is complete, the database runs using the Oracle Database software version of the target Oracle Home.

To patch a database by moving it to a different Oracle Home:

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli database move --oracleHome path_to_target_oracle_home --dbname
database_name
```

Where:

- `--oracleHome` identifies the path of the target Oracle home that uses the desired Oracle Database software version. Note that the target Oracle home must exist in your system prior to using the `database move` command.
- `--dbname` specifies the name of the database that is being moved.

Copyright © 2008, 2026, Oracle and/or its affiliates.

For example:

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0.0/
dbhome_2 --dbname xyz
```

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli database move`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.

Patching Oracle Grid Infrastructure

To apply a patch to Oracle Grid Infrastructure, use the `grid patch` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli grid patch --targetVersion target_software_version_number
```

Where `--targetVersion` identifies target software version that the Oracle Grid Infrastructure will be patched to.

For example:

```
dbaascli grid patch --targetVersion 19.11.0.0.0
```

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli grid patch`.

- Patching Oracle Grid Infrastructure (GI) Using GI Software Image
  To patch Oracle Grid Infrastructure (GI) using GI software image, use this procedure.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Patching Oracle Grid Infrastructure (GI) Using GI Software Image

To patch Oracle Grid Infrastructure (GI) using GI software image, use this procedure.

Oracle Grid Infrastructure can also be patched by first creating a patched software image, and then using that image to perform the patching operation. This provides the advantage that an image can be created ahead of time outside of the patching window. It also helps in conflict resolution as any conflicts among the patches are highlighted during the image creation process without impacting the patching window.

1. Create a patched software image.

   ```
   dbaascli grid patch --targetVersion <target_software_version_number> --
   createImage
   ```

   Once the patched software image creation is completed, the image can then be used for performing the patching operation.
2. Perform the patching operation.

   ```
   dbaascli grid patch --targetVersion <target_software_version_number> --
   imageLocation <location_of_patched_software_image>
   ```

Listing Available Software Images and Versions for Database and Grid Infrastructure

To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.

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`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-3 dbaascli cswlib showImages

```
[root@dg11lrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
      showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file location:
      /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
AM_46941.log

############
List of Available Database Images
#############

17.IMAGE_TAG=18.17.0.0.0
   VERSION=18.17.0.0.0
   DESCRIPTION=18c JAN 2022 DB Image

18.IMAGE_TAG=19.10.0.0.0
   VERSION=19.10.0.0.0
   DESCRIPTION=19c JAN 2021 DB Image

19.IMAGE_TAG=19.11.0.0.0
   VERSION=19.11.0.0.0
   DESCRIPTION=19c APR 2021 DB Image

20.IMAGE_TAG=19.12.0.0.0
  VERSION=19.12.0.0.0
  DESCRIPTION=19c JUL 2021 DB Image

21.IMAGE_TAG=19.13.0.0.0
  VERSION=19.13.0.0.0
  DESCRIPTION=19c OCT 2021 DB Image

Images can be downloaded using their image tags. For details, see help using
'dbaascli cswlib download --help'.
dbaascli execution completed
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli cswlib showImages
  To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.

Performing a Precheck Before Patching Databases and Grid Infrastructure

You can perform a prerequisites-checking operation (also called a "precheck") for the commands in this topic using the applicable precheck flag.

Running prechecks allows you to run only the precheck portion of the patching operation without performing actual patching. Oracle recommends running prechecks to discover software issues that could prevent successful patching.

Copyright © 2008, 2026, Oracle and/or its affiliates.

To perform patching prechecks, first, connect to a virtual machine in your Exadata Cloud@Customer instance as the `root` user.

- Precheck for Oracle Home Patching (In-Place Patching)
  Use the `--executePrereqs` flag with the `dbaascli dbhome patch` command.
- Precheck for Database Move Patching (Out-of-Place Patching)
  Use the `--executePrereqs` flag with the `dbaascli database move` command.
- Precheck for Oracle Grid Infrastructure Patching
  Use the `--executePrereqs` flag with the `dbaascli grid patch` command.

Precheck for Oracle Home Patching (In-Place Patching)

Use the `--executePrereqs` flag with the `dbaascli dbhome patch` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli dbhome patch --oracleHome dbhome_path --targetVersion
   Oracle_Database_version --executePrereqs
   ```

   Where:

   - `--oracleHome` identifies the path of the Oracle home to be prechecked.
   - `--targetVersion` specifies the target Oracle Database version to be patched to, specified as five numeric segments separated by periods (e.g. 19.12.0.0.0).
4. Exit the `root` user command shell:

   ```
   exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

Precheck for Database Move Patching (Out-of-Place Patching)

Use the `--executePrereqs` flag with the `dbaascli database move` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

3. Run the following command:

   ```
   dbaascli database move --oracleHome path_to_target_oracle_home --dbname
   database_name --executePrereqs
   ```

   Where:

   - `--oracleHome` identifies the path of the target Oracle Home that uses the desired Oracle Database software version. Note that the target Oracle Home must exist in your system prior to using the `database move` command.
   - `--dbname` specifies the name of the database that is being moved

4. Exit the `root` user command shell:

   ```
   exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.

Precheck for Oracle Grid Infrastructure Patching

Use the `--executePrereqs` flag with the `dbaascli grid patch` command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the following command:

   ```
   dbaascli grid patch --targetVersion target_software_version_number --
   executePrereqs
   ```

   Where `--targetVersion` identifies target software version that the Oracle Grid Infrastructure will be patched to, specified as five numeric segments separated by periods, for example, 19.12.0.0.0

4. Exit the `root` user command shell:

   ```
   exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Resuming or Rolling Back a Patching Operation

You can resume or revert a failed patching operation. Reverting a patch is known as a rollback.

- Resuming a Patch Operation
  To resume a patching operation, use the `--resume` flag with the original patching command.
- Rolling Back a Patch Operation
  Use the `--rollback` flag with the original patching command to roll back (revert) a patching operation.

Resuming a Patch Operation

To resume a patching operation, use the `--resume` flag with the original patching command.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the original patching command to resume a patching operation:
   For example:

   ```
   dbaascli dbhome patch --oracleHome /u02/app/oracle/product/19.0.0.0/
   dbhome_2 --targetVersion 19.9.0.0.0 --resume
   ```

4. Exit the `root` user command shell:

   ```
   exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Rolling Back a Patch Operation

Use the `--rollback` flag with the original patching command to roll back (revert) a patching operation.

1. Connect to the virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. Run the original patching command to roll back (revert) a patching operation:
   For example:

   ```
   dbaascli grid patch --targetVersion 19.11.0.0.0 --rollback
   ```

> **Note**
>
> - Resume and Rollback operations are supported for Oracle Home patching, Oracle Grid Infrastructure patching, and database move operations.
> - When resuming or rolling back a patching operation, you must run the resume or rollback command from the same node that was used to run the original patching command, and you must run the original command with the addition of the `--resume` or `--rollback` flag.

4. Exit the `root` user command shell:

   ```
   exit
   ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Updating Cloud Tooling Using dbaascli

To update the cloud tooling release for Oracle Exadata Database Service on Cloud@Customer, complete this procedure.

Cloud-specific tooling is used on the Exadata Cloud@Customer 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

1. Connect to a virtual machine as the `opc` user.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

   ```
   sudo -s
   ```

3. To update to the latest available cloud tooling release, run the following command:

   ```
   dbaascli admin updateStack
   ```

   The command takes care of updating the cloud tooling release on all the nodes of the cluster.

   For more details and other available options, refer to `dbaascli admin updateStack --help`.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer 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
- Duplicate an On-Premises Database

Using dbaascli to Duplicate a Cloud Database

You can create a duplicate database using `dbaascli`. This new database can be in the same cloud region as the source region or across the regions. The following steps describe how to create a duplicate database on cloud.

> **Note**
>
> If a database is configured with OCI Vault for TDE encryption and you want to duplicate a database, then refer to the following sections.

Prepare for duplication

Ensure that the following prerequisites are 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 © 2008, 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.

Duplicate an On-Premises Database

Using `dbaascli`, you can duplicate an on-prem database onto the cloud. This can be done with the `dbaascli database duplicate` command. This command creates a new database on the cloud, which is a duplicate of an on-prem database along with its data. While this process is going on, the on-prem database remains still operational. You can migrate your applications to the duplicated database on the cloud after due verification.

Prepare for duplication

The migration process includes the following prerequisites to be met.

- Make sure that there is a network path setup to access an on-prem database from the OCI node through the `EZConnect` string.
- If an on-prem database is configured with TDE, copy the TDE wallet file (`ewallet.p12`) to the OCI node, where you decide to run the `dbaascli` command.
- Create an Oracle home on the OCI node if required. The Oracle home version must be the same as the source or of a higher RU version.

Verify the necessary RPMs

This process requires a minimum `dbaastools` RPM version of 23.3.2.0.0 but updating to the latest `dbaastools` rpm is always recommended.

- To check the currently installed version, run:

  ```
  dbaascli --version
  DBAAS CLI version 23.3.2.0.0
  ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

- To apply the latest tools RPM, as the `root` user, run:

  ```
  # dbaascli admin updateStack
  ```

Run the prerequisite checks

To run the prerequisite checks, use the `--executePrereqs` command option. This will perform only the prerequisite checks without performing the actual Oracle Database duplication.

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle Home
Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --executePrereqs
```

Duplicate the database

Duplicate the database using the following command:

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle Home
Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet>
```

For example:

```
dbaascli database duplicate --sourceDBConnectionString xyzhost.oracle.com:1521/
dbuniquename.oracle.com --dbName orcl --oracleHome /u02/app/oracle/product/
19.0.0.0/dbhome_1 --sourceDBTDEWalletLocation /tmp/wallet_copy/tde --
waitForCompletion false
```

Upon successful completion of this command, the database is duplicated to Cloud and ready for sanity checks for application usage. Once verification is done, application connections can be migrated to the Cloud database.

Refer to `dbaascli database duplicate –help` for additional configuration options.

Few considerations for migration

- If you prefer to allocate multiple channels for RMAN duplicate, you could do so by specifying the `--rmanParallelism` argument.
- Exadata Cloud Service configures database memory as Automatic Shared Memory Management (ASMM). If your on-prem database is configured with different memory management, make sure to adjust memory parameter values accordingly on the OCI side by providing values for `--sgaSizeInMB` and `--pgaSizeInMB`.
- Verify that the on-prem database does not contain any deprecated or invalid initialization parameters.
- Database initialization parameters related to database storage (datafile location, redo location, recovery area destination, control file multiplexing) may be changed using the `--initParams` argument.
  For example, to override `db_create_online_log_dest` value for the duplicate database: `--initParams db_create_online_log_dest_1=+DATAC1,db_create_online_log_dest_2=+RECOC1`

Troubleshooting the database duplication

- `dbaascli` operation log file can be found under `/var/opt/oracle/log/<dbname>/database/duplicate`

Copyright © 2008, 2026, Oracle and/or its affiliates.

- One of the jobs of the duplicate is to run `dbca`. Its log file can be found under `/u02/app/oracle/cfgtoollogs/dbca` and `/u02/app/oracle/cfgtoollogs/dbca/<dbuniquename>`.

If the operation fails, you will have an option to resume the operation by providing the `--resume` argument to the same command. Alternatively, clean up the database using `dbaascli database delete –dbname <dbname> –force`, and then rerun the database duplicate command.

Release Notes

Review the changes made in various releases of `dbaascli`.

- Release 26.1.1.0.0 (260302)
- Release 25.4.1.0.0 (251125.0831, 251125.0837)
- Release 25.3.1.0.0 (250826)
- Release 25.2.1.0.0 (250612)
- Release 25.1.2.0.0 (250505)
- Release 25.1.1.0.0 (250220)
- Release 24.4.1.0.0 (241212)
- Release 24.3.2.0.0 (240912)
- Release 24.3.1.0.0 (240730)
- Release 24.2.1.0.0 (240620)
- Release 24.1.2.0.0 (240327, 240424, 240502)
- Release 24.1.1.0.0 (240118, 240219)
- Release 23.4.1.0.0 (231219)
- Release 23.3.2.0.0 (231115)
- Release 23.3.1.0.0 (230817, 231020)
- Release 23.2.1.0.0 (230708, 230724)
- Release 23.1.2.0.0 (230411, 230616)
- Release 23.1.1.0.1 (230302)
- Release 22.4.1.0.1 (221214)
- Release 22.3.1.1.0 (221003)
- Release 22.3.1.0.1 (220831)
- Release 22.2.1.1.0 (220713)
- Release 22.2.1.0.1 (220504)
- Release 22.1.1.1.0 (220301)
- Release 22.1.1.0.1 (220223)
- Release 21.4.1.1.0
- Release 21.3.1.2.0
- Release 21.3.1.1.0
- Release 21.3.1.0.1

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Release 21.2.1.x.x

Release 26.1.1.0.0 (260302)

This release includes:

- AHF 26.1.1
- syslens 25.4.3.0
- Various bug fixes and stability improvements.

Release 25.4.1.0.0 (251125.0831, 251125.0837)

This release includes:

- AHF 25.9.1
- syslens 25.3.3.0
- Various bug fixes and stability improvements.

Release 25.3.1.0.0 (250826)

- Includes AHF 25.6.2
- Includes syslens 25.2.3.0
- The `dbaascli dbhome patch` command now includes the `skipDBForDatapatch` option. **Purpose:** Allows you to specify a list of databases on which the datapatch step will be skipped during in-place DB home patching.
- Various bug fixes and stability improvements

Release 25.2.1.0.0 (250612)

- Includes AHF 25.2.4
- Includes syslens 25.1.4.0
- Various bug fixes and stability improvements

Release 25.1.2.0.0 (250505)

- Includes AHF 25.2.4
- Includes syslens 25.1.1.0
- Various bug fixes and stability improvements

Release 25.1.1.0.0 (250220)

- Includes AHF 24.11.0
- Includes syslens 24.4.2.0
- Various bug fixes and stability improvements

Release 24.4.1.0.0 (241212)

- Includes AHF 24.9.1

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Includes syslens 24.3.3.1
- Various bug fixes and stability improvements

Release 24.3.2.0.0 (240912)

- Includes AHF version 24.7
- Includes syslens version 24.3.1.0
- Various bug fixes and stability improvements

Release 24.3.1.0.0 (240730)

- TDE migration from `sqlnet.ora` to `wallet_root` on database upgrade to 19c release
- Grid patch in-place to use image-based patching as default mode
- Includes AHF version 24.6.1
- Includes syslens version 24.1.2.0
- Various bug fixes and stability improvements

Release 24.2.1.0.0 (240620)

- Added support for Oracle AI Database 26ai.
- Includes AHF version 24.4.3
- Includes syslens version 24.1.2.0
- Various bug fixes and stability improvements.

Release 24.1.2.0.0 (240327, 240424, 240502)

- Includes AHF version 24.1.1.
- Includes syslens version 2.6.8.0.
- Various bug fixes and stability improvements

Release 24.1.1.0.0 (240118, 240219)

- Includes AHF version 23.11.1.
- Includes syslens version 2.6.4.3.
- Various bug fixes and stability improvements
- (240219) Bug 36309260 fix applicable to DBCS agent version 21.2 or earlier: communication between OCI Control Plane and DBCS agent does not work in some VMs.

Release 23.4.1.0.0 (231219)

- Includes AHF version 23.9.5.
- Includes syslens version 2.6.4.2.
- Various bug fixes and stability improvements

Copyright © 2008, 2026, Oracle and/or its affiliates.

Release 23.3.2.0.0 (231115)

- Pluggable Database Operations
  - Added support to set custom key version OCID (Bring Your Own Key - BYOK) of OCI Vault during create and clone operations. For details, see respective PDB commands help.
- Grid Infrastructure (GI) Patching
  - Enhanced the patching workflow to improve patching time, especially in environments having high number of databases.
  - A new option `--patchInParallel` is introduced that can be used to perform patching remote nodes in parallel.
- Database Patching
  - Provided option to run datapatch on a specific node of cluster.
- Includes AHF version 23.7.7
- Includes syslens version 2.3.6.10
- Various bug fixes and stability improvements

Release 23.3.1.0.0 (230817, 231020)

- New `dbaascli` commands
  - `dbaascli gridHome create` - This command can be used to create a Grid Infrastructure home of a supported version. For details, see `dbaascli gridHome create --help`.
  - `dbaascli system getGridHomes` - This command gives details on the available Grid Infrastructure homes on the system. For details, see `dbaascli system getGridHomes --help`.
- Pluggable Database Operations
  - Improvements in the area of refreshable pluggagble database lifecycle.
- Database Backup and Recovery
  - Added support to configure backups on standby sites in case of Data Guard configurations. The backups configuration are Data Guard site-specific, that is, the change of roles (for example, with Data Guard switchover operation) will not impact the backup operations of the database on primary or standby sites. Backups, if configured on primary site or stand-by site, will continue regardless of the role-change.
- – Includes AHF version 23.5.2 - Release 23.3.1.0.0 (**230817**)
  - Includes AHF version 23.5.4 - Release 23.3.1.0.0 (**231020**)
- Includes syslens version 2.3.6.9
- Various bug fixes and stability improvements
- Additional critical product fixes (**231020**)

Release 23.2.1.0.0 (230708, 230724)

- Database Lifecycle related improvements

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - Introduced `dbaascli grid removeTCPSCert` to remove expired TCPS certificates. For details, see `dbaascli grid removeTCPSCert --help`.
  - Added option to exclude specific PDBs during database duplicate. For details, see `skipPDBs` argument in `dbaascli database duplicate --help`.
- Database Backup and Recovery
  - Changed the default for `FILES_PER_SET` to 64 for OSS backups. This can be changed with `dbaascli database backup --configure`. For details, see `dbaascli database backup --help`.
  - Archive log backups continue from the standby site after the role switchover in data guard environments.
  - For backups that are not managed by Oracle, the schedules for L0 and L1 backups are not created by default. They must be be created explicitly by using `dbaascli database backup --configure` command.
- – Includes AHF version 23.3.4 - Release 23.2.1.0.0 (**230708**)
  - Includes AHF version 23.3.5 - Release 23.2.1.0.0 (**230724**)
- Various bug fixes and stability improvements

Release 23.1.2.0.0 (230411, 230616)

- Database Lifecycle related improvements
  - Added option to create database templates (DBCA temapltes) to object store. DBCA templates can subsequently be used to create databases. For details, see `dbaascli database createTemplate --help`.
- Pluggable Database Operations
  - Introduced `dbaascli pdb refresh` to refresh a pluggable database that was created using manual refresh option. For details, see `dbaascli pdb refresh --help`.
  - Added option to convert refreshable pluggable database to a regular pluggable database. For details, see `dbaascli pdb open --help`.
  - Creation of a refreshable pluggable database now requires existing source database user for creation of database link to the source pluggable database. For details, see `dblinkUserName` argument in `dbaascli pdb remoteClone --help`.
- Includes AHF version 23.2.0
- Various bug fixes and stability improvements

Release 23.1.1.0.1 (230302)

- Database Lifecycle related improvements
  - Added support to create a duplicate database from a source database which uses OCI Vault Services for encryption key management.
- Includes AHF version 22.2.5
- Various bug fixes and stability improvements

Release 22.4.1.0.1 (221214)

- Pluggable Database Operations

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - Added option to not open the PDB at the end of relocate. For details, see `skipOpenPDB` argument in `dbaascli pdb relocate --help`. After using this option, the pdb relocate can be completed by running the command using `completePDBRelocate` argument.
  - Added option to clean up the relocated PDB metadata/services at the source location. For details, see `cleanupRelocatedPDB` argument in `dbaascli pdb delete --help`
- New `dbaascli` commands
  - `dbaascli database createTemplate` - This command can be used to create database templates (DBCA templates) that can subsequently be used to create databases. DBCA templates are widely used for creating a clone database with DBCA - a tool that is shipped with Oracle Database server software. For details, see `dbaascli database createTemplate --help`
  - Introduced `dbaascli tde rotateMasterKey` to rotate the master key for database encryption. For details, see `dbaascli tde rotateMasterKey --help`. The command `dbaascli tde rotate masterkey` is now deprecated.
- Database Lifecycle related improvements
  - Added support to use dbca templates in database creation workflows. For details, see `dbcaTemplateFilePath` argument in `dbaascli database create --help`
  - Improved performance for duplicate database creation. For details on how to create duplicate database, see `dbaascli database duplicate --help`
  - Added support to create a duplicate database from a source database which is not TDE-encrypted.
- TDE management
  - Introduced `dbaascli tde rotateMasterKey` to rotate the master key for database encryption. For details, see `dbaascli tde rotateMasterKey --help`. The command `dbaascli tde rotate masterkey` is now deprecated.
  - Revamped workflow for all TDE operations. For details, see `dbaascli tde --help`
- Grid Infrastructure (GI) Patching
  - Added support to allow parallel execution of patching operation on nodes. This option needs to be carefully exercised as it results into reduced database availability.
- Database Backup and Recovery
  - Revamped workflow for creating database from standalone backups
- Includes AHF version 22.2.4
- Various bug fixes and stability improvements

Release 22.3.1.1.0 (221003)

- New `dbaascli` commands
  - `dbaascli database getDetails` - This command shows the detailed information of a given database, for example, dbname, node information, pluggable databases information, and so on. For details, see `dbaascli database getDetails --help`.
- Pluggable Database Operations
  - Added support for creating pluggable databases as refreshable clone using `refreshablePDB` argument. For details, see `dbaascli pdb remoteClone --help`
- Various bug fixes and stability improvements

Copyright © 2008, 2026, Oracle and/or its affiliates.

Release 22.3.1.0.1 (220831)

- New database lifecycle commands
  - `dbaascli database addInstance` - This command can be used to add a database instance to one of the nodes of the cluster where database is not already configured. For details, see `dbaascli database addInstance --help`.
  - `dbaascli database deleteInstance` - This command can be used to delete a database instance from one of the nodes of the cluster where database is configured. For details, see `dbaascli database deleteInstance --help`.
  - `dbaascli database duplicate` - This command can be used to create a new database from an already existing database within a cluster, or across clusters, provided network connection exists between the clusters. For details, see `dbaascli database duplicate --help`.
- Cloud Software Library
  - Introduced `dbaascli cswlib listLocal` command to list images that are downloaded from software library locally on the system. For details, see `dbaascli cswlib listLocal --help`. The command `dbaascli dbimage list` is now deprecated.
  - Introduced `dbaascli cswlib deleteLocal` command to delete images that are downloaded from cloud software library. For details, see `dbaascli cswlib deleteLocal --help`. The command `dbaascli dbImage purge` is now deprecated.
- The log location for the command `dbaascli admin updateStack` has been changed to follow the convention of other `dbaascli` commands. The logs can be conveniently found under `/var/opt/oracle/log/admin/updateStack` directory. The earlier location was `/var/opt/oracle/log/tooling/Update`.
- `dbaascli` help is now cloud platform aware in that it will list help output for commands applicable for the cloud environment it is operating on.
- Added support for changing TDE password in dataguard environments. For details, see `dbaascli tde changePassword --help`. This support is currently not available for 11.2.0.4 release.
- Included AHF version 22.1.5.
- Revamped workflow for database upgrade operation.
- Revamped workflow for database home create operation.
- Various bug fixes and stability improvements

Release 22.2.1.1.0 (220713)

- New `dbaascli` commands:
  - `dbaascli dbHome getDatabases` - This command lists all the databases running from a given database Oracle home. The output is returned in JSON format to facilitate automation. For details, see `dbaascli dbHome getDatabases --help`.
  - `dbaascli database getPDBs` - This command lists all the pluggable databases of a given container database. The output is returned in JSON format to facilitate automation. For details, see `dbaascli database getPDBs --help`.
  - `dbaascli dbHome delete` - This command deletes a given database Oracle home. For details, see `dbaascli dbHome delete --help`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - `dbaascli dataguard prepareStandbyBlob` - This command generates a blob file containing various files that are required on the standby site for a Data Guard environment. For details, see `dbaascli dataguard prepareStandbyBlob --help`.
- Grid Infrastructure (GI) Patching:
  - New optimized workflow
  - Introduced a way to create the Grid Infrastructure (GI) software image prior to patching. This GI image can be subsequently used for performing the GI patching operation. The advantage of this approach is that it results in reduced patching window as the image is already prepared. The GI stack on the node is not brought down to create the image. For details, see `createImage` option in `dbaascli grid patch --help`
  - Introduced a way to perform the Grid Infrastructure patching through the use of user specified GI software image, created using `createImage` option of the `dbaascli grid patch` command. For details, see `imageLocation` option in `dbaascli grid patch --help`.
- Change Password support in Data Guard environment:
  - Added support to change password in Data Guard environments. For details, see `dbaascli database changePassword --help` and `dbaascli dataguard prepareStandbyBlob --help`
- Data Guard configuration:
  - Added support to update Data Guard Automation Attributes (in the `/var/opt/oracle/dg/dg.conf` file). For details, see `dbaascli dataguard --help.`
- Various bug fixes and stability improvements

Release 22.2.1.0.1 (220504)

- New `dbaascli` commands:
  - Introduced `dbaascli admin showLatestStackVersion` to show the latest **dbaastools** version available for customers to download and install. The installation of **dbaastools** RPM can be performed by using the command `dbaascli admin updateStack`. For details see the `dbaascli` Command Reference section.
- Cloud Software Library:
  - Deprecated the support for BP activation (`dbaascli cswlib activateBP`) as BPs (Bundle Patches) are now replaced with RUs ("Release Updates"). Cloud deployment consumes RUs in the form of software images, identified with `Image Tags`. It is therefore recommended to use image tags while interfacing with Cloud Software Library (`cswlib`) commands. For details, see `dbaasscli cswlib download –help`.
  - Eliminated the need to download non-CDB images to create non-CDB databases. Now users can create the non-CDB database using regular images. For details, see `createAsCDB` option in `dbaascli database create –help`.
- Non-CDB Database Creation:
  - Enhanced database creation workflow to create a non-CDB database using standard database software image. For details, see `createAsCDB` option in `dbaascli database create –help`.
- Database Home Patching:
  - New optimized workflow

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Grid Infrastructure Upgrade:
  - New optimized workflow
- Pluggable Database (PDB) Operations:
  - Deletion of PDB in Data Guard environments requires explicit acknowledgment to indicate that operations necessary on standby site are completed, by passing of additional argument `–allStandByPrepared`. For details, see `dbaascli pdb delete --help`.
- Provided rolling capability for database bounce operation. For details, see `dbaascli database bounce –help`.
- Various bug fixes and stability improvements

Release 22.1.1.1.0 (220301)

- New `dbaascli` commands:
  - Introduced `dbaascli system getDBHomes` to get all the database Oracle homes on the cluster. The output is returned in JSON format to facilitate automation.
  - Introduced `dbaascli dbhome getDetails` to get detailed information on a specific Oracle home. The output is returned in JSON format to facilitate automation.
- Cloud Software Library (cswlib):
  - Deprecated the support for `dbaascli cswlib list` command for cloud software library listing operations. The new command is `dbaascli cswlib showImages` that lists the images along with its of `ImageTag`. It is recommended to use `Image tags` to download the images from the cloud software library. For details on downloads using image tags, see `dbaascli cswlib download –help`.
  - Various bug fixes and stability improvements

Release 22.1.1.0.1 (220223)

- Grid Infrastructure Upgrade:
  - New optimized workflow
- Database Backup and Recovery:
  - Internal update to metadata repository for backup metadata
  - Introduced deprecation messages for `bkup_api` commands as they are now replaced with `dbaascli` commands. For details, see `dbaascli database backup --help` and `dbaascli database recover –help`
- Pluggable Database (PDB) Operations:
  - Relocate operation of PDB is now supported. For details, see `dbaascli pdb relocate –help`.
  - Revamped workflow for non-CDB to PDB conversion. For details, see `dbaascli database convertToPDB –help`.
- Encryption Key Management:
  - Transparent Data Encryption (TDE) heartbeat-specific initialization parameters are set to the cloud recommended values for databases with 'Customer Managed Keys'.
- Cloud Software Library Management:

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - Revamped software library download of artifacts through `imageTags`. It is recommended to use `imageTags` to download the database and grid software images. For details, see `dbaascli cswlib showimages` and `dbaascli cswlib download –help`
- Included AHF version 21.4.2
- Various bug fixes and stability improvements

Release 21.4.1.1.0

- Enabled encryption of the system level tablespaces (`SYSTEM`, `SYSAUX`, `UNDO`, and `TEMP`) for databases that will get created with this version of dbaastools onwards. This feature is enabled for Oracle Database version 19.6.0.0.0 and above.
- Grid Patching:
  - Prerequisite condition added to check for following file ownership to be owned by `grid` user.
    * `<gi_home>/suptools/tfa/release/tfa_home/jlib/jdev-rt.jar`
    * `<gi_home>/suptools/tfa/release/tfa_home/jlib/jewt4.jar`
- Database Patching:
  - Simultaneous `database move` operation is disallowed by default. A new option `–allowParallelDBMove` is introduced that can be used to override the default behavior for Oracle Database releases 12.2 and above.
  - Fixed issues related to move of standby databases being in `MOUNT` mode.
- Database Backup and Recovery:
  - Added new command-line options for database backup. For more details, refer to *dbaascli database backup* command reference.
  - Added new command-line options for database recovery. For more details, refer to *dbaascli database recover* command reference.
  - `bkup_api` usage for backup and recovery operations will be deprecated in future.
  - To align with the Oracle recommended practice of using `SYSBACKUP` administrative privilege for Backup and Recovery operations, cloud automation creates a common administrative user `C##DBLCMUSER` with `SYSBACKUP` role at the `CDB$ROOT` container level. Backup and Recovery operations are therefore performed with the user having the least required privileges. Credentials for this user are randomly generated and securely managed by cloud automation. If the user is not found or is `LOCKED` and `EXPIRED`, then cloud automation will recreate or unlock this user during the backup or recovery operation. This change in the cloud automation is made starting with dbaastools version 21.4.1.1.0.
- Enhanced `dbaascli resume` functionality to resume any previous session by specifying the `–sessionID <value>` argument to the resume command. The session ID is shared in the `dbaascli` output as well as in the logs.
- Enhanced `dbaascli help` output to show the command usage.
- Deprecated the usage of `dbaascli` shell (interactive session). This will be completely unsupported after March 2022. It is recommended to execute complete `dbaascli` commands on command prompt as suggested in all document examples.
- Included Autonomous Health Framework (AHF) version 21.2.8.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Various bug fixes and stability improvements.

Release 21.3.1.2.0

- Improved the timing of `dbaascli` operations with enhanced Control Plane metadata synchronization logic.
- Enhanced `dbaascli` logs to have millisecond-level information along with the associated thread.
- Introduced more prerequisite checks in database home patching and database move operations to catch potential failures scenarios with suggestions to corrective action.
- Database patching operations now retain the state of the databases to be same as it was prior to patching. For pluggable databases, pdb saved state is honored.
- Various bug fixes and stability improvements.

Release 21.3.1.1.0

- Added support to unlock PDB Admin user account as part of PDB creation, `localClone`, or `remoteClone` operation. For details, see option `--lockPDBAdminAccount` in `dbaascli pdb create --help`.
- Fixed an issue that updates the database resource registered with Oracle Grid Infrastructure in existing environments with the correct value of database name.
- Enhanced PDB lifecycle operations.
- Various bug fixes and stability improvements.

Release 21.3.1.0.1

- Support for the following `dbaascli` commands to be run as `oracle` user.
  - `dbaascli pdb bounce`
  - `dbaascli pdb close`
  - `dbaascli pdb connectString`
  - `dbaascli pdb create`
  - `dbaascli pdb delete`
  - `dbaascli pdb getDetails`
  - `dbaascli pdb list`
  - `dbaascli pdb localClone`
  - `dbaascli pdb open`
  - `dbaascli pdb remoteClone`
- Revamped out-of-place patching of database. For details, see `dbaascli database move –help`.
- Timing related enhancements in Oracle Grid Infrastructure patching workflow. For details, see `dbaascli grid patch –help`.
- Deprecated the support for `exadbcpatchmulti` / `dbaascli patch` for patching operations. The `dbaascli dbhome patch` and `dbaascli grid patch` commands are provided for

Copyright © 2008, 2026, Oracle and/or its affiliates.

patching operation for database homes and Oracle Grid Infrastructure. Refer to the *Patching Oracle Grid Infrastructure* and *Oracle Database Using dbaascli* section for details. Also see, *dbaascli Command Reference* section.

- Deprecated the support for `dbaascli` tools patch command to bring consistency in the `dbaascli` command conventions. The new command is `dbaascli admin updateStack`. For details, see section *Updating Cloud Tooling using dbaascli*.
- Ability to run `dbaascli` in disconnected mode for long running operations. Executing `dbaascli` command with `--waitForCompletion false` gets you a job ID that can be queried later to get the status of the operation, using `dbaascli job getStatus –jobid job_id`. This is useful for long running operations where users may want to get the control back immediately after command execution. In this release, this option is available only for `dbaascli database create` command. More commands will be added in subsequent releases to have this support. The help output for those commands will reflect the support of `--waitForCompletion` option.
- Deprecated the support for `dbaascli` shell. It is recommended that users run the complete `dbaascli` commands on the command prompt as suggested in all the document examples. Execution of just `dbaascli` will show the output of its usage help instead of entering into a `dbaascli` shell.
- Various bug fixes and stability improvements.

Release 21.2.1.x.x

- Redesigned Oracle Grid Infrastructure patching operation and added ability to resume from failed point, patch on subset of nodes, instance draining, and other enhancements. For details, see `dbaascli grid patch --help`. Also refer to the *Patching Oracle Grid Infrastructure and Oracle Database Using dbaascli* section.
- Deprecated the support for `exadbcpatchmulti` / `dbaascli patch` for patching operations. `dbaascli dbhome patch` and `dbaascli grid patch` commands are provided for patching operation for database homes and Oracle Grid Infrastructure. Refer to the *Patching Oracle Grid Infrastructure and Oracle Database Using dbaascli* section for details. Also see, *dbaascli Command Reference* section.
- Deprecated the support for `dbaascli tools patch` command to bring consistency in the command conventions. The new command is `dbaascli admin updateStack`.
- Redesigned PDB management APIs for create, local clone, and remote clone operations. For details, see `dbaascli pdb --help`.
- Redesigned database delete API. For details, see `dbaascli database delete --help`.
- Revamped dbhome creation (support for custom software image, scale-out operation). For details, see `dbaascli dbhome create --help`.
- Support for database creation on subset of cluster nodes. For details, see `dbaascli database create --help`.
- Ability to run `dbaascli` in disconnected mode for long running operations. Executing `dbaascli` command with `--waitForCompletion false` gets you a job ID that can be queried later to get the status of the operation, using `dbaascli job getStatus –jobid job_id`. This is useful for long running operations where users may want to get the control back immediately after command execution. In this release, this option is available only for `dbaascli database create` command. More commands will be added in subsequent releases to have this support. The help output for those commands will reflect the support of `--waitForCompletion` option.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Enhanced dbhome patching experience with introduction of multiple options like `skipPDBs`, `continueWithDowntime`, and so on. For details, see `dbaascli dbhome patch --help`.
- Support for better diagnostic collection. For details, see `dbaascli diag collect --help`.
- Minor improvements in the area of database upgrade automation.
- Various bug fixes and stability improvements.

dbaascli Command Reference

You must use `dbaascli` to create databases and integrate them with the cloud automation framework.

`dbaascli` is a cloud native interface that can take DBCA templates as inputs, calls the functionality of DBCA to create databases, and then calls OCI APIs to integrate the database into the cloud automation framework. Customers using DBCA in scripts today can update their existing scripts to call `dbaascli` instead of DBCA. If `dbaascli` cannot be used due to a particular feature of DBCA being unavailable in dbaascl, then customers should open a My Oracle Support (MOS) request to add that functionality to `dbaascli`.

- Administration and Configuration
  This section covers essential tasks for managing and configuring Oracle Database environments. It includes commands such as `dbaascli admin updateStack` for installing or updating the `dbaastools` RPM, and `dbaascli job getStatus` for checking the status of specific jobs.
- CPU Scaling
  This section focuses on adjusting CPU resources in a VM cluster. It includes commands like `dbaascli cpuscale get_status` to check the status of current or past scaling requests, and `dbaascli cpuscale update` to increase or decrease the number of CPU cores allocated to a virtual machine, allowing flexible resource management based on workload demands.
- Cloud Software Library (CSWLIB) Management
  This section provides tools for managing software images in Exadata Database Service on Cloud@Customer environments. Commands like `dbaascli cswlib deleteLocal` allow for the removal of local images, while `dbaascli cswlib download` enables the download of new software images. You can also view available images locally with `dbaascli cswlib listLocal` or check all available Database and Grid Infrastructure images using `dbaascli cswlib showImages`. These commands help manage and maintain software libraries efficiently.
- Database Management
  This section covers comprehensive tasks for managing Oracle Databases. It includes commands for creating (`dbaascli database create`), deleting (`dbaascli database delete`), and upgrading databases (`dbaascli database upgrade`). Other key tasks include adding and deleting instances (`dbaascli database addInstance`, `dbaascli database deleteInstance`), managing backups (`dbaascli database backup`), and handling database recovery (`dbaascli database recover`). You can also modify database parameters, manage pluggable databases, patch databases, and convert non-CDB databases to PDBs. These commands ensure efficient control over the entire database lifecycle.
- Data Guard Management
  This section focuses on managing Oracle Data Guard configurations and operations. It includes commands like `dbaascli dataguard prepareStandbyBlob` to generate a blob file for setting up a standby site, and `dbaascli dataguard updateDGConfigAttributes` to update Data Guard automation attributes across all cluster nodes. These commands

Copyright © 2008, 2026, Oracle and/or its affiliates.

streamline the setup and maintenance of Data Guard environments for high availability and disaster recovery.

- Database Home Management
  This section provides tools for managing Oracle Database homes. It includes commands such as `dbaascli dbhome create` to create a new Oracle Database home and `dbaascli dbHome delete` to remove an existing one. You can also view detailed information about a specific Oracle home with `dbaascli dbHome getDetails` and check which databases are running from a given Oracle home using `dbaascli dbhome getDatabases`. These commands ensure efficient management of database environments.
- Diagnostics and Health Checks
  This section covers tools for maintaining the health and diagnosing issues in Oracle Database environments. Commands like `dbaascli diag collect` are used to gather diagnostic data, while `dbaascli diag healthCheck` allows you to run health checks to identify potential issues. These tools help ensure the system's stability and performance by proactively monitoring and addressing concerns.
- Grid Infrastructure Management
  This section focuses on managing Oracle Grid Infrastructure, which supports clustering and high availability. Key tasks include configuring a Grid Infrastructure home (`dbaascli gridHome create`), upgrading the Grid Infrastructure (`dbaascli grid upgrade`), and managing TCPS (Transport Layer Security) certificates by configuring (`dbaascli grid configureTCPS`), removing (`dbaascli grid removeTCPSCert`), or rotating them (`dbaascli grid rotateTCPSCert`). These commands ensure efficient setup, maintenance, and security of the Grid Infrastructure.
- Patching and Upgrading
  This section provides tools for updating and maintaining Oracle environments through patching and upgrades. It includes commands such as `dbaascli grid patch` to patch Oracle Grid Infrastructure, `dbaascli dbHome patch` to apply patches to Oracle homes, and `dbaascli database move` to move databases between homes during upgrades or patching processes. These commands help ensure systems remain secure, stable, and up-to-date.
- Pluggable Database (PDB) Management
  This section covers the management of Pluggable Databases (PDBs) within a container database (CDB). It includes commands for creating (`dbaascli pdb create`), deleting (`dbaascli pdb delete`), and cloning PDBs (`dbaascli pdb localClone`, `dbaascli pdb remoteClone`). You can manage PDB states with commands to open, close, or bounce PDBs, and retrieve connection details (`dbaascli pdb getConnectString`). Additional commands support backup, recovery, and relocating PDBs, ensuring comprehensive control over PDB lifecycle and operations.
- System Management
  This section focuses on overseeing and managing Oracle homes within the system. It includes commands like `dbaascli system getDBHomes` to view details about all Oracle Database homes and `dbaascli system getGridHomes` to list the details of all Grid Infrastructure homes. These commands provide essential information for maintaining and troubleshooting the overall system environment.
- Transparent Data Encryption (TDE) Management
  This section encompasses the management of Transparent Data Encryption (TDE) for securing database data. It includes commands for handling encryption keys and keystores, such as `dbaascli tde addSecondaryHsmKey` to add secondary HSM keys, `dbaascli tde rotateMasterKey` to rotate the master key, and `dbaascli tde encryptTablespacesInPDB` to encrypt tablespaces within a pluggable database. You can also convert between FILE and HSM-based TDE (`dbaascli tde fileToHsm`, `dbaascli tde hsmToFile`), manage key versions, and retrieve key details with various commands. These tools ensure effective encryption management and data security.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Deprecated dbaascli commands
  The `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.

Administration and Configuration

This section covers essential tasks for managing and configuring Oracle Database environments. It includes commands such as `dbaascli admin updateStack` for installing or updating the `dbaastools` RPM, and `dbaascli job getStatus` for checking the status of specific jobs.

- dbaascli admin updateStack
  To install or update a dbaastools RPM, use the `dbaascli admin updateStack` command.
- dbaascli job getStatus
  To view the status of a specified job, use the `dbaascli job getStatus` 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 Exadata Database Service on Cloud@Customer 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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `--resume` option resumes the previous execution of the `updateStack` command if it was interrupted or incomplete.

Q: What is the purpose of the --prechecksOnly option?

A: The `--prechecksOnly` option runs only the prechecks for the operation without actually performing the installation or update.

Q: How is the --nodes parameter used?

A: The `--nodes` parameter specifies a comma-delimited list of nodes on which the RPM should be installed. If not provided, the RPM will be installed on all cluster nodes.

Q: What should I do if I encounter issues with the dbaascli admin updateStack command?

A: Ensure you are running the command as the root user and that you are connected to an Exadata Cloud Infrastructure virtual machine. Check if there are any specific error messages and consult the command documentation or Oracle support if needed.

Q: How do I connect to an Exadata Cloud Infrastructure virtual machine to use the dbaascli admin updateStack command?

A: You need to use SSH to connect to the virtual machine. Refer to the section on "Connecting to a Virtual Machine with SSH" in the documentation for detailed instructions.

Example Use Cases

Example 1: Installing or updating the dbaastools RPM on all nodes

```
dbaascli admin updateStack
```

Installs or updates the `dbaastools` RPM on all nodes of the Exadata Cloud@Customer environment.

Example 2: Running prechecks only before installing or updating the RPM

```
dbaascli admin updateStack --prechecksOnly
```

Runs only the prechecks for the `dbaastools` RPM update, without actually performing the installation. It ensures that all prerequisites are satisfied before proceeding with the update.

Example 3: Resuming a previously interrupted updateStack operation

```
dbaascli admin updateStack --resume
```

Resumes a previous `dbaastools` RPM update operation that was interrupted or did not complete successfully.

Example 4: Installing or updating dbaastools on specific nodes

```
dbaascli admin updateStack --nodes node1,node2
```

Installs or updates the `dbaastools` RPM on the specified nodes *node1* and *node2* only, without affecting other nodes in the cluster.

Example 5: Resuming the updateStack process on specific nodes

```
dbaascli admin updateStack --resume --nodes node3,node4
```

Resumes the update process for `dbaastools` on the specific nodes node3 and node4 only, if the previous execution was interrupted.

Copyright © 2008, 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 Cloud@Customer system by using a Secure Shell (SSH) connection.

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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, *Connecting to a Virtual Machine with SSH*.

Syntax

```
dbaascli job getStatus --jobID
```

Where:

- `--jodID` specifies the job ID

Frequently Asked Questions

Q: What is the dbaascli job getStatus command used for?

A: The `dbaascli job getStatus` command is used to view the status

Example Use Cases

Example 1: Checking the status of a specific job using job ID

`dbaascli job getStatus --jobID` *`12345`*

Checks the status of the job with the ID *12345*. The output will show the current state of the job (e.g., in progress, completed, or failed).

Example 2: Checking the status of a patching job using job ID

`dbaascli job getStatus --jobID` *`98765`*

Retrieves the status of a patching job with the ID *98765* to see whether the patch has been applied successfully or is still running.

Example 3: Checking the status of a database backup job

`dbaascli job getStatus --jobID` *`45678`*

Checks the status of a database backup job with the ID *45678*. The output will provide details about the progress or completion of the backup.

Example 4: Checking a long-running job's progress

`dbaascli job getStatus --jobID` *`23456`*

Check the progress of a long-running job (ID *23456*) to see if it's still executing or has finished.

Example 5: Viewing the status of a database creation job

Copyright © 2008, 2026, Oracle and/or its affiliates.

```
dbaascli job getStatus --jobID 67890
```

Checks the status of a database creation job identified by the job ID *67890*.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

CPU Scaling

This section focuses on adjusting CPU resources in a VM cluster. It includes commands like `dbaascli cpuscale get_status` to check the status of current or past scaling requests, and `dbaascli cpuscale update` to increase or decrease the number of CPU cores allocated to a virtual machine, allowing flexible resource management based on workload demands.

- dbaascli cpuscale get_status
  To check the status of current or last scale request performed when network connectivity between the Control Plane Server and OCI region is disrupted, use the `dbaascli cpuscale get_status` command.
- dbaascli cpuscale update
  To scale up or down the CPU core count for a virtual machine in a VM cluster when network connectivity between the Control Plane Server and OCI region is disrupted, use the `dbaascli cpuscale update` command.

dbaascli cpuscale get_status

To check the status of current or last scale request performed when network connectivity between the Control Plane Server and OCI region is disrupted, use the `dbaascli cpuscale get_status` command.

Prerequisites

Run the command as the `root` user.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, *Connecting to a Virtual Machine with SSH*.

Syntax

Displays various command execution states as it progresses from `scheduled`, `running`, and finally to `success` or `failure`.

```
dbaascli cpuscale get_status
```

Frequently Asked Questions

Q: What is the dbaascli cpuscale get_status command used for?

A: The `dbaascli cpuscale get_status` command is used to check the status of the current or last CPU scale request, especially when network connectivity between the Control Plane Server and OCI region is disrupted.

Q: What are the prerequisites for using the dbaascli cpuscale get_status command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

Q: What should I do if I encounter issues running the dbaascli cpuscale get_status command?

A: Ensure you are running the command as the `root` user and that you are connected to an Exadata Cloud@Customer virtual machine. If issues persist, consult the command documentation or seek support from Oracle.

Q: What if the command shows a status of failure?

A: If the command shows a failure status, check the detailed logs for error messages and troubleshoot based on the specific error. You may need to address network issues or review the scale request details.

Example Use Cases

Example 1: Checking the status of the most recent CPU scaling operation

```
dbaascli cpuscale get_status
```

Checks the status of the current or last CPU scaling request. It will provide information on whether the scaling is scheduled, running, or has completed with success or failure.

Example 2: Checking the status after a scaling request failed

A scaling operation was requested but encountered network issues between the Control Plane Server and the OCI region.

```
dbaascli cpuscale get_status
```

Checks the status of the scaling request. Since the scaling process failed due to network issues, the output will provide details about the failure state.

Example 3: Checking the status when scaling is in progress

A CPU scaling operation is ongoing, and the user wants to monitor its progress.

```
dbaascli cpuscale get_status
```

Checks the current status, showing that the scaling request is in the "running" state. It allows the user to track the operation until it completes or fails.

Example 4: Checking the status after successful scaling completion

A scaling operation was performed and completed successfully.

```
dbaascli cpuscale get_status
```

Checks the status and confirms that the scaling operation has successfully completed. It reports the final state as "success."

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli cpuscale update

To scale up or down the CPU core count for a virtual machine in a VM cluster when network connectivity between the Control Plane Server and OCI region is disrupted, use the `dbaascli cpuscale update` command.

To scale OCPUs up or down in a VM cluster in disconnected mode, run the `dbaascli cpuscale update` and `dbaascli cpuscale get_status` commands from any node inside a VM cluster to change the CPU core count for that cluster. If you have more than one VM cluster, then run a separate command from any node inside each VM cluster you want to scale up or down. These commands are designed to not work if issued during the normal connected mode and will time out after 600 seconds (10 minutes).

Prerequisites

Run the command as the `root` user.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, *Connecting to a Virtual Machine with SSH*.

Syntax

Exadata Database Service on Cloud@Customer is considered to be in a **Disconnected** mode when there is a loss of connectivity with the DBaaS control plane running on Oracle Cloud Infrastructure (OCI).

```
dbaascli cpuscale update --coreCount coreCount --message message
```

Where:

- `--coreCount` specifies the target number of CPUs you want to scale up or down for each VM in the cluster
- `--message` optionally, you can include a message for your reference

Frequently Asked Questions

Q: What is the dbaascli cpuscale update command used for?

A: The `dbaascli cpuscale update` command is used to scale up or down the CPU core count for a virtual machine in a VM cluster when network connectivity between the Control Plane Server and Oracle Cloud Infrastructure (OCI) is disrupted.

Q: What are the prerequisites for using the dbaascli cpuscale update command?

A: Before using this command, ensure you are running it in disconnected mode, meaning there is a loss of connectivity with the DBaaS control plane on OCI. Run the command from any node inside the VM cluster, and note that it will time out after 600 seconds (10 minutes) if used in connected mode. The command must be executed as the root user.

Q: How do I connect to a virtual machine to use this command?

A: To use the `dbaascli cpuscale update` command, you need to connect to an Exadata Cloud@Customer virtual machine using SSH. Refer to "Connecting to a Virtual Machine with SSH" for detailed instructions.

Q: What does the --coreCount option specify in the command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `--coreCount` option specifies the target number of CPUs you want to scale up or down for each VM in the cluster.

Q: Can I include a message with the dbaascli cpuscale update command?

A: Yes, you can include an optional message using the `--message` option for your reference.

Q: How do I check the status of a CPU scale operation?

A: To check the status of a CPU scale operation, use the `dbaascli cpuscale get_status` command. This should be run from any node inside the VM cluster as well.

Q: What happens if I run the dbaascli cpuscale update command in connected mode?

A: The command is designed to not work in connected mode and will time out after 600 seconds (10 minutes). It should only be used in disconnected mode.

Q: How do I scale CPU cores for multiple VM clusters?

A: If you have more than one VM cluster, you need to run the `dbaascli cpuscale update` command separately from any node inside each VM cluster that you want to scale up or down.

Example Use Cases

Example 1: Scaling up CPU cores to 20

The VM cluster is running with 16 cores, and you want to increase it to 20.

```
dbaascli cpuscale update --coreCount 20 --message "Scaling up for increased
demand"
```

Scales the CPU core count up to 20, and includes a message "Scaling up for increased demand" for reference.

Example 2: Scaling down CPU cores to 8

The VM cluster is currently using 12 cores, but you want to reduce the count to 8 to save resources.

```
dbaascli cpuscale update --coreCount 8 --message "Reducing CPU for maintenance
period"
```

Reduces the number of CPU cores to 8 and provides a message for future reference about why the scaling operation was performed.

Example 3: Scaling CPU without a message

You need to scale the CPU cores from 32 to 24, but no additional message is necessary.

```
dbaascli cpuscale update --coreCount 24
```

This command scales the core count down to 24 without any message. The operation will be carried out with the default logging of actions.

Example 4: Verifying the status after scaling CPU

After running the scaling command, you want to check if the update was successful.

```
dbaascli cpuscale get_status
```

Checks the status of the current or last scaling request, allowing you to confirm whether the scale-up or scale-down operation was successful.

Example 5: Attempting to scale when VM is already at maximum cores

Copyright © 2008, 2026, Oracle and/or its affiliates.

The VM cluster already has the maximum allowed CPU cores (48), but an attempt is made to scale up.

```
dbaascli cpuscale update --coreCount 50 --message "Attempt to scale beyond limit"
```

Fail since the VM cluster cannot exceed its maximum allowed cores. The status will reflect the failure after attempting to scale to 50 cores.

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Cloud Software Library (CSWLIB) Management

This section provides tools for managing software images in Exadata Database Service on Cloud@Customer environments. Commands like `dbaascli cswlib deleteLocal` allow for the removal of local images, while `dbaascli cswlib download` enables the download of new software images. You can also view available images locally with `dbaascli cswlib listLocal` or check all available Database and Grid Infrastructure images using `dbaascli cswlib showImages`. These commands help manage and maintain software libraries efficiently.

- dbaascli cswlib deleteLocal
  To delete the local image, use the `dbaascli cswlib deleteLocal` command.
- dbaascli cswlib download
  To download available software images and make them available in your Exadata Database Service on Cloud@Customer 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 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-4 dbaascli cswlib deletelocal

```
dbaascli cswlib deletelocal --imagetag 19.15.0.0.0
DBAAS CLI version MAIN
Executing command cswlib deletelocal --imagetag 19.15.0.0.0
Job id: 8b3e71de-4b81-4832-b49c-7f892179bb4f
Log file location: /var/opt/oracle/log/cswLib/deleteLocal/
dbaastools_2022-07-18_10-00-02-AM_73658.log
dbaascli execution completed
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli cswlib download

To download available software images and make them available in your Exadata Database Service on Cloud@Customer environment, use the `dbaascli cswlib download` command.

Prerequisites

Run the command as the `root` user.

Copyright © 2008, 2026, Oracle and/or its affiliates.

To use the utility, you must connect to an Exadata Database Service on Cloud@Customer virtual machine.

See, *Connecting to a Virtual Machine with SSH*.

Syntax

```
dbaascli cswlib download --version | --imageTag
[--product]
```

Where:

- `--version` specifies an Oracle home image version
- `--imageTag` specifies the image tag of the image
- `--product` specifies the image type. Valid values: `database` or `grid`

Frequently Asked Questions

Q: What is the purpose of the dbaascli cswlib download command?

A: The `dbaascli cswlib download` command is used to download available software images and make them available in your Exadata Cloud Infrastructure.

Q: What are the prerequisites for running the dbaascli cswlib download command?

A: You must run the command as the `root` user. Additionally, you need to be connected to an Exadata Cloud Infrastructure virtual machine.

Q: How do I connect to the virtual machine required for this command?

A: You need to use SSH to connect to the Exadata Cloud Infrastructure virtual machine. Detailed instructions can be found in the documentation under "Connecting to a Virtual Machine with SSH."

Q: What does the --version option specify in the command?

A: The `--version` option specifies the Oracle home image version that you want to download.

Q: How do I use the --imageTag option in the dbaascli cswlib download command?

A: The `--imageTag` option is used to specify the image tag of the software image you want to download.

Q: What is the purpose of the --product option in the command?

A: The `--product` option specifies the type of image you want to download. The valid values are database or grid.

Q: Can I download both database and grid images simultaneously?

A: No, you must specify either `database` or `grid` using the `--product` option, so each download operation is specific to one type of image.

Q: What happens if I do not specify a version or image tag?

A: The command will likely fail or prompt you for the required information since the `--version` or `--imageTag` options are necessary to identify the specific software image to download.

Q: Is it necessary to specify both --version and --imageTag together?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-5 dbaascli cswlib download --product --imageTag

```
dbaascli cswlib download --product database --imageTag 19.14.0.0.0
```

Example 7-6 dbaascli cswlib download --version 19.9.0.0.0

```
dbaascli cswlib download --product database --imageTag 19.14.0.0.0
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli cswlib listLocal

To view the list of locally available Database and Grid Infrastructure images, use the `dbaascli cswlib listLocal` command.

Run the command as the `root` user.

Syntax

```
dbaascli cswLib listLocal [--product <value>]
```

Where:

- `--product` identifies Oracle home product type. Valid values: `database` or `grid`.

Frequently Asked Questions

Q: What is the purpose of the dbaascli cswlib listLocal command?

A: The `dbaascli cswlib listLocal` command is used to view the list of locally available Database and Grid Infrastructure images on your system.

Q: What are the prerequisites for running the dbaascli cswlib listLocal command?

A: The command must be run as the `root` user to have the necessary permissions to access and list the available images.

Q: How do I specify which type of images to list using this command?

A: Use the `--product` option to specify the type of Oracle home images you want to list. The valid values are `database` or `grid`.

Q: What does the --product option represent in the dbaascli cswlib listLocal command?

A: The `--product` option identifies the Oracle home product type, allowing you to filter the list of available images to either `database` or `grid` types.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-7 dbaascli cswlib listlocal

```
dbaascli cswlib listlocal
DBAAS CLI version MAIN
Executing command cswlib listlocal
Job id: bc4f047c-0a34-4d4d-a1ea-21ddc2a9c627
Log file location: /var/opt/oracle/log/cswLib/listLocal/
dbaastools_2022-07-18_10-29-53-AM_16077.log
############ List of Available Database Images  #############
1.IMAGE_TAG=12.2.0.1.220419
  IMAGE_SIZE=5GB
  VERSION=12.2.0.1.220419
  DESCRIPTION=12.2 APR 2022 DB Image
2.IMAGE_TAG=18.16.0.0.0
  IMAGE_SIZE=6GB
  VERSION=18.16.0.0.0
  DESCRIPTION=18c OCT 2021 DB Image
3.IMAGE_TAG=19.14.0.0.0
  IMAGE_SIZE=5GB
  VERSION=19.14.0.0.0
  DESCRIPTION=19c JAN 2022 DB Image
dbaascli execution completed
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli cswlib showImages

To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.

Run the command as the `root` user.

Syntax

```
dbaascli cswlib showImages
[--product]
```

Where:

- `--product` identifies Oracle home product type. Valid values: `database` or `grid`.

Frequently Asked Questions

Q: What is the purpose of the dbaascli cswlib showImages command?

A: The `dbaascli cswlib showImages` command is used to view the list of available Database and Grid Infrastructure images that can be downloaded or managed within your Oracle Exadata Database Service environment.

Q: What are the prerequisites for running the dbaascli cswlib showImages command?

A: The command must be run as the `root` user to ensure you have the necessary permissions to view the available images.

Q: How do I filter the images listed by this command?

A: You can filter the images by specifying the `--product` option with either `database` or `grid` to list only the images related to that product type.

Q: What does the --product option represent in the dbaascli cswlib showImages command?

A: The `--product` option identifies the Oracle home product type, allowing you to filter the list of images to either database or grid.

Q: Can I view both database and grid images in a single command execution?

A: No, you need to run the command twice with different `--product` values (`database` and `grid`) to view both types of images.

Q: What happens if I do not specify the --product option in the command?

A: If the `--product` option is not specified, the command may list all available images or it may prompt you to specify the product type, depending on your environment configuration.

Q: When should I use the dbaascli cswlib showImages command?

A: Use this command when you want to view the list of Database or Grid Infrastructure images that are available for download or deployment in your Oracle Exadata Database Service environment.

Q: Is there any difference between dbaascli cswlib showImages and dbaascli cswlib listLocal commands?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-8 dbaascli cswlib showImages

```
dbaascli cswlib showImages
```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Database Management

This section covers comprehensive tasks for managing Oracle Databases. It includes commands for creating (`dbaascli database create`), deleting (`dbaascli database delete`), and upgrading databases (`dbaascli database upgrade`). Other key tasks include adding and deleting instances (`dbaascli database addInstance`, `dbaascli database deleteInstance`), managing backups (`dbaascli database backup`), and handling database recovery (`dbaascli database recover`). You can also modify database parameters, manage pluggable databases, patch databases, and convert non-CDB databases to PDBs. These commands ensure efficient control over the entire database lifecycle.

- dbaascli database addInstance
  To add the database instance on the specified node, use the `dbaascli database addInstance` command.
- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.
- dbaascli database bounce
  To shut down and restart a specified Exadata Database Service on Cloud@Customer 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- dbaascli database createTemplate
  Use this command to create database templates (DBCA templates) that can subsequently be used to create databases.
- dbaascli database delete
  To delete an Oracle Database, use the `dbaascli database delete` command.
- dbaascli database deleteInstance
  To delete the database instance on the specified node, use the `dbaascli database deleteInstance` command.
- dbaascli database duplicate
  To create a database from an active database, use the `dbaascli database duplicate` command.
- dbaascli database getDetails
  This command shows the detailed information of a given database e.g. dbname, node information, pluggable databases information etc.
- dbaascli database getPDBs
  To view the list of all pluggable databases in a container database, use the `dbaascli database getPDBs` command.
- dbaascli database modifyParameters
  To modify or reset initialization parameters for an Oracle Database, use the `dbaascli database modifyParameters` command.
- dbaascli database recover
  To recover a database, use the `dbaascli database recover` command.
- dbaascli database runDatapatch
  To patch an Oracle Database, use the `dbaascli database runDatapatch` command.
- dbaascli database start
  To start an Oracle Database, use the `dbaascli database start` 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 database addInstance

To add the database instance on the specified node, use the `dbaascli database addInstance` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database addInstance --dbname <value> --node <value> [--newNodeSID
<value>]
```

Where:

- `--dbname` specifies Oracle Database name
- `--node` specifies the node name for the database instance

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--newNodeSID` specifies SID for the instance to add in the new node

Frequently Asked Questions

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Q: Will this command cause database downtime?

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli database backup

To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database backup --dbname <value>
        {
            --list
                {
                    [--backupType <value>]
                    | [--json <value>]
                }
            | --start [--level0] [--level1]
                {
                    [--archival --tag <value>]
                    | [--archivelog]
                }
            | --delete --backupTag <value>
            | --status --uuid <value> [--json <value>]
            | --getBackupReport
                {
                    --tag <value>
                    | --latest
                }
                --json <value>
            | --configure
                {
                    --configFile <value>
                    | --enableRTRT
                    | --disableRTRT
                    | --disableCatalog
                    | --deleteImmutableConfiguration
                }
            | --getConfig
                {
                    [--configFile <value>]
                    | [--showOldParams]
                }
            | --validate [--untilTime <value>]
            | --showHistory [--all]
            | --getSchedules
        }
```

Where:

```
--dbname specifies Oracle Database name
--list returns database backup information
           [--backupType | --json]
           [--backupType specifies backupType (REGULAR-L0 | REGULAR-L1 | ARCHIVELOG |
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
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.

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:

Copyright © 2008, 2026, Oracle and/or its affiliates.

```
dbaascli database backup --dbname <value> --start --level0
```

Q: How do I perform an incremental backup using the dbaascli database backup command?

A: To perform a Level-1 incremental backup, use this syntax:

```
dbaascli database backup --dbname <value> --start --level1
```

Q: What is the difference between Level-0 and Level-1 backups?

A: A Level-0 backup is a full backup of the database, while a Level-1 backup is an incremental backup that only captures changes made since the last Level-0 or Level-1 backup.

Q: Can I take an archival backup using this command?

A: Yes, you can create an archival backup by using the `--archival` option along with the `--start` command:

```
dbaascli database backup --dbname <value> --start --archival --tag <backup_tag>
```

Q: How do I delete an existing archival backup?

A: To delete an archival backup, use the following syntax:

```
dbaascli database backup --dbname <value> --delete --backupTag <tag_value>
```

Q: How can I check the status of a specific backup using the command?

A: You can check the status of a backup by using the `--status` option with the `--uuid` parameter, like this:

```
dbaascli database backup --dbname <value> --status --uuid <backup_uuid>
```

Q: How do I list all backups for a database?

A: To list all available backups for a specific database, use the `--list` option:

```
dbaascli database backup --dbname <value> --list
```

For JSON output, add the `--json` option:

```
dbaascli database backup --dbname <value> --list --json <file_name>
```

Q: How can I retrieve a backup report?

A: You can get a backup report using the `--getBackupReport` option, either for a specific tag or for the latest backup:

```
dbaascli database backup --dbname <value> --getBackupReport --tag <backup_tag> --
json <file_name>
```

Or to retrieve the latest report:

```
dbaascli database backup --dbname <value> --getBackupReport --latest --json
<file_name>
```

Q: How do I configure database backup settings?

A: Use the *--configure* option to specify the backup configuration file or to enable/disable Real-Time Redo Transport (RTRT):

```
dbaascli database backup --dbname <value> --configure --configFile <config_file>
```

To enable RTRT:

Copyright © 2008, 2026, Oracle and/or its affiliates.

```
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 7-9 Examples

- To change the archive log retention period follow the below steps:

  ```
  dbaascli database backup --getConfig --dbname <dbname>
  ```

  This will generate a backup config file `.cfg`.

  Update `bkup_archlog_fra_retention` value in this config file.

  Run the configure command:

  ```
  dbaascli database backup --configure --dbname <dbname> --configfile
  <config file generated above>
  ```

- To get backup configuration for a database *myTestDB*:

  ```
  dbaascli database backup --dbName myTestDB --getConfig --configFile /tmp/
  configfile_1.txt
  ```

- To set backup configuration for a database *myTestDB* by modifying the config file with configuration details:

  ```
  dbaascli database backup --dbName myTestDB --configure --configFile /tmp/
  configfile_1_modified.txt
  ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

- 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 Exadata Database Service on Cloud@Customer database, use the `dbaascli database bounce` command.

Prerequisites

Run the command as the `oracle` user.

Syntax

```
dbaascli database bounce
[--dbname][--rolling <value>]
```

Where:

- `--dbname` specifies the name of the database
- `--rolling` specifies `true` or `false` to bounce the database in a rolling manner. Default value is `false`.

The command performs a database shutdown in immediate mode. The database is then restarted and opened. In Oracle Database 12c or later, all of the PDBs are also opened.

Frequently Asked Questions

Q: What is the purpose of the dbaascli database bounce command?

A: The `dbaascli database bounce` command is used to shut down and restart an Oracle Database in Exadata Cloud Infrastructure. It supports restarting the database in a rolling manner, ensuring minimal disruption.

Q: What are the prerequisites for running the dbaascli database bounce command?

A: The command must be run as the `oracle` user, which has the required privileges to shut down and restart the database.

Q: What does the --dbname option specify in this command?

A: The `--dbname` option specifies the name of the Oracle Database that you want to shut down and restart.

Q: What is the --rolling option used for in the dbaascli database bounce command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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?

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 7-10 dbaascli database bounce

```
dbaascli database bounce --dbname dbname
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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?

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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?

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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-11 dbaascli database changePassword

```
dbaascli database changepassword --dbname db19
```

dbaascli database convertToPDB

To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.

Syntax

```
dbaascli database convertToPDB --dbname <value> [--cdbName <value>] [--
executePrereqs]
        {
            [--copyDatafiles]
            | [--backupPrepared]
        }
        [--targetPDBName <value>] [--waitForCompletion <value>] [--resume [--
sessionID <value>]]
```

Where:

- `--dbname` specifies the name of Oracle Database
- `--cdbName` specifies the name of the target CDB in which the PDB will be created. If the CDB does not exist, then it will be created in the same Oracle home as the source non-CDB
- `--executePrereqs` specifies to run only the pre-conversion checks
- `--copyDatafiles` specifies to create a new copy of the data files instead of using the ones from the source database
- `--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

Frequently Asked Questions

Q: What does the dbaascli database convertToPDB command do?

A: The `dbaascli database convertToPDB` command converts a specified non-CDB Oracle Database into a Pluggable Database (PDB) within a Container Database (CDB).

Q: How do I specify which database to convert?

A: Use the `--dbname` option to specify the name of the non-CDB Oracle Database that you want to convert. For example:

```
dbaascli database convertToPDB --dbname <db_name>
```

Q: How do I specify the target CDB for the PDB conversion?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Use the `--cdbName` option to specify the name of the target Container Database (CDB) in which the new PDB will be created. If the CDB does not exist, a new one will be created in the same Oracle home as the source non-CDB.

Q: What does the --executePrereqs option do?

A: The `--executePrereqs` option runs pre-conversion checks to ensure that the database is ready for conversion. No changes will be made to the database during this step.

Q: How do I copy the data files during the conversion?

A: You can use the `--copyDatafiles` option to create a new copy of the data files rather than using the original files from the source database.

Q: What is the purpose of the --keepSourceDB option?

A: The `--keepSourceDB` option allows you to preserve the original source non-CDB after the conversion operation is completed. If you do not use this option, the source database will be removed after conversion.

Q: How do I confirm that a backup is prepared before conversion?

A: Use the `--backupPrepared` flag to acknowledge that you have taken a proper backup of the non-CDB database before performing the conversion. This step is crucial to avoid data loss.

Q: How do I specify a custom name for the new PDB?

A: Use the `--targetPDBName` option to provide a specific name for the new PDB that will be created as part of the conversion. For example:

```
dbaascli database convertToPDB --dbname <db_name> --targetPDBName <pdb_name>
```

Q: Can I run the conversion in the background?

A: Yes, you can use the `--waitForCompletion` option to specify whether the operation should be run in the background. Use `false` to run in the background and `true` to wait for the operation to complete before proceeding. The default is `true`.

Q: How do I resume a previously interrupted conversion?

A: You can use the `--resume` option to resume a previously interrupted conversion process. Optionally, you can provide the `--sessionID` to resume a specific session.

Q: What happens if I don't specify a CDB name?

A: If the `--cdbName` option is not provided, the system will create the new PDB in the same Oracle home as the source non-CDB database.

Q: Can I continue the conversion after an interruption without knowing the session ID?

A: Yes, using the `--resume` option without specifying a session ID will resume the last known session. If you want to resume a specific session, you can provide the `--sessionID`.

Q: What does the --sessionID option do?

A: The `--sessionID` option is used in conjunction with `--resume` to specify a particular session to resume, in case there were multiple interrupted sessions.

Q: Is it mandatory to have a backup before converting a non-CDB to a PDB?

A: While the `--backupPrepared` flag is optional, it is highly recommended to take a backup of the non-CDB before performing the conversion to PDB. This ensures that you can restore the database in case of any issues during the conversion.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-12 dbaascli database convertToPDB

To run pre-conversion prechecks:

```
dbaascli database convertToPDB --dbname ndb19 --cdbname cdb19 --
backupPrepared --executePrereqs
```

To run a full conversion with a copy of the data files from the non-CDB:

```
dbaascli database convertToPDB --dbname tst19 --cdbname cdb19 --copyDatafiles
```

dbaascli database create

To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Use this command to create Oracle Database version 12.1.0.2 or higher with the release update JAN 2021 or higher. For databases with lower versions, it is recommended to use the OCI Console based API.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database create --dbName {--oracleHome | --oracleHomeName}
[--dbUniqueName <value>]
[--dbSID <value>]
[--createAsCDB <value>]
[--pdbName <value>]
[--pdbAdminUserName <value>]
[--dbCharset <value>]
[--dbNCharset <value>]
[--dbLanguage <value>]
[--dbTerritory <value>]
[--sgaSizeInMB <value>]
[--pgaSizeInMB <value>]
[--datafileDestination <value>]
[--fraDestination <value>]
[--fraSizeInMB <value>]
[--nodeList <value>]
[--tdeConfigMethod <value>]
[--kmsKeyOCID <value>]
{
            [--resume [--sessionID <value>]]
            | [--revert [--sessionID <value>]]
        }
[--executePrereqs]
[--honorNodeNumberForInstance <value>]
[--lockPDBAdminAccount <value>]
[--dbcaTemplateFilePath <value>]
[--waitForCompletion]
```

Where:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--dbname` specifies the name of the database
- `--oracleHome` specifies the location of the Oracle home
- `--oracleHomeName` specifies the name of the Oracle home
- `--dbUniqueName` specifies database unique name
- `--dbSID` specifies the SID of the database
- `--createAsCDB` specifies `true` or `false` to create database as CDB or Non-CDB
- `--pdbName` specifies the name of the PDB
- `--pdbAdminUserName` specify PDB admin user name
- `--dbCharset` specifies database character set
- `--dbNCharset` specifies database national character set
- `--dbLanguage` specifies the database language
- `--dbTerritory` specifies the database territory
- `--sgaSizeInMB` specifies the `sga_target` value in megabyte unit
- `--pgaSizeInMB` specifies the `pga_aggregate_target` value in megabyte unit
- `--datafileDestination` specifies the ASM disk group name to use for database datafiles
- `--fraDestination` specifies ASM disk group name to use for database Fast Recovery Area
- `--fraSizeInMB` specifies the Fast Recovery Area size value in megabyte unit
- `--nodeList` specifies a comma-delimited list of nodes for the database
- `--tdeConfigMethod` specifies TDE configuration method. Valid values: `FILE`, `KMS`
- `--kmsKeyOCID` specifies KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE
- `--resume` resumes the previous execution
- `--revert` rolls back the previous run
- `--sessionID` resumes or reverts to a specific session ID.
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--honorNodeNumberForInstance` specifies `true` or `false` to indicate instance name to be suffixed with the cluster node numbers. Default value: `true`
- `--lockPDBAdminAccount` specifies `true` or `false` to lock the PDB admin user account. Default value is `true`
- `--dbcaTemplateFilePath` specifies the absolute path of the dbca template name to create the database.
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`

Frequently Asked Questions

Q: What does the dbaascli database create command do?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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?

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Use the `--nodeList` option to specify a comma-separated list of nodes where the database should be created.

Q: How can I resume or revert a previous database creation attempt?

A: Use the `--resume` option to resume the previous execution or the `--revert` option to roll back the previous run. You can also specify a `--sessionID` to resume or revert a specific session.

Q: What does the --executePrereqs option do?

A: The `--executePrereqs` option only runs the prerequisites for the database creation operation, without actually creating the database. Use `yes` or `no` to enable or disable this option.

Q: Can I specify a custom DBCA template for the database creation?

A: Yes, use the `--dbcaTemplateFilePath` option to provide the absolute path of the DBCA template file that should be used to create the database.

Q: Can I run the database creation operation in the background?

A: Yes, you can use the `--waitForCompletion` option to specify whether the command should wait for the database creation to complete (`true`) or run the operation in the background (`false`).

Q: What happens if I don't specify the --dbUniqueName option?

A: If you don't specify a unique name for the database using `--dbUniqueName`, the system will automatically generate one based on the provided `--dbName`.

Q: Can I lock the PDB admin account during the creation of a CDB?

A: Yes, you can use the `--lockPDBAdminAccount` option and set it to true to lock the PDB admin account after database creation. By default, this value is set to `true`.

Example 7-13 dbaascli database create

```
dbaascli database create --dbName db19 --oracleHomeName myhome19 --dbSid
db19sid --nodeList node1,node2 --createAsCDB true
```

dbaascli database createTemplate

Use this command to create database templates (DBCA templates) that can subsequently be used to create databases.

Run the command as the `root` or `oracle` user.

Syntax

Create a new DBCA template from the specified database.

```
dbaascli database createTemplate --dbname <value>
 {
   --templateLocation <value> | --uploadToObjectStorage --
objectStorageLoginUser <value> --objectStorageBucketName <value> [--
objectStorageUrl <value>]
 }
 [--templateName <value>] [--rmanParallelism <value>]
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Where:

- `--dbname` specifies the name of the database
- `--templateLocation` specifies the template name
- `--uploadToObjectStorage`: specifies to upload the template to Object Storage
  - `--objectStorageLoginUser`: specifies the Object Storage login user
  - `--objectStorageBucketName`: specifies the Object Storage bucket name
  - `--objectStorageUrl`: specifies the Object Storage URL
- `--templateName`: specifies the name of the template
- `--rmanParallelism` specifies the parallelsim value

Frequently Asked Questions

Q: What is the purpose of the dbaascli database createTemplate command?

A: The `dbaascli database createTemplate` command is used to create database templates (DBCA templates) from a specified database, which can later be used to create new databases.

Q: How do I specify the name of the database for which I want to create a template?

A: Use the `--dbname` option to specify the name of the Oracle Database from which the template will be created. For example:

```
dbaascli database createTemplate --dbname <db_name>
```

Q: How do I specify where the template should be saved?

A: Use the `--templateLocation` option to specify the location where the DBCA template will be saved. For example:

```
dbaascli database createTemplate --dbname <db_name> --templateLocation /path/to/
template
```

Q: Can I upload the template to Object Storage directly?

A: Yes, you can use the `--uploadToObjectStorage` option to upload the DBCA template to Object Storage. You will need to specify the Object Storage login user and bucket name with the `--objectStorageLoginUser` and `--objectStorageBucketName` options, respectively.

Q: How do I specify the Object Storage login details when uploading the template?

A: Use the following options to specify Object Storage details:

`--objectStorageLoginUser`: Specifies the Object Storage login user.

`--objectStorageBucketName`: Specifies the Object Storage bucket name.

`--objectStorageUrl`: (Optional) Specifies the URL of the Object Storage if it is different from the default one.

For example:

```
dbaascli database createTemplate --dbname <db_name> --uploadToObjectStorage --
objectStorageLoginUser <user> --objectStorageBucketName <bucket_name>
```

Q: How do I specify a custom name for the DBCA template?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Use the `--templateName` option to specify a custom name for the DBCA template. For example:

```
dbaascli database createTemplate --dbname <db_name> --templateName
<template_name>
```

Q: What is the --rmanParallelism option used for?

A: The `--rmanParallelism` option specifies the level of parallelism for RMAN operations during the template creation process. For example:

```
dbaascli database createTemplate --dbname <db_name> --rmanParallelism 4
```

Q: What happens if I don't specify the --templateLocation or --uploadToObjectStorage options?

A: If you do not specify a template location using `--templateLocation` or choose to upload to Object Storage using `--uploadToObjectStorage`, the command will not know where to store the created template, and it will fail to complete.

Q: Can I use both --templateLocation and --uploadToObjectStorage at the same time?

A: No, you should choose either `--templateLocation` to save the template locally or `--uploadToObjectStorage` to upload it to Object Storage, but not both.

Q: Is the --objectStorageUrl option mandatory when uploading the template to Object Storage?

A: No, the `--objectStorageUrl` option is optional. If not specified, the default Object Storage URL will be used. You only need to specify this if using a custom Object Storage URL.

Q: What user privileges are required to run the dbaascli database createTemplate command?

A: The command must be run either as the `root` or `oracle` user.

Q: Can I resume a previously failed template creation process?

A: No, the `dbaascli database createTemplate` command does not support resuming a failed process. You will need to restart the command from the beginning.

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>]]
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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?

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
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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?

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 7-14 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

Copyright © 2008, 2026, Oracle and/or its affiliates.

Frequently Asked Questions

Q: What is the purpose of the dbaascli database deleteInstance command?

A: The `dbaascli database deleteInstance` command is used to delete a specific Oracle Database instance on a specified node in an Exadata Cloud Infrastructure environment.

Q: How do I specify which Oracle Database instance to delete?

A: You can specify the Oracle Database instance to delete by using the `--dbname` option to provide the database name and the `--node` option to provide the node name. For example:

```
dbaascli database deleteInstance --dbname <db_name> --node <node_name>
```

Q: Can I delete the instance even if the node is unreachable?

A: Yes, you can use the `--continueOnUnreachableNode` option to proceed with the deletion, even if the specified node is unreachable. For example:

```
dbaascli database deleteInstance --dbname <db_name> --node <node_name> --
continueOnUnreachableNode
```

Q: What happens if the node specified is unreachable during the delete instance operation?

A: If the node is unreachable and the `--continueOnUnreachableNode` option is not used, the operation will fail. If the option is used, the operation will continue even if the node cannot be accessed.

Q: How do I delete a database instance from a specific node?

A: Use the following command to delete a database instance from a specific node:

```
dbaascli database deleteInstance --dbname <db_name> --node <node_name>
```

Q: What user privileges are required to run the dbaascli database deleteInstance command?

A: The command must be run as the `root` user.

Q: Can I delete an instance without specifying the node?

A: No, the `--node` option is required to specify which node the database instance should be deleted from.

Q: What does the --continueOnUnreachableNode option do?

A: The `--continueOnUnreachableNode` option allows the operation to proceed even if the specified node cannot be reached, ensuring that the instance deletion continues in scenarios where the node might be down.

Q: Is it possible to delete multiple database instances at once using this command?

A: No, the `dbaascli database deleteInstance` command is used to delete a single database instance on a specified node at a time. You would need to run the command separately for each instance you want to delete.

Example 7-15 database deleteinstance

```
database deleteinstance --node test-node
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli database duplicate

To create a database from an active database, use the `dbaascli database duplicate` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database duplicate --dbName <value> --sourceDBConnectionString
<value>
        {
            --oracleHome <value>
            | --oracleHomeName <value>
        }
[--dbSID <value>]
[--dbUniqueName <value>]
[--sgaSizeInMB <value>]
[--pgaSizeInMB <value>]
[--datafileDestination <value>]
[--fraDestination <value>]
[--fraSizeInMB <value>]
[--sourceDBWalletLocation <value>]
[--nodeList <value>]
        {
            [--resume [--sessionID <value>]]
            | [--revert [--sessionID <value>]]
        }
[--rmanParallelism <value>]
[--rmanSectionSizeInGB <value>]
[--tdeConfigMethod <value>]
[--kmsKeyOCID <value>]
[--sourceDBTdeConfigMethod <value>]
[--sourceDBKmsKeyOCID <value>]
[--executePrereqs <value>]
[--waitForCompletion <value>]
[--skipPDBs <value>]
```

Where:

- `--dbName` specifies Oracle Database name
- `--sourceDBConnectionString` specifies source database connection string in the format of *<scan_name>:<scan_port>/<database_service_name>*
- `--oracleHome` specifies Oracle home location
- `--oracleHomeName` specifies Oracle home name
- `--dbSID` specifies database SID
- `--dbUniqueName` specifies database unique name
- `--sgaSizeInMB` specifies `sga_target` value in mega byte unit
- `--pgaSizeInMB` specifies `pga_aggregate_target` value in mega byte unit

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--datafileDestination` specifies ASM disk group name to use for database datafiles
- `--fraDestination` specifies ASM disk group name to use for database fast recovery area
- `--fraSizeInMB` specifies fast recovery area size value in mega byte unit
- `--sourceDBWalletLocation` specifies source database TDE wallet file location. This is required to duplicate database from active database
- `--nodeList` specifies a comma-delimited list of nodes for the database
- `--resume` specifies to resume the previous execution
  - `--sessionID` specifies to resume a specific session ID
- `--revert` specifies to rollback the previous execution
  - `--sessionID` specifies to rollback a specific session ID
- `--rmanParallelism` specifies parallelsim value
- `--rmanSectionSizeInGB` specifies RMAN section size in GB
- `--tdeConfigMethod` specifies TDE configuration method. Allowed values are `FILE` and `KMS`.
- `--kmsKeyOCID` specifies KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE.
- `--sourceDBTdeConfigMethod` specifies source database TDE configuration method. Allowed values are `FILE` and `KMS`.
- `--sourceDBKmsKeyOCID` specifies source database KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE.
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes|no`
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true|false`
- `--skipPDBs` specifies a comma-delimited list of source database PDB names, which needs to be excluded for the duplicate database operation. Example: pdb1,pdb2...

Frequently Asked Questions

Q: What is the purpose of the 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>
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

or

`--oracleHomeName <value>`

Q: What is the purpose of the --sourceDBWalletLocation option?

A: The `--sourceDBWalletLocation` option specifies the location of the source database TDE wallet file, which is required to duplicate the database from an active source database.

Q: Can I skip duplicating specific PDBs from the source database?

A: Yes, you can use the `--skipPDBs` option to specify a comma-delimited list of PDB names that should be excluded from the duplicate operation. For example:

`--skipPDBs pdb1,pdb2`

Q: How do I configure TDE for the new database?

A: Use the `--tdeConfigMethod` option to specify the TDE configuration method (`FILE` or `KMS`). If you choose `KMS`, you can provide the `KMS` key OCID using the `--kmsKeyOCID` option. For example:

`--tdeConfigMethod FILE`

or

`--tdeConfigMethod KMS --kmsKeyOCID <value>`

Q: What does the --executePrereqs option do?

A: The `--executePrereqs` option specifies whether to run only the prerequisite checks for the operation. Valid values are `yes` to run prereqs only, or `no` to proceed with the full operation.

Q: How can I resume a previously interrupted duplicate operation?

A: Use the `--resume` option along with the `--sessionID` option to resume a previously interrupted duplicate operation. For example:

`--resume --sessionID <value>`

Q: What does the --waitForCompletion option do?

A: The `--waitForCompletion` option specifies whether to wait for the operation to complete. Setting this to `true` waits for completion, while `false` runs the operation in the background. For example:

`--waitForCompletion true`

Q: What is the purpose of the --rmanParallelism option?

A: The `--rmanParallelism` option specifies the parallelism value for RMAN (Recovery Manager) during the duplication process. This can improve the speed of the duplication operation by using multiple parallel processes.

Q: How do I specify the size of the SGA and PGA for the new database?

A: Use the `--sgaSizeInMB` and `--pgaSizeInMB` options to specify the sizes of the SGA (System Global Area) and PGA (Program Global Area) in megabytes, respectively. For example:

`--sgaSizeInMB <value>`

`--pgaSizeInMB <value>`

Q: What does the --revert option do?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-16 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>
```

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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Syntax

```
dbaascli database getPDBs --dbname <value>
```

Where:

- `--dbname` specifies the name of the container database

Frequently Asked Questions

Q: What is the purpose of the dbaascli database getPDBs command?

A: The `dbaascli database getPDBs` command is used to list all the pluggable databases (PDBs) within a specified container database (CDB).

Q: How do I specify the container database for the getPDBs command?

A: You use the `--dbname` option to specify the name of the container database. For example:

```
--dbname <value>
```

Q: Do I need to run the dbaascli database getPDBs command as a specific user?

A: Yes, you must run the command as either the `root` user or the `oracle` user.

Q: Can I view PDBs in a non-CDB database using the getPDBs command?

A: No, the getPDBs command is only applicable to container databases (CDBs). You cannot use this command for non-CDB databases.

Q: What is the format of the output from the dbaascli database getPDBs command?

A: The command returns a list of all PDBs within the specified container database. The output typically includes PDB names, statuses, and other relevant details about each pluggable database.

Q: Can this command be used for multiple databases at once?

A: No, the `dbaascli database getPDBs` command works with a single container database at a time, specified by the `--dbname` option.

Q: Is it necessary to shut down the database to use the getPDBs command?

A: No, the `getPDBs` command does not require the database to be shut down. It can be run while the container database is operational.

Example 7-17 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Syntax

```
dbaascli database modifyParameters --dbname <value>
{
--setParameters <values>[--instance <value>] [--backupPrepared] [--
allowBounce]|
--resetParameters <values> [--instance <value>] [--backupPrepared] [--
allowBounce]
}
--responseFile
[--backupPrepared]
[--instance]
[--allowBounce]
[--waitForCompletion]
```

Where:

- `--dbname` specifies the name of the database.
- `--setParameters` specifies a comma-delimited list of parameters to modify with new values. For example: `parameter1=valueA,parameter2=valueB`, and so on. For blank values use parameter1=valueA,parameter2='',etc.
- `--resetParameters` specifies a comma-delimited list of parameters to be reset to their corresponding default values. For example, `parameter1,parameter2`, and so on.
- `--instance` specifies the name of the instance on which the parameters will be processed. If not specified, then the operation will be performed at the database level.
- `--backupPrepared` acknowledges that a proper database backup is in place prior to modifying critical or sensitive parameters.
- `--allowBounce` grants permission to bounce the database in order to reflect the changes on applicable static parameters.
- `--waitForCompletion` specify false to run the operation in background. Valid values : true| false.]

Frequently Asked Questions

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

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

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

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

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

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

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

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

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

Q: Can I modify parameters using a response file?

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

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

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

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

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

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

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

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

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

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

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

Example 7-18 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>
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
                    | --untilSCN <value>
                    | --latest
                }
            | --status --uuid <value>
        }
```

Where:

```
--dbname: Oracle Database name.
      --start | --status
--start: Begins database recovery.
      --untilTime | --untilSCN | --latest
      --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.
--status
      --uuid <value>
```

Frequently Asked Questions

Q: What is the purpose of the dbaascli database recover command?

A: The `dbaascli database recover` command is used to recover an Oracle Database from backups stored in a backup storage destination.

Q: How do I specify which database to recover?

A: You can specify the database you want to recover using the `--dbname` option followed by the database name. For example:

`--dbname <database_name>`

Q: What are the recovery options available with the dbaascli database recover command?

A: The recovery options are:

`--untilTime`: Recover the database to a specific time.

`--untilSCN`: Recover the database to a specific System Change Number (SCN).

`--latest`: Recover the database to the last known state.

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:

Copyright © 2008, 2026, Oracle and/or its affiliates.

`--latest`

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`).

Q: Is there a way to recover the database without specifying a time or SCN?

A: Yes, you can recover the database to its last known state using the `--latest` option, which doesn't require specifying a time or SCN.

Q: Can I perform a partial recovery?

A: Yes, you can recover the database to a specific point in time or SCN using the `--untilTime` or `--untilSCN` options, respectively.

Example 7-19 Examples

- To recover the database *myTestDb* to latest:

  ```
  dbaascli database recover --dbname myTestDb --start --latest
  ```

- To query the status of recovery request submitted with `uuid` *2508ea18be2911eb82d0020017075151*:

  ```
  dbaascli database recover --dbname myTestDb --status --uuid
  2508ea18be2911eb82d0020017075151
  ```

dbaascli database runDatapatch

To patch an Oracle Database, use the `dbaascli database runDatapatch` command.

Prerequisites

- Before performing a `runDatapatch` operation, ensure that all of the database instances associated with the database are up and running.
- Run the command as the `root` user.

Syntax

```
dbaascli database runDatapatch --dbname
[--resume]
    [--sessionID]
[--skipPdbs | --pdbs]
[--executePrereqs]
[--patchList]
[--skipClosedPdbs]
[--rollback]
```

Where:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--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?

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:

Copyright © 2008, 2026, Oracle and/or its affiliates.

`--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:

```
--sessionID 67890

dbaascli database runDatapatch --dbname db19
```

dbaascli database start

To start an Oracle Database, use the `dbaascli database start` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli database start
[--dbname]
[--mode]
```

Where:

- `--dbname` specifies the name of the database
- `--mode` specifies mount or nomount to start database in the corresponding mode

The command starts and opens the database. In Oracle Database 12c or later, all of the PDBs are also opened.

Frequently Asked Questions

Q: What is the purpose of the dbaascli database start command?

A: The `dbaascli database start` command is used to start an Oracle Database.

Q: What must be done before running the dbaascli database start command?

A: The command must be run as the `root` user.

Q: How do I specify the database I want to start?

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

`--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 7-20 dbaascli database start

```
dbaascli database start --dbname dbname --mode mount
```

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`

Copyright © 2008, 2026, Oracle and/or its affiliates.

The command performs a database shutdown in immediate mode. No new connections or new transactions are permitted. Active transactions are rolled back, and all connected users are disconnected.

Frequently Asked Questions

Q: What is the purpose of the dbaascli database stop command?

A: The `dbaascli database stop` command is used to stop an Oracle Database.

Q: What are the prerequisites for using the dbaascli database stop command?

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

Q: How do I specify which database to stop?

A: You can specify the database by using the `--dbname` option followed by the name of the database. For example:

```
--dbname myDatabase
```

Q: What are the valid shutdown modes for the dbaascli database stop command?

A: The valid shutdown modes are:

```
abort
```

```
immediate
```

```
normal
```

```
transactional
```

Q: What is the default shutdown mode if no mode is specified?

A: If no mode is specified, the database will be shut down in `immediate` mode by default.

Q: What happens in immediate shutdown mode?

A: In `immediate` mode, no new connections or transactions are permitted, active transactions are rolled back, and all connected users are disconnected.

Q: How can I stop the database in abort mode?

A: To stop the database in abort mode, use the `--mode` option with abort. For example:

```
--mode abort
```

Q: What does normal mode do when stopping the database?

A: In normal mode, the database allows current user sessions to complete and then stops without affecting active transactions.

Q: What is transactional mode used for in the 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 © 2008, 2026, Oracle and/or its affiliates.

Example 7-21 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 © 2008, 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 © 2008, 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 7-22 dbaascli database upgrade pre-upgrade requisite checks

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

Data Guard Management

This section focuses on managing Oracle Data Guard configurations and operations. It includes commands like `dbaascli dataguard prepareStandbyBlob` to generate a blob file for setting up a standby site, and `dbaascli dataguard updateDGConfigAttributes` to update Data Guard automation attributes across all cluster nodes. These commands streamline the setup and maintenance of Data Guard environments for high availability and disaster recovery.

- 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 dataguard failover
  To perform a manual failover to the standby database, use the `dataguard failover` command.
- dbaascli dataguard reinstate
  To reinstate a failed database as a standby database after a failover, use the `dataguard reinstate` command.
- dbaascli dataguard switchover
  To perform a switchover to the standby database, use the `dataguard switchover` command.
- dbaascli dataguard prepareForStandby
  To create an Oracle Standby Database, use the `dbaascli dataguard prepareForStandby` command as the **first step**.
- dbaascli dataguard configureStandby
  To create a new Standby Database, use the `dbaascli dataguard configureStandby` command as a second step after the `prepareForStandby` step.
- dbaascli dataguard registerStandby
  To register a newly created standby database with all existing standby databases and to primary database, use the `dbaascli dataguard registerStandby` command as a third step after `configureStandby` step.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- dbaascli dataguard deregisterStandby
  During standby deletion, run the `dbaascli dataguard deregisterStandby` command before deleting the database on the standby cluster to deregister the standby database from the Oracle Data Guard Broker configuration.
- dbaascli dataguard configureAWR
  To enable or disable Automatic Workload Repository (AWR) configuration on your Active Data Guard standby, use the `dbaascli dataguard configureAWR` command.
- dbaascli dataguard updateConfiguration
  To update the transport mode or protection mode or both the parameters of a Data Guard environment, use the `dbaascli dataguard updateConfiguration` command.

dbaascli 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:

```
--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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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.

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
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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 dataguard failover

To perform a manual failover to the standby database, use the `dataguard failover` command.

Run this command as the `oracle` user on the target standby database.

Syntax

```
dbaascli dataguard failover --dbname <value> [--useImmediateFailover] [--
executePrereqs] [--waitForCompletion <value>] [--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--useImmediateFailover` use this flag when the Oracle Data Guard configuration in a warning or error state.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- Performing a Manual Failover Operation Using the dbaascli Utility
  To perform a manual failover to the standby database, use the `dataguard failover` command.

Performing a Manual Failover Operation Using the dbaascli Utility

To perform a manual failover to the standby database, use the `dataguard failover` command.

1. Connect to the virtual machine in the Oracle Data Guard configuration that will host the new primary database as the `opc` user.

   For detailed instructions, see Connecting to a Virtual Machine with SSH.
2. Start a root-user command shell and then switch to the `oracle` user:

   ```
   $ sudo -s
   # su - oracle
   $
   ```
3. Initiate the failover to the standby database.

   ```
   $ dbaascli dataguard failover --dbname <db-name>
   ```
4. Return to being the `root` user.

   ```
   $ exit
   #
   ```
5. Exit the root-user command shell and disconnect from the virtual machine.

   ```
   # exit
   $ exit
   ```

dbaascli dataguard reinstate

To reinstate a failed database as a standby database after a failover, use the `dataguard reinstate` command.

Run this command as the `oracle` user on where reinstate is required (that is failed standby database).

Syntax

```
dbaascli dataguard reinstate --dbname <value> [--primaryDBUniqueName <value>]
[--executePrereqs] [--waitForCompletion <value>] [--resume [--sessionID
<value>]]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--primaryDBUniqueName` specifies database unique name of the current primary database in the Oracle Data Guard setup.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--executePrereqs` runs the prerequisite checks and report the results.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.

To determine when a member should be reinstated in a Data Guard (DG) configuration:

Monitor the `dgmgrl show database` output for the following ORA errors:

- On the new primary cluster:
  ORA-16661: The standby database must be reinstated
- On the old primary cluster:
  ORA-16623: Member detected role change

These messages indicate that a failover has occurred. To restore full synchronization within the Data Guard configuration, the former primary must be reinstated.

- Reinstating a Failed Primary Database Using the dbaascli Utility
  To reinstate a failed primary database after a failover, use the `dataguard reinstate` command.

Reinstating a Failed Primary Database Using the dbaascli Utility

To reinstate a failed primary database after a failover, use the `dataguard reinstate` command.

1. Connect to one of the virtual machines in the Oracle Data Guard configuration as the `oracle` user.

   For detailed instructions, see Connecting to a Virtual Machine with SSH.
2. Initiate reinstatement of the failed primary database.

   ```
   $ dbaascli dataguard reinstate --dbname <db-name>
   ```

   In a multiple standby setup (Data Guard group), it is recommended to specify the `--primaryDBUniqueName` argument.

   ```
   dbaascli dataguard reinstate --dbname <db-name> --primaryDBUniqueName
   <primary-DB-unique-name>
   ```
3. Disconnect from the virtual machine.

   ```
   $ exit
   ```

dbaascli dataguard switchover

To perform a switchover to the standby database, use the `dataguard switchover` command.

Run this command as the `oracle` user.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Syntax

```
dbaascli dataguard switchover --dbname <value> [--targetStandbyDBUniqueName
<value>] [--executePrereqs] [--enableDGDebug] [--waitForCompletion <value>]
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--targetStandbyDBUniqueName` specifies the standby database unique name to change the role from standby to primary database.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--enableDGDebug` enables the traces while performing the operation.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- Performing a Switchover Operation Using the dbaascli Utility
  To perform a switchover to the standby database, use the `dataguard switchover` command.

Performing a Switchover Operation Using the dbaascli Utility

To perform a switchover to the standby database, use the `dataguard switchover` command.

1. Connect to the virtual machine in the Oracle Data Guard configuration that will host the new primary database as the `opc` user.

   For detailed instructions, see Connecting to a Virtual Machine with SSH.

2. Start a root-user command shell and then switch to the `oracle` user.

   ```
   $ sudo -s
   # su - oracle
   $
   ```

3. Initiate the switchover to the standby database.

   ```
   $ dbaascli dataguard switchover --dbname <db-name>
   ```

   In a multiple standby setup (Data Guard group), it is recommended to specify the *--targetStandbyDBUniqueName* argument.

   ```
   dbaascli dataguard switchover --dbname <db-name> --
   targetStandbyDBUniqueName <target-standby-DB-unique-name>
   ```

4. Return to being the `root` user.

   ```
   $ exit
   #
   ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

5. Exit the root-user command shell and disconnect from the virtual machine.

```
# exit
$ exit
```

dbaascli dataguard prepareForStandby

To create an Oracle Standby Database, use the `dbaascli dataguard prepareForStandby` command as the **first step**.

Run this command as the `root` user on the primary database. At the end of the command run, a standby BLOB file is created. You must copy this file to the standby database system to proceed with the `configureStandby` step.

> **Note**
>
> For Disaster Recovery (DR) configurations on Exadata Cloud@Customer (ExaDB-C@C), you must use the Oracle Cloud Infrastructure (OCI) Console or the OCI SDK to set up Data Guard. The `dbaascli` utility is not supported for this use case and should not be used.

Syntax

```
dbaascli dataguard prepareForStandby --dbname <value> --standbyDBUniqueName
<value>  --standbyDBDomain | --noDBDomain  --standbyScanIPAddresses <Standby
SCAN IP Addresses> [ --standbyScanPort ] [ --standbyServiceName ] [  --
primaryScanIPAddresses ] [ --primaryScanPort ] [--executePrereqs] [--resume
[--sessionID <value>]] [--revert [--sessionID <value>]] [--waitForCompletion]
[--skipDRConfiguration]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--standbyDBUniqueName` specifies the standby database unique name for which the primary database will be configured.
- `--standbyDBDomain` specifies the standby database domain for which the primary database will be configured.
- `--noDBDomain` specifies not to use the database domain name for standby database.
- `--standbyScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the standby database SCAN listener, or the SCAN name of the standby database.
- `--standbyScanPort` specifies the corresponding SCAN port number of the standby database.
- `--standbyServiceName` specifies the name of the standby database service for which the primary database will be configured.
- `--primaryScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the primary database SCAN listener, or the SCAN name of the primary database.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--primaryScanPort` specifies the corresponding SCAN port number of the primary database.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--revert` rolls back the previous operation.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- `--skipDRConfiguration` specifies whether to skip Disaster Recovery (DR) configuration as part of the standby database setup. Valid values: `true` (skip DR configuration) or `false` (configure DR).
- Performing PrepareForStandby Operation Using the dbaascli Utility
  To prepare the primary database for creating a new standby database, use the `dbaascli dataguard prepareForStandby` command.

Performing PrepareForStandby Operation Using the dbaascli Utility

To prepare the primary database for creating a new standby database, use the `dbaascli dataguard prepareForStandby` command.

1. Connect to the Virtual Machine where you want to host the primary database as the `opc` user.

   For detailed instructions, see Connecting to a Virtual Machine with SSH.
2. Start a root-user command shell.

   ```
   $ sudo -s
   ```
3. Run the `prepareForStandby` command. Enter the `SYS` password when prompted.

   ```
   dbaascli dataguard prepareForStandby --dbName <db name> --
   standbyDBUniqueName <standby db unique name> --standbyDBDomain <standby db
   domain> --standbyScanIPAddresses  <comma-delimieted list of standby SCAN
   IP addresses> --primaryScanIPAddresses <comma-delimited list of primary
   SCAN IP addresses> --standbyScanPort <standby SCAN listener port>
   ```

   Upon completion, the command displays the location where the standby BLOB file is created.
4. Exit the root-user command shell and disconnect from the virtual machine.

   ```
   #  exit
   ```

dbaascli dataguard configureStandby

To create a new Standby Database, use the `dbaascli dataguard configureStandby` command as a second step after the `prepareForStandby` step.

Run this as the `root` user on the standby cluster.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Syntax

```
dbaascli dataguard configureStandby --dbname <value>  --oracleHome <value> |
--oracleHomeName <value> --standbyDBUniqueName <value> [--standbyDBDomain
<value>] | [--noDBDomain] --primaryScanIPAddresses <value> --primaryScanPort
<value> --primaryServiceName <value> --protectionMode <value> --transportType
<value> --activeDG <value> [--standbyBlobFromPrimary <value>] | [--
standbyDBInfoJsonLocation <value>] [--standbyScanIPAddresses <value>] [--
standbyScanPort <value>] [--standbySID <value>] [--nodeList <value>] [--
skipAWRConfiguration] [--primaryDBOCID <value>] [--sgaSizeInMB <value>] [--
pgaSizeInMB <value>] [--datafileDestination <value>] [--fraDestination
<value>] [--redoLogDestination <value>] [--fraSizeInMB <value>] [--
tdeKeyStoreType <value> [--tdeKeyOCID <value>]] [--tdeKeyOCID <value>] [--
executePrereqs]  [--resume [--sessionID <value>]] | [--revert [--sessionID
<value>]] --waitForCompletion <value>] [--enableFIPS <value>] [--
skipDRConfiguration] [--okvServer <value> --okvAdminUserName <value> [--
okvServerRestPort <value>]] [--okvWalletName <value>]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--oracleHome` specifies the Oracle home path.
- `--oracleHomeName` specifies the Oracle home name.
- `--standbyDBUniqueName` specifies the database unique name for the standby database.
- `--standbyDBDomain` specifies the standby database domain for which the primary database will be configured.
- `--noDBDomain` specifies not to use the database domain name for standby database.
- `--primaryScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the primary database SCAN listener, or the SCAN name of the primary database.
- `--primaryScanPort` specifies the corresponding SCAN port number of the primary database service.
- `--primaryServiceName` specifies the name of the primary database service for which the standby database will be configured.
- `--protectionMode` specifies the Data Guard protection mode to be set when configuring the standby database. Valid values: `MAX_PERFORMANCE|MAX_AVAILABILITY`.
- `--transportType` specifies the Data Guard transport type to be set when configuring the standby database. Valid values: `ASYNC|SYNC`.
- `--activeDG` specifies if the Data Guard configuration will be active or not. Valid values: `true|false`.
- `--standbyBlobFromPrimary` specifies the location of the standby BLOB file, which is prepared from the primary database. This is required only for standby operations.
- `--standbyDBInfoJsonLocation` specifies the location of the info file generated from the primary database for exporting additional metadata. This option is required only for standby operations.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--standbyScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the standby database SCAN listener, or the SCAN name of the standby database.
- `--standbyScanPort` specifies the corresponding SCAN port number of the standby database.
- `--standbySID` specifies the standby database SID for standby configuration.
- `--nodeList` specifies a list of nodes where the standby database is expected to run, including nodes that are already running or configured.
- `--skipAWRConfiguration` specifies whether to skip Oracle AWR configuration as part of the standby database setup. Valid values: `true` (skip AWR configuration) or `false` (configure AWR).
- `--primaryDBOCID` specifies the resource OCID value corresponding to the primary database.
- `--sgaSizeInMB` specifies `sga_target` value in MB.
- `--pgaSizeInMB` specifies the `pga_aggregate_target` value in MB.
- `--datafileDestination` specifies the storage location to use for database datafiles.
- `--fraDestination` specifies the storage location to use for database fast recovery area.
- `--redoLogDestination` specifies the storage location to use for the redo log files.
- `--fraSizeInMB` specifies fast recovery area size value in MB.
- `--tdeKeyStoreType` specifies the TDE keystore type. Valid values: `FILE|KMS|AZURE|GOOGLE|AWS|OKV`
- `--tdeKeyOCID` specifies `KMS/AZURE/GOOGLE/AWS` key OCID to use for TDE. This is applicable only if `KMS/AZURE/GOOGLE/AWS` is selected for TDE keystore type.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--revert` rolls back the previous operation.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- `--enableFIPS` specifies whether to enable FIPS. Set to `false` to disable it. Valid values: `true|false`.
- `--skipDRConfiguration` specifies whether to skip Disaster Recovery (DR) configuration as part of the standby database setup. Valid values: `true` (skip DR configuration) or `false` (configure DR).
- `--okvServer` specifies the Oracle Key Vault server. Comma-delimited list of multiple IP addresses.
- `--okvAdminUserName` specifies the Oracle Key Vault admin user name.
- `--okvServerRestPort` specifies the REST port number for Oracle Key Vault.
- `--okvWalletName` specifies the Oracle Key Vault wallet name.
- Performing configureStandby Operation Using the dbaascli Utility
  To create a standby database, use the `dbaascli dataguard configureStandby` command.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Performing configureStandby Operation Using the dbaascli Utility

To create a standby database, use the `dbaascli dataguard configureStandby` command.

1. Connect to the Virtual Machine where you want to host the standby database as the `opc` user.

   For detailed instructions, see Connecting to a Virtual Machine with SSH.

2. Start a root-user command shell.

   ```
   $ sudo -s
   ```

3. Run the `configureStandby` command.

   ```
   dbaascli dataguard configureStandby --standbySID <standby SID> --activeDG
   <true|false> --dbName <db name> --standbyDBUniqueName <standby db unique
   name> --standbyScanName <comma-delimited list of standby SCAN IP
   addresses> --protectionMode <MAX_PERFORMANCE|MAX_AVAILABILITY> --
   standbyScanPort <standby SCAN port> --oracleHome <oracle home path> --
   standbyDBDomain <standby db domain name>. --primaryServiceName <primary db
   service name> --transportType <ASYNC|SYNC> --primaryScanPort <primary db
   SCAN port> --primaryScanIPAddresses <comma-delimited list primary db SCAN
   IP addresses> --standbyBlobFromPrimary <standby BLOB from primary>
   ```

4. Enter the primary database `SYS`, `TDE`, and `AWR` passwords when prompted.

   Upon successful completion of the command, the standby database starts and becomes operational.

dbaascli dataguard registerStandby

To register a newly created standby database with all existing standby databases and to primary database, use the `dbaascli dataguard registerStandby` command as a third step after `configureStandby` step.

Run this command as the `root` user on the primary cluster. Additionally, in a multiple-standby setup, run the command on all standby clusters except the newly created standby database cluster.

Syntax

```
dbaascli dataguard registerStandby --dbname <value> --standbyDBUniqueName
<value>  --standbyDBDomain <value> | --noDBDomain --standbyScanIPAddresses
<value> [--standbyScanPort <value>] [--standbyServiceName <value>] [--
executePrereqs] [--resume [--sessionID <value>]] | [--revert [--sessionID
<value>]] [--waitForCompletion <value>]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--standbyDBUniqueName` specifies the database unique name of the standby database to be registered with the Oracle Data Guard Broker configuration.
- `--standbyDBDomain` specifies the standby database domain for which the primary database will be configured.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--noDBDomain` specifies not to use the database domain name for standby database.
- `--standbyScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the standby database SCAN listener, or the SCAN name of the standby database.
- `--standbyScanPort` specifies the corresponding SCAN port number of the standby database.
- `--standbyServiceName` specifies the name of the standby database service for which the primary database will be configured.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--revert` rolls back the previous operation.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- Performing registerStandby Operation Using the dbaascli Utility
  To register the specified standby database with the Oracle Data Guard Broker configuration, use the `dbaascli dataguard registerStandby` command.

Performing registerStandby Operation Using the dbaascli Utility

To register the specified standby database with the Oracle Data Guard Broker configuration, use the `dbaascli dataguard registerStandby` command.

For single standby use cases, the `registerStandby` command must be run only on the primary cluster, as there is a one-to-one association between the primary and the standby.

However, in configurations with multiple standby databases, you must run the `registerStandby` command on both the primary cluster and all existing standby clusters—excluding the new standby database being added.

For example, consider a setup with two standby databases: *stdby1* and *stdby2*, where *stdby2* is the new standby to be registered. In this case, run the `registerStandby` command on the primary cluster and on *stdby1*, but not on *stdby2*.

In summary, when adding a new standby database to an existing Oracle Data Guard configuration, run the `registerStandby` command on the primary and on all other previously registered standby clusters, except the new standby being added.

1. Connect to the primary cluster of the Data Guard configuration as the `opc` user.

   For detailed instructions, see Connecting to a Virtual Machine with SSH.

2. Start a root-user command shell.

   ```
   $ sudo -s
   ```

3. Run the `registerStandby` command.

   ```
   dbaascli dataguard registerStandby --dbname <db-name> --
   standbyDBUniqueName <standby-DB-unique-name> --standbyDBDomain <standby-DB-
   domain>
   ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Upon successful completion of the command, the specified standby database will be registered with the Oracle Data Guard Broker configuration.

4. Repeat steps 1 through 3, as performed on the primary cluster, on all existing standby clusters in the Oracle Data Guard Broker configuration, except the one being registered.

dbaascli dataguard deregisterStandby

During standby deletion, run the `dbaascli dataguard deregisterStandby` command before deleting the database on the standby cluster to deregister the standby database from the Oracle Data Guard Broker configuration.

Run this command as the `root` user on the primary cluster. However, in the context of multiple standby databases, this command needs to be executed on all the standby clusters except the target standby.

Syntax

```
dbaascli dataguard deregisterStandby --dbname <value> --standbyDBUniqueName
<value> [--executePrereqs] [--resume [--sessionID <value>]] [--
waitForCompletion <value>]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--standbyDBUniqueName` specifies the database unique name of the standby database to be deregistered from the Oracle Data Guard Broker configuration.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- Performing deregisterStandby Operation Using the dbaascli Utility
During standby deletion, run the `dbaascli dataguard deregisterStandby` command before deleting the database on the standby cluster to deregister the standby database from the Oracle Data Guard Broker configuration.

Performing deregisterStandby Operation Using the dbaascli Utility

During standby deletion, run the `dbaascli dataguard deregisterStandby` command before deleting the database on the standby cluster to deregister the standby database from the Oracle Data Guard Broker configuration.

For single standby use cases, the `deregisterStandby` command must be run only on the primary cluster, as there is a one-to-one association between the primary and the standby.

However, in configurations with multiple standby databases, you must run the `deregisterStandby` command on both the primary cluster and all existing standby clusters—excluding the standby database which is being deregistered currently.

For example, consider a setup with two standby databases: *stdby1* and *stdby2*, where *stdby2* is to be deregistered. In this case, run the `deregisterStandby` command on the primary cluster and on `stdby1`, but not on `stdby2`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

In summary, while deleting a standby database from an existing Oracle Data Guard configuration, run the `deregisterStandby` command on the primary and on all other existing standby clusters before the delete database operation on desired standby cluster.

1. Connect to the primary cluster of the Oracle Data Guard configuration as the `opc` user.

   For detailed instructions, see Connecting to a Virtual Machine with SSH.

2. Start a root-user command shell.

   ```
   $ sudo -s
   ```

3. Run the `deregisterStandby` command.

   ```
   dbaascli dataguard deregisterStandby --dbname <db-name> --
   standbyDBUniqueName <standby-DB-unique-name>
   ```

   Upon successful completion of the command, the specified standby database will be deregistered (removed) from the Oracle Data Guard Broker configuration.

4. Repeat steps 1 through 3, as performed on the primary cluster, on all existing standby clusters in the Oracle Data Guard Broker configuration, except the one being deregistered.

dbaascli dataguard configureAWR

To enable or disable Automatic Workload Repository (AWR) configuration on your Active Data Guard standby, use the `dbaascli dataguard configureAWR` command.

Run this command as the `root` user on the Active Data Guard standby cluster where you want to enable or disable AWR configuration. Use this command if AWR was not configured during the standby addition process.

Syntax

```
dbaascli dataguard configureAWR --dbname <value> { --action <value> | --
enable | --disable } [--executePrereqs] [--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--action` specifies whether to enable or disable AWR. Use `--action enable` to enable AWR and `--action disable` to disable it.
  The `--action` argument is retained for backward compatibility. However, it is recommended to use `--enable` or `--disable`, as they provide the same functionality but are more explicit
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- Performing configureAWR Operation Using the dbaascli Utility
  To configure AWR on an ADG standby database, use the `dbaascli dataguard configureAWR` command.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Performing configureAWR Operation Using the dbaascli Utility

To configure AWR on an ADG standby database, use the `dbaascli dataguard configureAWR` command.

1. Connect to the Virtual Machine where your ADG standby database is hosted as the `opc` user.

   For detailed instructions, see Connecting to a Virtual Machine with SSH.

2. Start a root-user command shell.

   ```
   $ sudo -s
   ```

3. Run the `configureAWR` command.

   To enable AWR, run:

   ```
   $ dbaascli dataguard configureAWR --dbname <db-name> --enable
   ```

   To disable AWR, run:

   ```
   $ dbaascli dataguard configureAWR --dbname <db-name> --disable
   ```

4. Enter the `SYS` and `AWR` passwords when prompted.

   Upon successful completion of the command, the ADG standby database would have been configured with AWR

dbaascli dataguard updateConfiguration

To update the transport mode or protection mode or both the parameters of a Data Guard environment, use the `dbaascli dataguard updateConfiguration` command.

Run this as the `root` user.

When the update transport mode command is run on the primary, only the transport mode of the primary database is updated. To update the transport mode of a standby database, the command must be run separately on that standby.

In contrast, when the update protection mode command is run on the primary, the protection mode is updated for both the primary and standby databases. The protection mode can also be updated from the standby side, in which case both the primary and standby databases are updated.

When updating the transport or protection mode from the primary, the system checks the current modes on both the primary and standby databases and proceeds with the update only if all required conditions are met.

Syntax

```
dbaascli dataguard updateConfiguration --dbname <value> [--protectionMode
<value>] [--transportType <value>] [--standbyDGType <value>] [--
executePrereqs] [--resume [--sessionID <value>]] [--waitForCompletion <value>]
```

Where:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--dbname` specifies the Oracle Database name.
- `--protectionMode` specifies the Data Guard protection mode to be set when configuring the standby database. Valid values: `MAX_PERFORMANCE|MAX_AVAILABILITY`.
- `--transportType` specifies the Data Guard transport type to be set when configuring the standby database. Valid values: `ASYNC|SYNC`.
- `--standbyDGType` specifies the standby database Data Guard type to be set. Valid values: ADG|DG.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- Performing updateConfiguration Operation Using the dbaascli Utility
  To update transport mode and protection mode or both the parameters, use the `dbaascli dataguard updateConfiguration` command.

Performing updateConfiguration Operation Using the dbaascli Utility

To update transport mode and protection mode or both the parameters, use the `dbaascli dataguard updateConfiguration` command.

1. Connect to the Virtual Machine where your ADG standby database is hosted as the `opc` user.

   For detailed instructions, see Connecting to a Virtual Machine with SSH.
2. Start a root-user command shell.

   ```
   $ sudo -s
   ```
3. Run the `updateConfiguration` command.

   ```
   $ dbaascli dataguard updateConfiguration --dbname <db-name> --
   protectionMode MAX_PERFORMANCE|MAX_AVAILABILITY --transportType ASYNC|SYNC
   --standbyDGType ADG|DG.
   ```

   Upon successful completion of the command, the specified Data Guard will be configured with specified transport mode or protection mode.

Database Home Management

This section provides tools for managing Oracle Database homes. It includes commands such as `dbaascli dbhome create` to create a new Oracle Database home and `dbaascli dbHome delete` to remove an existing one. You can also view detailed information about a specific Oracle home with `dbaascli dbHome getDetails` and check which databases are running from a given Oracle home using `dbaascli dbhome getDatabases`. These commands ensure efficient management of database environments.

- dbaascli dbhome create
  To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- dbaascli dbHome delete
  To delete a given Oracle Database home, use the `dbaascli dbHome delete` command.
- dbaascli dbhome getDatabases
  To view information about all Oracle Databases running from a given database Oracle home, use the `dbaascli dbHome getDatabases` command. Specify either the Oracle home location or Oracle home name.
- dbaascli dbHome getDetails
  To view information about a specific Oracle home, use the `dbaascli dbHome getDetails` command. Specify either the Oracle home location or Oracle home name.

dbaascli dbhome create

To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbhome create --version <value>
[--oracleHome <value>]
[--oracleHomeName <value>]
[--enableUnifiedAuditing <value>]
[--imageTag <value>]
[--ImageLocation <value>
```

Where:

- `--version` specifies the version of Oracle Home specified as five numeric segments separated by periods, for example, 19.12.0.0.0
- `--oracleHome` specifies the location of Oracle home
- `--oracleHomeName` specifies user-defined Oracle home name. If not provided, then the default name will be used
- `--enableUnifiedAuditing` specifies `true` or `false` to enable or disable unified auditing link option in Oracle home
- `--imageTag` specifies Oracle home image tag
- `--imageLocation` - path of the image to be used.
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true` or `false`.

Frequently Asked Questions

Q: What is the purpose of the dbaascli dbhome create command?

A: The `dbaascli dbhome create` command is used to create a new Oracle Database home with the desired version.

Q: What are the prerequisites for running the dbaascli dbhome create command?

A: The command must be run as the `root` user.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: How do I specify the Oracle Database version while creating a new Oracle Home?

A: Use the `--version` option followed by the Oracle Database version in the format of five numeric segments separated by periods, such as 19.11.0.0.0.

Q: What does the --oracleHome option specify?

A: The `--oracleHome` option specifies the location where you want to install the Oracle Home. If not provided, the default location will be used.

Q: Can I assign a custom name to the new Oracle Home?

A: Yes, you can use the `--oracleHomeName` option to specify a user-defined name for the Oracle Home. If not specified, a default name will be used.

Q: How do I enable or disable Unified Auditing in the new Oracle Home?

A: Use the `--enableUnifiedAuditing` option and specify `true` to enable or `false` to disable Unified Auditing for the Oracle Home.

Q: What does the --imageTag option do?

A: The `--imageTag` option specifies the Oracle Home image tag, which can be used in cases where the image tag differs from the version.

Q: What is an example of using the dbaascli dbhome create command with version and image tag?

A: An example of the command with version and image tag is:

```
dbaascli dbhome create --version 19.8.0.0.0 --imageTag 19.8.0.0.0
```

This creates an Oracle Home for version 19.8.0.0.0 with the corresponding image tag.

Q: What happens if I don't provide the --oracleHome or --oracleHomeName options?

A: If `--oracleHome` is not provided, the Oracle Home will be installed in the default location. If `--oracleHomeName` is not specified, a default name will be assigned to the Oracle Home.

Q: How can I verify if the Oracle Home creation was successful?

A: After running the command, check the output logs for any success messages or errors. You can also verify the Oracle Home by navigating to the specified location or using Oracle tools like `orainstRoot.sh`.

Q: Is it possible to create multiple Oracle Homes with different versions on the same system?

A: Yes, you can create multiple Oracle Homes with different versions by specifying different values for the `--version` and `--oracleHomeName` options.

Q: What should I do if the Oracle Home creation fails?

A: Check the output logs for detailed error messages. Verify that you have the correct version format, required permissions, and sufficient disk space. Correct any issues and try running the command again.

Example 7-23 dbaascli dbhome create

```
dbaascli dbhome create --version 19.11.0.0.0
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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

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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What happens if I do not provide the --resume or --sessionID options?

A: If the `--resume` or `--sessionID` options are not provided, the command will initiate a new deletion process instead of resuming an interrupted one.

Q: Is there any way to confirm the deletion of the Oracle Home after running the command?

A: You can verify the deletion by checking the output logs for success messages and ensuring the Oracle Home directory is no longer present at the specified location.

Q: Can I delete an Oracle Home that is currently in use by a running database?

A: No, the Oracle Home should not be in use by any running databases or services during the deletion process. Make sure to stop any related databases before running the delete command.

Q: What should I do if the dbaascli dbHome delete command fails?

A: Review the output logs for any error messages. Ensure that the Oracle Home is not in use, verify the correct Oracle Home location or name, and confirm that you have the necessary permissions. After resolving any issues, rerun the command or use the `--resume` option if needed.

Q: Can I delete multiple Oracle Homes at once using the dbaascli dbHome delete command?

A: No, the command only allows you to delete one Oracle Home at a time by specifying either the `--oracleHome` or `--oracleHomeName` option.

Q: What is an example of deleting an Oracle Home by its name?

A: Here is an example of deleting an Oracle Home by name:

dbaascli dbHome delete --oracleHomeName myOracleHome

This command deletes the Oracle Home with the name myOracleHome.

Q: What is an example of deleting an Oracle Home by its location?

A: Here is an example of deleting an Oracle Home by specifying its location:

dbaascli dbHome delete --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1

This command deletes the Oracle Home located at /u01/app/oracle/product/19.0.0/dbhome_1.

Q: Can I cancel the deletion process once it has started?

A: No, once the deletion process has started, it cannot be canceled. Ensure that the Oracle Home is ready for deletion before executing the command.

dbaascli dbhome getDatabases

To view information about all Oracle Databases running from a given database Oracle home, use the `dbaascli dbHome getDatabases` command. Specify either the Oracle home location or Oracle home name.

Run the command as the `root` user.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Syntax

```
dbaascli dbHome getDatabases
{ --oracleHomeName value | --oracleHome value }
```

Where:

- `--oracleHomeName` specifies user-defined Oracle home name
- `--oracleHome` specifies the location (path) of Oracle home

Frequently Asked Questions

Q: What is the purpose of the dbaascli dbHome getDatabases command?

A: The `dbaascli dbHome getDatabases` command is used to view information about all Oracle Databases running from a specified Oracle Database home.

Q: How can I specify the Oracle Database home to check?

A: You can specify the Oracle Database home by using one of the following options:

`--oracleHomeName <value>` to specify the user-defined name of the Oracle home.

`--oracleHome <value>` to specify the full location (path) of the Oracle home.

Q: What is the difference between --oracleHomeName and --oracleHome options?

A:

`--oracleHomeName` refers to a user-defined name for the Oracle home.

`--oracleHome` refers to the physical location (or directory path) of the Oracle home on the system.

Q: How do I run the dbaascli dbHome getDatabases command?

A: To run the command, use the following syntax:

```
dbaascli dbHome getDatabases --oracleHomeName <value>
```

or

```
dbaascli dbHome getDatabases --oracleHome <value>
```

Ensure to run the command as the `root` user.

Q: Can I specify both the Oracle home name and Oracle home location in the same command?

A: No, you can only specify either `--oracleHomeName` or `--oracleHome` in a single command execution. Choose one option based on how you identify the Oracle home.

Q: What kind of information does the 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:

Copyright © 2008, 2026, Oracle and/or its affiliates.

```
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.

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 7-24 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Syntax

```
dbaascli dbHome getDetails
{ --oracleHomeName value | --oracleHome value }
```

Where:

- `--oracleHomeName` specifies user-defined Oracle home name
- `--oracleHome` specifies the location of Oracle home

Frequently Asked Questions

Q: What is the purpose of the dbaascli dbHome getDetails command?

A: The `dbaascli dbHome getDetails` command is used to view detailed information about a specific Oracle home on the system.

Q: How do I specify the Oracle home I want to get details about?

A: You can specify the Oracle home using one of the following options:

`--oracleHomeName <value>` to specify the user-defined name of the Oracle home.

`--oracleHome <value>` to specify the full location (path) of the Oracle home.

Q: What is the difference between --oracleHomeName and --oracleHome?

A:

`--oracleHomeName` is the user-defined name for an Oracle home.

`--oracleHome` refers to the full directory path where the Oracle home is located.

Q: How do I run the dbaascli dbHome getDetails command?

A: To run the command, use the following syntax:

```
dbaascli dbHome getDetails --oracleHomeName <value>
```

or

```
dbaascli dbHome getDetails --oracleHome <value>
```

Make sure to run the command as the `root` user.

Q: Can I specify both --oracleHomeName and --oracleHome in the same command?

A: No, you can only use one option per command execution. You must either specify the Oracle home name or the Oracle home location, not both.

Q: What information does the dbaascli dbHome getDetails command return?

A: The command provides detailed information about the specified Oracle home, such as its version, status, and any other configuration details associated with the Oracle home.

Q: What is an example of using the 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
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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?

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 7-25 dbaascli dbHome getDetails - using Oracle home location

```
dbaascli dbHome getDetails --oracleHome /u02/app/home_db19c/
```

Example 7-26 dbaascli dbHome getDetails - using Oracle home name

```
dbaascli dbHome getDetails --oracleHomeName home_db19c
```

Diagnostics and Health Checks

This section covers tools for maintaining the health and diagnosing issues in Oracle Database environments. Commands like `dbaascli diag collect` are used to gather diagnostic data, while `dbaascli diag healthCheck` allows you to run health checks to identify potential issues.

Copyright © 2008, 2026, Oracle and/or its affiliates.

These tools help ensure the system's stability and performance by proactively monitoring and addressing concerns.

- dbaascli diag collect
  To collect diagnostics, use the `dbaascli diag collect` command.
- dbaascli diag healthCheck
  To run diagnostic health checks, use the `dbaascli diag healthCheck` command.

dbaascli diag collect

To collect diagnostics, use the `dbaascli diag collect` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli diag collect [--components <value>] [--startTime <value>] [--endTime
<value>] [--nodes <value>] [--dbNames <value>]
        {
            [--objectStoreBucketUri <value>]
            | [--destLocation <value>]
        }
        [--waitForCompletion <value>]
```

Where:

- `--components` specifies a list of components for log collection.
  Valid values:
  - `db`
  - `gi`
  - `os`
  - `dbaastools`
  - `all`
- `--startTime` specifies the start time for log collection. Valid date and time format: `YYYY-MM-DDTHH24:MM:SS`
- `--endTime` specifies the end time for log collection. Valid date and time format: `YYYY-MM-DDTHH24:MM:SS`
- `--nodes` specifies a comma-delimited list of nodes to collect logs
- `--dbNames` specifies the database name for which to collect logs. You can specify only one database name.
- `--objectStoreBucketURI` specifies an Object Storage service pre-authenticated request (PAR) URL used to upload collected logs. Logs are collected from Guest VM. For more information, see *Using Pre-Authenticated Requests*.
- `--destLocation` specifies the location on Guest VM to collect logs. Default: `/var/opt/oracle/dbaas_acfs`
- `--waitForCompletion` Values: `true|false`. Default `true`. Specify `false` to run in the background.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Frequently Asked Questions

Q: What is the purpose of the dbaascli diag collect command?

A: The `dbaascli diag collect` command is used to collect diagnostic logs for Oracle Database and associated components such as OS, Grid Infrastructure (GI), and DBaaS tools.

Q: How do I specify which components to collect diagnostics for?

A: You can specify the components using the `--components` option. Valid values are:

`db` (for database logs)

`gi` (for Grid Infrastructure logs)

`os` (for operating system logs)

`dbaastools` (for DBaaS tools logs)

`all` (for collecting logs for all components)

Q: How can I collect logs for a specific time range?

A: Use the following options to specify a time range:

`--startTime <value>` to set the start time for log collection.

`--endTime <value>` to set the end time for log collection.

The time must be in the format: `YYYY-MM-DDTHH24:MM:SS`.

Q: What format should the start and end times be in?

A: Both `--startTime` and `--endTime` should be in the format `YYYY-MM-DDTHH24:MM:SS`. For example, *2024-09-01T15:30:00*.

Q: How do I specify which nodes to collect diagnostics from?

A: You can use the `--nodes` option to specify a comma-delimited list of nodes. For example:

`--nodes node1,node2`

Q: How do I collect logs for a specific database?

A: Use the `--dbNames` option to specify the name of the database for which you want to collect logs. Only one database name can be specified at a time.

Q: How can I store the collected logs in Object Storage?

A: Use the `--objectStoreBucketUri` option to specify the pre-authenticated request (PAR) URL for your Object Storage bucket where the logs will be uploaded.

Q: Can I collect logs to a local directory instead of Object Storage?

A: Yes, you can use the `--destLocation` option to specify a directory on the guest VM where the logs will be collected. The default location is `/var/opt/oracle/dbaas_acfs`.

Q: What happens if I don’t specify a destination for the logs?

A: If no destination is specified, the collected logs will be stored in the default location `/var/opt/oracle/dbaas_acfs` on the guest VM.

Q: What does the --waitForCompletion option do?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `--waitForCompletion` option specifies whether to wait for the command to finish before returning control to the user. The default value is `true`. If you specify `false`, the operation runs in the background.

Q: How can I run the log collection in the background?

A: Set the `--waitForCompletion` option to `false` to run the log collection process in the background:

```
dbaascli diag collect --waitForCompletion false
```

Q: Can I resume a previous log collection session with this command?

A: No, the `dbaascli diag collect` command does not support resuming previous sessions. You will need to start a new log collection process.

Q: How can I ensure that logs are uploaded directly to the Object Storage?

A: You can provide a valid `--objectStoreBucketUri` (Object Storage pre-authenticated request URL) where the logs will be uploaded after collection.

Q: Can I collect logs for multiple databases at once?

A: No, you can only specify one database name at a time using the `--dbNames` option.

Q: What should I do if I want to collect all available logs for all components?

A: Use `--components` all to collect logs from all components, including the database, Grid Infrastructure, operating system, and DBaaS tools.

Q: What is an example command to collect logs for the database component from a specific time range?

A: Here's an example command:

```
dbaascli diag collect --components db --startTime 2024-09-01T12:00:00 --endTime
2024-09-01T14:00:00 --dbname orcl
```

Q: What is an example command to collect logs and upload them to Object Storage?

A: Here's an example command:

```
dbaascli diag collect --components db --objectStoreBucketUri https://
objectstorage.example.com/n/namespace-string/b/bucket-name/o/PAR-URL
```

Q: What is the default behavior if the --components option is not specified?

A: If you do not specify the `--components` option, the command may not know which logs to collect and could fail. It is recommended to always specify the components you want to collect logs from.

Q: Can I specify both the --objectStoreBucketUri and --destLocation options in the same command?

A: No, you must choose one destination—either Object Storage via `--objectStoreBucketUri` or a local directory via `--destLocation`.

Q: What should I do if I encounter an error while using the dbaascli diag collect command?

A: Double-check that you have provided valid component names, date/time formats, and destination options. Also, ensure you are running the command as the `root` user.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Using Pre-Authenticated Requests

dbaascli diag healthCheck

To run diagnostic health checks, use the `dbaascli diag healthCheck` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli diag healthCheck
[--destLocation]
[--nodes]
[--objectStoreBucketURI]
```

Where:

- `--destLocation` specifies the location on Guest VM to collect logs. Default: `/var/opt/oracle/dbaas_acfs`
- `--nodes` specifies a comma-delimited list of nodes to collect logs
- `--objectStoreBucketURI` specifies an Object Storage service pre-authenticated request (PAR) URL used to upload collected logs. Logs are collected from Guest VM. For more information, see *Using Pre-Authenticated Requests*.

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Yes, you can specify multiple nodes using the `--nodes` option in a comma-delimited format. For example: `--nodes node1,node2`.

Q: What is an example command to run a health check on a specific node?

A: Here's an example command to run the health check on a specific node:

```
dbaascli diag healthCheck --nodes node1
```

Q: How can I store the logs in Object Storage instead of on the local machine?

A: You can provide a pre-authenticated request (PAR) URL using the `--objectStoreBucketURI` option to store the logs in Object Storage.

Q: Can I specify both --destLocation and --objectStoreBucketURI at the same time?

A: Yes, you can specify both `--destLocation` for local storage and `--objectStoreBucketURI` to upload logs to Object Storage.

Q: What should I do if I encounter an error while running the dbaascli diag healthCheck command?

A: Ensure that you are running the command as the root user and have provided valid options for `--destLocation`, `--nodes`, or `--objectStoreBucketURI`. Verify that the node names are correct if specified.

Q: Can I run the health check in the background?

A: The `dbaascli diag healthCheck` command does not have an explicit background mode, but you can run it in the background by appending & at the end of the command.

Q: What happens if I don't provide the --nodes option?

A: If the `--nodes` option is not provided, the health check will be performed on all nodes in the cluster by default.

Q: Can I resume a previous health check session using this command?

A: No, the `dbaascli diag healthCheck` command does not support resuming previous sessions. You must initiate a new health check each time.

Q: What is an example command to run a health check and upload logs to Object Storage?

A: Here's an example command:

```
dbaascli diag healthCheck --objectStoreBucketURI https://
objectstorage.example.com/n/namespace-string/b/bucket-name/o/PAR-URL
```

Q: What is the default behavior if I don't specify --destLocation or --objectStoreBucketURI?

A: If neither `--destLocation` nor `--objectStoreBucketURI` is specified, the health check logs will be collected in the default directory `/var/opt/oracle/dbaas_acfs` on the local machine.

Q: Can I limit the health check to specific components or logs?

A: No, the `dbaascli diag healthCheck` command does not allow you to specify individual components or logs. It performs a general diagnostic health check for the system.

Q: What should I verify before running the dbaascli diag healthCheck command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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

Grid Infrastructure Management

This section focuses on managing Oracle Grid Infrastructure, which supports clustering and high availability. Key tasks include configuring a Grid Infrastructure home (`dbaascli gridHome create`), upgrading the Grid Infrastructure (`dbaascli grid upgrade`), and managing TCPS (Transport Layer Security) certificates by configuring (`dbaascli grid configureTCPS`), removing (`dbaascli grid removeTCPSCert`), or rotating them (`dbaascli grid rotateTCPSCert`). These commands ensure efficient setup, maintenance, and security of the Grid Infrastructure.

- dbaascli gridHome create
  To configure Grid Infrastructure home, use the `dbaascli gridHome create` command.
- dbaascli grid configureTCPS
  To configure TCPS for the existing cluster, use the `dbaascli grid configureTCPS` command.
- dbaascli grid 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 Infrastrucure from one major version to another, use the `dbaascli grid upgrade` command.

dbaascli gridHome create

To configure Grid Infrastructure home, use the `dbaascli gridHome create` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli gridHome create --version value [--resume [--sessionID value]] [--
waitForCompletion value]
```

Where:

- `--version` specifies the Grid home version
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Frequently Asked Questions

Q: What is the purpose of the dbaascli gridHome create command?

A: The `dbaascli gridHome create` command is used to configure a Grid Infrastructure home for Oracle Grid infrastructure in an Exadata Cloud@Customer environment.

Q: What is the prerequisite for running the dbaascli gridHome create command?

A: The command must be run as the `root` user.

Q: How do I specify the version of Grid Infrastructure home to be created?

A: Use the `--version` option to specify the Grid home version. This is a mandatory option when creating the Grid home.

Q: Can I resume a previous dbaascli gridHome create session?

A: Yes, you can use the `--resume` option to resume a previous session. Optionally, you can specify the session ID using the `--sessionID` option to resume a specific session.

Q: What does the --resume option do in the dbaascli gridHome create command?

A: The `--resume` option allows you to resume a previously interrupted or incomplete operation.

Q: How do I run the operation in the background?

A: You can set the `--waitForCompletion` option to false to run the operation in the background. The valid values for this option are `true` (default) or `false`.

Q: What is the default behavior if --waitForCompletion is not specified?

A: If `--waitForCompletion` is not specified, the operation will run in the foreground, and the command will wait for the operation to complete before returning control to the user.

Q: What is the purpose of the --sessionID option?

A: The `--sessionID` option is used to specify the ID of a previous session that you want to resume, in case of an incomplete or interrupted operation.

Q: Can I use the dbaascli gridHome create command to upgrade an existing Grid home?

A: No, this command is used specifically for configuring a new Grid Infrastructure home, not for upgrading an existing one.

Q: What is an example command to create a Grid home with version 19.9.0.0.0?

A: Here's an example command:

```
dbaascli gridHome create --version 19.9.0.0.0
```

Q: What should I do if the dbaascli gridHome create command is interrupted?

A: You can resume the operation by using the `--resume` option. If you have the session ID, you can provide it using the `--sessionID` option to resume a specific session.

Q: Can I specify the --resume option without a session ID?

A: Yes, you can use the `--resume` option without specifying a session ID. In this case, the system will attempt to resume the most recent session.

Q: What happens if I specify --waitForCompletion false?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: If you specify `--waitForCompletion false`, the operation will run in the background, allowing you to continue using the command line while the operation completes.

Q: Is it possible to track the progress of a background operation?

A: The dbaascli command does not provide a direct way to track background operations. You may need to check system logs or the session status manually.

Q: Can I specify an invalid Grid home version when creating it?

A: No, the `--version` option must contain a valid Grid Infrastructure version. If an invalid version is provided, the command will return an error.

Q: How can I determine the Grid home version to use for creation?

A: You can refer to Oracle's documentation or consult with your database administrator to determine the correct Grid home version to use for your environment.

Q: What should I verify before running the dbaascli gridHome create command?

A: Ensure that you are logged in as the `root` user and that the version you want to use for Grid Infrastructure is compatible with your environment.

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]
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

Where:

- `--pkcs12WalletFilePath` specifies the absolute path of the certificate file, which is in the `pkcs12` wallet format
- `--caCertChain` concatenated list of certs, containing intermediate CA's and root CA certs
- `--precheckOnly` specifies `yes` to run only the prechecks for this operation. Valid values: `yes` or `no`.
- `--serverCert` specifies the path of PEM certificate to use or rotate for TCPS configuration.
- `--privateKey` specifies the path of the private key file of the certificate.
- `--certType` type of the cert to be added to the Grid Infrastructure wallet. Accepted values are: `SELF_SIGNED_CERT`, `CA_SIGNED_CERT`, or `PKCS12_CERT`. Default: `SELF_SIGNED_CERT`
- `--privateKeyPasswordProtected` specifies if the private key is password protected or not. Valid values: `true` or `false`. Default: `true`.

Frequently Asked Questions

Q: What is the purpose of the dbaascli grid configureTCPS command?

A: The `dbaascli grid configureTCPS` command is used to configure Transport Layer Security (TCPS) for the existing cluster in an Oracle Exadata environment.

Q: What is the prerequisite for running the 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What does the --precheckOnly option do?

A: The `--precheckOnly` option specifies whether to run only the prechecks for the TCPS configuration operation. Accepted values are `yes` or `no`.

Q: What does the --serverCert option specify?

A: The `--serverCert` option specifies the path to the PEM certificate that will be used or rotated for the TCPS configuration.

Q: How do I specify the private key for TCPS configuration?

A: Use the `--privateKey` option to specify the path to the private key file associated with the server certificate.

Q: What are the accepted values for the --certType option?

A: The accepted values for the `--certType` option are:

`SELF_SIGNED_CERT`

`CA_SIGNED_CERT`

`PKCS12_CERT`

The default value is `SELF_SIGNED_CERT`.

Q: Is the private key password protected by default?

A: Yes, the `--privateKeyPasswordProtected` option is set to `true` by default, indicating that the private key is password protected. You can set it to `false` if the private key is not password protected.

Q: Can I run a precheck before configuring TCPS?

A: Yes, you can run only the prechecks for the operation by setting the `--precheckOnly` option to yes. This helps to validate the environment before making changes.

Q: What happens if I provide an incorrect path for the PKCS12 wallet file?

A: If the `--pkcs12WalletFilePath` contains an incorrect path, the command will fail, and the TCPS configuration will not proceed.

Q: What should I do if the private key is password protected?

A: If the private key is password protected, ensure that the `--privateKeyPasswordProtected` option is set to `true` (which is the default).

Q: Can I specify my own CA-signed certificate for TCPS configuration?

A: Yes, you can specify your own CA-signed certificate by using the `--serverCert` and `--privateKey` options, and by setting the `--certType` to `CA_SIGNED_CERT`.

Q: What is an example command to configure TCPS using a self-signed certificate?

A: Here's an example:

```
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`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-27 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
```

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `dbaascli grid removeTCPSCert` command is used to remove existing TCPS certificates from the Grid Infrastructure wallet in an Oracle Exadata environment.

Q: What is the prerequisite for running the dbaascli grid removeTCPSCert command?

A: The command must be run as the `root` user.

Q: What does the --subject option specify in the dbaascli grid removeTCPSCert command?

A: The `--subject` option specifies the subject of the certificate to be removed from the Grid Infrastructure wallet.

Q: What is the purpose of the --userCert flag?

A: The `--userCert` flag indicates that the certificate to be removed is a user certificate.

Q: When should I use the --trustedCert flag?

A: Use the `--trustedCert` flag when removing a trusted certificate from the Grid Infrastructure wallet.

Q: What does the --requestedCert flag do?

A: The `--requestedCert` flag indicates that the certificate being removed is a requested certificate.

Q: What does the --serialNumber option specify?

A: The `--serialNumber` option specifies the serial number of the certificate to be removed. It is useful for uniquely identifying a certificate when there are multiple certificates with the same subject.

Q: What is the purpose of the --executePrereqs option?

A: The `--executePrereqs` option runs prerequisite checks before removing the certificate and reports the results, ensuring that the environment is properly prepared for the operation.

Q: What does the --resume option do?

A: The `--resume` option resumes the removal operation if it was previously interrupted.

Q: How do I specify a session ID when resuming an interrupted operation?

A: Use the `--sessionID` option to specify the session ID of the interrupted operation that you want to resume.

Q: What does the --bounceListeners flag do?

A: The `--bounceListeners` flag is used to restart the Grid Infrastructure listener and scan listener after the TCPS certificate is removed.

Q: Can I remove a TCPS certificate without bouncing the listeners?

A: Yes, the `--bounceListeners` flag is optional. If you don’t specify it, the listeners will not be bounced automatically.

Q: How can I ensure that the operation will run safely?

A: You can use the `--executePrereqs` option to perform prerequisite checks before running the command, ensuring that everything is in order before the removal process.

Q: What should I do if I need to remove a specific user certificate by serial number?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Use the `--subject` option to specify the certificate's subject, the `--userCert` flag to indicate that it is a user certificate, and the `--serialNumber` option to specify the certificate's serial number.

Q: Can I remove multiple certificates at once?

A: No, the command is designed to remove a single certificate at a time based on the provided subject and other parameters.

Q: What happens if the certificate removal process is interrupted?

A: You can resume the operation using the `--resume` option along with the `--sessionID` of the interrupted process.

Q: Do I need to run the command as the root user?

A: Yes, the `dbaascli grid removeTCPSCert` command must be run as the `root` user to have the necessary privileges for removing TCPS certificates.

Q: How can I identify the certificate I want to remove?

A: You can identify the certificate by its subject, and optionally, by its serial number to ensure you're targeting the correct certificate for removal.

Q: What is an example command to remove a trusted certificate?

A: Here's an example:

```
dbaascli grid removeTCPSCert --subject "CN=example_cert" --trustedCert
```

dbaascli grid rotateTCPSCert

To rotate TCPS certificates, use the dbaascli grid rotateTCPSCert command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli grid rotateTCPSCert
[--pkcs12WalletFilePath]
[--caCertChain]
[--precheckOnly]
[--serverCert]
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

Where:

- `--pkcs12WalletFilePath` specifies the absolute path of the certificate file, which is in the `pkcs12` wallet format
- `--caCertChain` concatenated list of certs, containing intermediate CA's and root CA certs
- `--precheckOnly` specifies `yes` to run only the prechecks for this operation. Valid values: `yes` or `no`.
- `--serverCert` specifies the path of PEM certificate to use or rotate for TCPS configuration.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--privateKey` specifies the path of the private key file of the certificate.
- `--certType` type of the cert to be added to the Grid Infrastructure wallet. Accepted values are: `SELF_SIGNED_CERT`, `CA_SIGNED_CERT`, or `PKCS12_CERT`. Default: `SELF_SIGNED_CERT`
- `--privateKeyPasswordProtected` specifies if the private key is password protected or not. Valid values: `true` or `false`. Default: `true`.

Frequently Asked Questions

Q: What is the purpose of the dbaascli grid rotateTCPSCert command?

A: The `dbaascli grid rotateTCPSCert` command is used to rotate TCPS (Transport Layer Security Protocol) certificates in the Grid Infrastructure wallet in Oracle Exadata environments.

Q: What is the prerequisite for running the dbaascli grid rotateTCPSCert command?

A: The command must be run as the `root` user.

Q: What does the --pkcs12WalletFilePath option specify?

A: The `--pkcs12WalletFilePath` option specifies the absolute path to the certificate file in the PKCS12 wallet format for the TCPS configuration.

Q: What is the purpose of the --caCertChain option?

A: The `--caCertChain` option specifies a concatenated list of certificates, including intermediate CA and root CA certificates, for the TCPS configuration.

Q: What does the --precheckOnly option do?

A: The `--precheckOnly` option allows you to run prechecks without making any actual changes. The valid values are "`yes`" to run only the prechecks and "`no`" to proceed with the rotation.

Q: How is the --serverCert option used?

A: The `--serverCert` option specifies the path to the PEM (Privacy Enhanced Mail) server certificate that is used or rotated for the TCPS configuration.

Q: What does the --privateKey option specify?

A: The `--privateKey` option specifies the path to the private key file corresponding to the server certificate used for TCPS rotation.

Q: What are the valid values for the --certType option?

A: The `--certType` option accepts the following values for specifying the type of certificate to be added to the Grid Infrastructure wallet:

`SELF_SIGNED_CERT` (default)

`CA_SIGNED_CERT`

`PKCS12_CERT`

Q: What does the --privateKeyPasswordProtected option do?

A: The `--privateKeyPasswordProtected` option indicates whether the private key is password-protected. Valid values are `true` (default) and `false`

Q: Can I run the dbaascli grid rotateTCPSCert command without rotating the certificates?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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`).

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 7-28 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
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli grid upgrade

To upgrade Oracle Grid Infrastrucure from one major version to another, use the `dbaascli grid upgrade` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli grid upgrade --version
[--resume]
[--executePrereqs]
[--containerURL]
[--softwareOnly]
[--targetHome]
[--revert]
```

Where:

- `--version` specifies the target version
- `--resume` resumes the previous run
- `--executePrereqs` runs prereqs for Grid Infrastrucure upgrade
- `--containerUrl` specifies the custom URL for fetching Grid Infrastrucure 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What is the purpose of the --containerURL option?

A: The `--containerURL` option specifies a custom URL to fetch the Grid Infrastructure software image for the upgrade.

Q: What does the --softwareOnly option do?

A: The `--softwareOnly` option installs only the Grid Infrastructure software without configuring or upgrading the Grid environment.

Q: When would you use the --targetHome option?

A: The `--targetHome` option specifies the path to the existing target Grid home where the upgrade will be performed.

Q: What happens if the upgrade fails?

A: If the upgrade fails, you can use the `--revert` option to roll back the upgrade to its previous state.

Q: Can I perform a Grid Infrastructure upgrade in stages?

A: Yes, by using the `--softwareOnly` option, you can first install the software and then later complete the full upgrade, allowing for staged upgrades.

Q: How do I use the dbaascli grid upgrade command to upgrade only the software?

A: Use the following syntax to upgrade only the software:

```
dbaascli grid upgrade --version <target_version> --softwareOnly
```

Q: Can I check for upgrade prerequisites without performing the upgrade?

A: Yes, you can run only the prerequisite checks using:

```
dbaascli grid upgrade --version <target_version> --executePrereqs
```

Q: How can I upgrade Grid Infrastructure using a custom container URL?

A: You can specify the URL for fetching the Grid Infrastructure image as follows:

```
dbaascli grid upgrade --version <target_version> --containerURL <custom_url>
```

Q: How can I resume a previously interrupted upgrade process?

A: To resume a previously interrupted or failed upgrade, use:

```
dbaascli grid upgrade --version <target_version> --resume
```

Q: What does the --revert option do in the dbaascli grid upgrade command?

A: The `--revert` option rolls back a failed or interrupted Grid Infrastructure upgrade to its original state.

Q: Can I perform a full upgrade without configuring the Grid Infrastructure immediately?

A: Yes, you can first install only the software using the `--softwareOnly` option and then configure it later.

Q: What should I do if an upgrade fails and I want to undo the changes?

A: Use the `--revert` option to roll back the failed upgrade:

```
dbaascli grid upgrade --version <target_version> --revert
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-29 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
```

Patching and Upgrading

This section provides tools for updating and maintaining Oracle environments through patching and upgrades. It includes commands such as `dbaascli grid patch` to patch Oracle Grid Infrastructure, `dbaascli dbHome patch` to apply patches to Oracle homes, and `dbaascli database move` to move databases between homes during upgrades or patching processes. These commands help ensure systems remain secure, stable, and up-to-date.

- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.
- Miscellaneous

dbaascli grid patch

To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli grid patch
 {
            --targetVersion <value>
            | --targetHome <value>
        }
        [--executePrereqs] [--nodeList <value>] [--continueWithDbDowntime] [--
drainTimeoutInSeconds <value>] [--containerURL <value>] [--imageFile <value>]
[--patchInParallel]
        {
            [--resume [--sessionID <value>]]
            | [--rollback [--sessionID <value>]]
        }
        [--waitForCompletion <value>]
```

Where:

- `--targetVersion` specifies the target version of Oracle Home specified as five numeric segments separated by periods (e.g. 19.12.0.0.0)

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--targetHome` specifies the fully qualified path of the target Grid Infrastructure home for the out of place patching
- `--containerURL` specifies custom URL for fetching Grid Infrastructure image
- `--executePrereqs` option to run prereqs
- `--nodeList` specifies a comma-delimited list of nodes if patching has to be performed on a subset of nodes
- `--patchInParallel` specifies to perform patching remote nodes in parallel
- `--rollback` specifies to roll back patched Oracle home
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--continueWithDbDowntime` continues patching with database downtime. This option can be used in environments wherein there is only 1 active instance up and the patching operation can be continued even with a downtime.
- `--drainTimeoutInSeconds` specifies the time (in seconds) to complete the resource draining while stopping the database
- `--createImage` creates an image from a copy of the active Grid home, patched to the specified target version
  - `--createImageDir` specifies fully qualified path of the directory where the image is to be created
- `--imageFile` specifies fully qualified path of the image to be used
- `--patchInParallel` performs the patching of the remote nodes in parallel
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true|false`

Frequently Asked Questions

Q: What does the dbaascli grid patch command do?

A: The `dbaascli grid patch` command is used to patch Oracle Grid Infrastructure to a specified minor version.

Q: Do I need special permissions to run the dbaascli grid patch command?

A: Yes, you need to run the `dbaascli grid patch` command as the `root` user.

Q: Can I specify a target version when patching Oracle Grid Infrastructure?

A: Yes, you can specify the target version using the `--targetVersion` option.

Q: How do I specify the target version for the patch?

A: Use the `--targetVersion` option followed by the version number in the format `19.12.0.0.0`.

Q: What does the --containerURL option do in the dbaascli grid patch command?

A: The `--containerURL` option allows you to specify a custom URL for fetching the Grid Infrastructure image.

Q: What is the purpose of the --executePrereqs option?

A: The `--executePrereqs` option is used to run prerequisite checks before applying the patch.

Q: How can I patch a subset of nodes using the dbaascli grid patch command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Use the `--nodeList` option followed by a comma-delimited list of node names to patch only a subset of nodes.

Q: What happens if I use the --rollback option?

A: The `--rollback` option will roll back the patched Oracle home to its previous state.

Q: Can I resume a previous patching session?

A: Yes, you can use the `--resume` option to resume the last patching session. If you have a specific session ID, you can specify it with the `--sessionID` option.

Q: What is the --continueWithDbDowntime option used for?

A: The `--continueWithDbDowntime` option allows you to continue patching even if there is database downtime, typically used in environments where there is only one active instance.

Q: How do I create an image from a patched Grid home?

A: Use the `--createImage` option to create an image. You can specify the directory where the image should be created using the `--createImageDir` option.

Q: What is the purpose of the --imageFile option?

A: The `--imageFile` option allows you to specify the fully qualified path of the image file to be used for patching.

Q: How can I run the dbaascli grid patch command in the background?

A: You can use the `--waitForCompletion` option set to `false` to run the operation in the background.

Q: Can I use a custom URL to fetch the patch image?

A: Yes, you can use the `--containerURL` option to specify a custom URL for fetching the Grid Infrastructure image.

Q: How do I specify which nodes to patch if I don't want to patch all of them?

A: You can specify the nodes you want to patch using the `--nodeList` option with a comma-separated list of node names.

Q: What should I do if I need to roll back a patch?

A: Use the `--rollback` option in the `dbaascli grid patch` command to roll back the patch.

Q: How do I handle a patching operation if my environment only has one active instance and I need to continue with downtime?

A: Use the `--continueWithDbDowntime` option to continue patching even with database downtime.

Q: Can I create an image of the patched Grid home?

A: Yes, you can use the `--createImage` option to create an image of the patched Grid home. If needed, specify the directory where the image should be saved using `--createImageDir`.

Q: What should I do if I want to resume a patching session after an interruption?

A: Use the `--resume` option to resume the patching session. If you know the session ID, you can specify it with `--sessionID`.

Q: What if the patching process fails midway?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: If the patching process fails, you can use the `--resume` option to restart the process. You can also use the `--rollback` option to revert to the previous state.

Q: How can I ensure all prerequisites are met before patching?

A: Use the `--executePrereqs` option to run all prerequisite checks before applying the patch.

Q: Can I perform patching in the background to avoid holding up the terminal?

A: Yes, you can use the `--waitForCompletion false` option to run the patching process in the background.

Q: How can I create a Grid home image after patching?

A: Use the `--createImage` option to create a new image from the patched Grid home. Specify the directory using `--createImageDir` if needed.

Q: How do I use an existing image file for patching?

A: You can use the `--imageFile` option to specify the fully qualified path to the image file you want to use for patching.

Q: What should I do if I want to avoid database downtime during patching?

A: Ensure that your environment has more than one active instance running. You can avoid using the `--continueWithDbDowntime` option, which is meant for environments with only one active instance.

Q: How do I know the progress of a patch running in the background?

A: If you run the patch with `--waitForCompletion false`, you can check the status of the background job using operating system commands like `ps` or check the logs located in the Grid home.

Q: Is it possible to patch to a higher major version using dbaascli grid patch?

A: No, `dbaascli grid patch` only allows patching to a minor version of the current major version. For major upgrades, you will need to follow a different upgrade process.

Q: Can I skip specific prerequisite checks during patching?

A: No, when you use `--executePrereqs`, all prerequisite checks will be executed. However, you can review the results of the prerequisite checks and manually handle any issues before proceeding.

Q: What should I do if the patching process is stuck or hanging?

A: If the patching process is unresponsive, you can stop it using operating system commands and then resume using the `--resume` option. If that doesn't work, try using the `--rollback` option to revert the patch.

Q: Can I automate the patching process across multiple clusters?

A: Yes, using scripts that include the `dbaascli grid patch` command with appropriate options, you can automate patching across different clusters.

Q: Where can I find logs for the patching process?

A: Logs are typically located in the Oracle Grid home logs directory or the default location specified during setup. You can monitor these logs for details about the patching process.

Q: Is it possible to create a silent patch process with no user interaction?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Yes, by specifying all necessary options in the command and running it in the background (`--waitForCompletion false`), you can create a non-interactive patching process.

Q: Can I check for available patch updates before applying a patch?

A: The `dbaascli grid patch` command itself does not have an option to list available patches. However, you can use Oracle's standard methods such as Oracle Support to identify the latest patches.

Q: Can I use dbaascli to patch multiple Oracle Homes?

A: No, the `dbaascli grid patch` command is designed to patch a specific Oracle Grid Infrastructure home at a time. You would need to run the command separately for each home.

Q: Is there a way to prevent downtime completely when patching Grid Infrastructure?

A: To minimize downtime, ensure that your environment has multiple active database instances (RAC configuration) so that patching can be done node by node. The `--continueWithDbDowntime` option should not be used in this case.

Q: How do I handle patching for RAC One Node environments?

A: In RAC One Node environments, you need to be cautious with the `--continueWithDbDowntime` option, as there may be only one active instance. Review the Oracle documentation for specific patching guidelines for RAC One Node.

Q: Can I view the session history of previous patches?

A: The dbaascli utility does not provide a direct way to view session history. However, logs of previous patching sessions can be found in the Grid home logs directory.

Example Use Cases

Example 1: Basic Grid Patching

```
dbaascli grid patch --targetVersion 19.12.0.0.0
```

Patches the Oracle Grid Infrastructure to version 19.12.0.0.0.

Example 2: Patching with a Custom Container URL

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --containerURL https://
example.com/custom/url
```

Patches Grid Infrastructure to version 19.12.0.0.0, using a custom container URL to fetch the Grid Infrastructure image.

Example 3: Patching with Prerequisite Checks

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --executePrereqs
```

Patches Grid Infrastructure to version 19.12.0.0.0 after running the prerequisite checks.

Example 4: Patching on a Subset of Nodes

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --nodeList node1,node2,node3
```

Patches Grid Infrastructure to version 19.12.0.0.0 on the specified nodes (node1, node2, and node3).

Example 5: Rolling Back the Patch

```
dbaascli grid patch --rollback
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Rolls back the last applied patch on the Oracle Grid Infrastructure.

Example 6: Resuming a Previous Patch Operation

```
dbaascli grid patch --resume
```

Resumes the previous patching operation from where it was stopped.

Example 7: Resuming a Patch Operation with a Specific Session ID

```
dbaascli grid patch --resume --sessionID 12345
```

Resumes the patching operation using session ID 12345.

Example 8: Patching with Database Downtime Allowed

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --continueWithDbDowntime
```

Patches the Grid Infrastructure to version 19.12.0.0.0 while allowing downtime of the database if needed.

Example 9: Creating a Patched Image

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --createImage --createImageDir /
path/to/dir
```

Creates an image of the patched Grid home (version 19.12.0.0.0) and stores it in the specified directory.

Example 10: Using an Existing Image File

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --imageFile /path/to/image/
file.zip
```

Patches Grid Infrastructure to version 19.12.0.0.0 using an existing image file located at /path/to/image/file.zip.

Example 11: Running the Patching Operation in the Background

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --waitForCompletion false
```

Patches Grid Infrastructure to version 19.12.0.0.0 and runs the operation in the background.

Example 12: Combination of Prerequisites, Custom URL, and Subset of Nodes

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --executePrereqs --containerURL
https://example.com/custom/url --nodeList node1,node2
```

Patches Grid Infrastructure to version 19.12.0.0.0, runs prerequisite checks, uses a custom URL for the image, and applies the patch only on node1 and node2.

Example 13: Creating a Patched Image with an Existing Image File

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --createImage --createImageDir /
path/to/dir --imageFile /path/to/existing/image.zip
```

Creates a patched image and stores it in the specified directory while using an existing image file for the patch.

Example 14: Verifying Prerequisites Without Patching

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --executePrereqs
```

Verifies whether all prerequisites are met for patching to version 19.12.0.0.0 without actually applying the patch.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 15: Running Patch and Ignoring Prerequisite Failures

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --continueWithDbDowntime --
executePrereqs
```

Runs the patch even if some prerequisite checks fail. This is useful in scenarios where downtime is allowed, and certain prerequisites can be ignored.

Example 16: Checking Patch Logs for Issues

```
tail -f /u01/app/grid/logs/grid_patch.log
```

Monitors the patch log in real time to diagnose any issues during the patching process.

Example 17: Applying the Patch in a Parallel Environment

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --nodeList node1,node2 --
waitForCompletion false
```

Patches Grid Infrastructure on a subset of nodes (node1 and node2) and runs the process in the background.

Example 18: Using a Specific Image File from an External Source

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --imageFile /mnt/images/
grid_patch_19.12.zip
```

Patches Grid Infrastructure using a pre-downloaded image file located on an external storage device.

Example 19: Running Patch with a Custom Session ID

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --resume --sessionID 67890
```

Resumes a patching operation that was interrupted, using session ID 67890.

Example 20: Scheduling Patching to Run at a Later Time

```
echo "dbaascli grid patch --targetVersion 19.12.0.0.0" | at 02:00
```

Schedules the patching command to run at 2:00 AM using the at command in Linux.

Example 21: Specifying Timeout for Completion

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --waitForCompletion true --
continueWithDbDowntime --timeout 7200
```

Patches Grid Infrastructure while allowing downtime, but waits up to 7200 seconds (2 hours) for completion before timeout.

Example 22: Creating a Custom Image for Another Environment

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --createImage --createImageDir /
backups/images/grid_patch
```

Creates a custom image of the patched Grid Infrastructure to store in the `/backups/images/grid_patch` directory for use in other environments.

Example 23: Patch Recovery After Interruption

```
dbaascli grid patch --resume --continueWithDbDowntime
```

Recovers and resumes the patching process if it was interrupted, with database downtime allowed.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 24: Combining Prerequisites Check with Background Execution

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --executePrereqs --
waitForCompletion false
```

Checks prerequisites and runs the patch in the background.

Example 25: Skipping Image Creation for Faster Patching

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --patchInParallel --
continueWithDbDowntime --waitForCompletion false
```

Patches Grid Infrastructure to version 19.12.0.0.0 in parallel across nodes, with database downtime allowed, and without creating an image to speed up the process.

Example 26: Monitoring Patch Progress Through Logs

```
tail -f /u01/app/grid/logs/grid_patch_progress.log
```

Monitors the log file for patching progress in real-time, providing insights into each step of the patching process.

Example 27: Patching with Custom Drain Timeout

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --drainTimeoutInSeconds 3600 --
continueWithDbDowntime
```

Patches Grid Infrastructure and sets a custom timeout of 1 hour (3600 seconds) to allow graceful resource draining during database downtime.

Example 28: Applying a Patch to Specific Nodes with Prerequisite Checks

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --nodeList node1,node4 --
executePrereqs
```

Patches only nodes node1 and node4 to version 19.12.0.0.0 and runs the prerequisite checks beforehand.

Example 29: Patching Without Waiting for Completion

```
dbaascli grid patch --targetVersion 19.12.0.0.0 --waitForCompletion false
```

Begins patching the Grid Infrastructure to version 19.12.0.0.0 in the background, allowing other tasks to be performed without waiting for the process to complete.

Example 30: Reapplying a Failed Patch After a Drain Timeout Issue

```
dbaascli grid patch --resume --drainTimeoutInSeconds 7200
```

Resumes the previous patching session and extends the resource draining timeout to 2 hours (7200 seconds) in case it failed due to insufficient time in the previous attempt.

Example 31: Viewing Patch Logs in Real-Time with Specific Session ID

```
tail -f /u01/app/grid/logs/grid_patch_12345.log
```

Monitors the log file for the patching session with session ID 12345 in real-time.

Example 32: Patching to a New Target Home

```
dbaascli grid patch --targetHome /u01/app/grid_home_19c --executePrereqs
```

Performs an out-of-place patch to a new Oracle Grid home located at /u01/app/grid_home_19c, with prerequisite checks.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 33: Stopping a Background Patch Job

```
ps -ef | grep dbaascli | grep patch | awk '{print $2}' | xargs kill -9
```

Stops a background patch job by finding and killing the associated process ID (PID).

Example 34: Verifying Patch Completion Without Logs

```
dbaascli grid status --targetVersion 19.12.0.0.0
```

Verifies if the patch to version 19.12.0.0.0 has been successfully applied by checking the current Grid Infrastructure version status.

dbaascli dbHome patch

To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbHome patch

{
    --oracleHome <value>
    | --oracleHomeName <value>
 }
        [--imageFilePath <value>] [--executePrereqs] [--nodes <value>]
        {
            [--resume [--sessionID <value>]]
            | [--rollback [--sessionID <value>]]
        }
[--skipDatapatch]
[--skipClosedPDBs]
[--skipPDBs <value>]
[--continueWithDbDowntime]
[--skipUnreachableNodes]
[--drainTimeoutInSeconds <value>]
[--waitForCompletion <value>]
[--nonRolling]
[--skipDatapatchForDB <value>]
```

Where:

- `--oracleHome` specifies the path of Oracle home
- `--oracleHomeName` specifies the name of Oracle home
- `--targetVersion` specifies the target version of Oracle Home specified as five numeric segments separated by periods, for example, 19.12.0.0.0.
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--continueWithDbDowntime` continues patching with database downtime. This option can be used in environments wherein there is only one active instance up and the patching operation can be continued even with a downtime.
- `--skipUnreachableNodes` skips operation on unreachable nodes
- `--nodes` specifies a comma-delimited list of nodes if patching has to be performed on a subset of nodes
- `--executePrereqs` runs prereqs
- `--skipDatapatch` skips running `datapatch` on the databases
- `--imageFilePath` specifies the absolute path of the image file to be used
- `--skipPDBs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: *cdb1:pdb1*,*cdb2:pdb2*, and so on
- `--skipClosedPdbs` skips running `datapatch` on closed PDBs
- `--rollback` rolls back patched Oracle home.
- `--waitForCompletion` specifies false to run the operation in background. Valid values : `true|false`
- `--drainTimeoutInSeconds` specifies time (in seconds) to complete the resource draining while stopping the database
- `--skipUnreachableNodes` skips operation on unreachable nodes
- `[--nonRolling]` patches the dbhome on all nodes in parallel. Note that this would cause a downtime in the database.
- `[--skipDatapatchForDB <value>]` skips running datapatch on the specified comma-delimited list of databases. Example: *db1*,*db2*,*db3*...

Frequently Asked Questions

Q: What is the dbaascli dbHome patch command used for?

A: The `dbaascli dbHome patch` command is used to patch Oracle home from one patch level to another.

Q: Do I need special permissions to run the dbaascli dbHome patch command?

A: Yes, you need to run the command as the `root` user.

Q: How do I specify the Oracle home path or name for the patch?

A: Use the `--oracleHome` option to specify the path of the Oracle home, or `--oracleHomeName` to specify the name of the Oracle home.

Q: How can I define the target version for the patch?

A: Use the `--targetVersion` option followed by the version number in the format `19.12.0.0.0`.

Q: What does the --resume option do?

A: The `--resume` option allows you to resume a previous patching session.

Q: How do I specify a particular session ID when resuming a patch?

A: Use the `--sessionID` option to specify the session ID of the patching session you want to resume.

Q: What is the --continueWithDbDowntime option used for?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `--continueWithDbDowntime` option allows patching to continue even if there is database downtime, useful in environments with only one active instance.

Q: How can I skip patching on unreachable nodes?

A: Use the `--skipUnreachableNodes` option to skip operations on nodes that are unreachable.

Q: How do I patch only specific nodes in a cluster?

A: Use the `--nodes` option followed by a comma-delimited list of node names to patch a subset of nodes.

Q: What is the --executePrereqs option for?

A: The `--executePrereqs` option runs prerequisite checks before applying the patch.

Q: How can I skip running datapatch on the databases?

A: Use the `--skipDatapatch` option to skip the datapatch process during patching.

Q: Can I specify a custom location for the database image?

A: Yes, use the `--imageLocation` option to specify a custom location for the database image.

Q: What does the --skipPDBs option do?

A: The `--skipPDBs` option allows you to skip running datapatch on a specified comma-delimited list of pluggable databases (PDBs).

Q: How can I skip datapatch on closed PDBs?

A: Use the `--skipClosedPDBs` option to skip datapatch on PDBs that are closed.

Q: What happens if I use the --rollback option?

A: The `--rollback` option will revert the Oracle home to its previous state before the patch was applied.

Q: How do I specify the Oracle home path for patching?

A: Use the `--oracleHome` option followed by the path to the Oracle home directory.

Q: How can I patch an Oracle home by its name rather than the path?

A: Use the `--oracleHomeName` option followed by the name of the Oracle home.

Q: How do I resume a patching operation if it was interrupted?

A: Use the `--resume` option along with the `--sessionID` option to resume a specific interrupted session.

Q: Can I continue the patching process if the database is down?

A: Yes, use the `--continueWithDbDowntime` option to continue patching even if the database is down.

Q: What should I do if some nodes are unreachable during the patching process?

A: Use the `--skipUnreachableNodes` option to bypass the unreachable nodes.

Q: How can I apply the patch to only certain nodes?

A: Specify the nodes you want to patch using the `--nodes` option with a comma-separated list of node names.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: How do I check prerequisites before applying the patch?

A: Use the `--executePrereqs` option to run prerequisite checks before applying the patch.

Q: What should I do if I want to avoid applying datapatch during the patching process?

A: Use the `--skipDatapatch` option to skip the datapatch step.

Q: How can I specify a different location for the database image used in the patching process?

A: Use the `--imageLocation` option to specify a custom location for the image.

Q: What if I need to skip datapatch on certain PDBs?

A: Use the `--skipPDBs` option to skip datapatch on a specified comma-delimited list of PDBs.

Q: Can I skip datapatch on PDBs that are not currently open?

A: Yes, use the `--skipClosedPDBs` option to skip datapatch on closed PDBs.

Q: What should I do if the patching fails midway?

A: You can use the `--rollback` option to revert to the previous state or try resuming the patching process with the `--resume` option.

Q: How can I check if all prerequisites are met before applying the patch?

A: Run the patch command with the `--executePrereqs` option to ensure all prerequisites are satisfied.

Q: What if the patching operation does not complete successfully and I need to retry?

A: Use the `--resume` option to retry the patching operation from where it left off. If needed, you can specify a `--sessionID` to resume a specific session.

Q: How can I verify if the patch was successfully applied?

A: You can verify the patching process by checking the Oracle home version using the `opatch lsinventory` command after the patch is completed.

Q: Can I run the patching command in a dry run mode to preview actions?

A: No, the `dbaascli dbHome patch` command does not have a dry-run feature. However, you can use the `--executePrereqs` option to run prerequisite checks before actually applying the patch.

Q: Is it possible to apply multiple patches in one run?

A: The `dbaascli dbHome patch` command only allows for one target version at a time. You would need to run the command separately for each patch version.

Q: How do I handle patching if the environment uses multiple Oracle homes?

A: You can specify the Oracle home you want to patch using either the `--oracleHome` or `--oracleHomeName` options, depending on whether you're specifying the path or the name of the Oracle home.

Q: Can I skip both PDB and CDB datapatching in one command?

A: Yes, you can combine the `--skipPDBs` and `--skipDatapatch` options to skip datapatching for both PDBs and the CDB in a single patch run.

Q: Can I apply a patch and rollback immediately if it causes issues?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Yes, after applying a patch, you can use the `--rollback` option to revert to the previous patch level if any issues arise.

Q: Can I patch multiple Oracle homes simultaneously?

A: No, you need to run the `dbaascli dbHome patch` command individually for each Oracle home.

Q: How do I track the progress of the patching operation?

A: During the patching process, the command provides output messages that show the progress. You can also check the log files for detailed information.

Q: Can I run patching in parallel on a clustered environment?

A: Patching operations can be applied to a subset of nodes using the `--nodes` option. However, simultaneous patching should be handled carefully, and you should ensure no overlapping sessions.

Q: How can I identify which patches are available for my Oracle home?

A: You can check the available patches via the Oracle support portal or by running the `opatch lsinventory` command to see the current patches applied to your Oracle home.

Q: Can I specify a timeout for draining resources when stopping the database during patching?

A: Yes, you can use the `--drainTimeoutInSeconds` option to specify the time in seconds for resource draining when stopping the database.

Q: What happens if the patch fails on one of the nodes in a multi-node environment?

A: You can use the `--skipUnreachableNodes` option to skip the failed node and continue the patching process on the remaining nodes. You can then address the issue on the failed node separately.

Q: How can I make the patching process run in the background?

A: Use the `--waitForCompletion` option with a value of `false` to allow the patching process to run in the background. This way, you don't need to wait for the process to complete interactively.

Q: Can I perform a rollback operation on a subset of nodes in a clustered environment?

A: Yes, you can use the `--nodes` option along with the `--rollback` option to roll back the patching on a specific set of nodes.

Q: What if I need to update the image location after starting the patch process?

A: The `--resume` option does not allow changing the image location. However, you can stop the session and start a new patch process with the updated `--imageLocation`.

Q: Is there a way to check which session IDs are available for resuming a patch?

A: You can check the log files or use Oracle Cloud tools to identify active or paused patching sessions and their session IDs.

Q: Can I limit the downtime during patching?

A: If you need to limit downtime, use the `--continueWithDbDowntime` option carefully. This allows you to proceed even when downtime is expected but requires planning for minimal service impact.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example Use Cases

Example 1: Basic Oracle Home Patching by Oracle Home Path

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0
```

Patches the Oracle home located at `/u01/app/oracle/product/19.0.0/dbhome_1` to version 19.12.0.0.0.

Example 2: Patching by Oracle Home Name

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0
```

Patches Oracle home named `DB_HOME_NAME` to version 19.12.0.0.0.

Example 3: Resuming a Previous Patch Operation

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
resume
```

Resumes the previous patching operation for the Oracle home located at `/u01/app/oracle/product/19.0.0/dbhome_1`.

Example 4: Resuming a Patch with a Specific Session ID

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --resume --sessionID 12345
```

Resumes the patching operation for Oracle home DB_HOME_NAME using session ID 12345.

Example 5: Patching with Database Downtime Allowed

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --continueWithDbDowntime
```

Patches the Oracle home located at `/u01/app/oracle/product/19.0.0/dbhome_1` to version 19.12.0.0.0 while allowing database downtime.

Example 6: Skipping Unreachable Nodes

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --skipUnreachableNodes
```

Patches Oracle home to version 19.12.0.0.0 while skipping any unreachable nodes.

Example 7: Patching a Subset of Nodes

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --nodes node1,node2
```

Patches Oracle home to version 19.12.0.0.0 on node1 and node2 only.

Example 8: Running Prerequisite Checks Before Patching

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0
--executePrereqs
```

Patches Oracle home `DB_HOME_NAME` to version 19.12.0.0.0 after running prerequisite checks.

Example 9: Skipping the Datapatch Step

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --skipDatapatch
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Patches Oracle home to version 19.12.0.0.0 without running datapatch on the databases.

Example 10: Using an Image File for Patching

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --imageFilePath /path/to/image/file.zip
```

Patches Oracle home to version 19.12.0.0.0 using an image file located at `/path/to/image/file.zip`.

Example 11: Skipping Specific PDBs During Datapatch

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --skipPDBs cdb1:pdb1,cdb2:pdb2
```

Patches Oracle home to version 19.12.0.0.0 and skips running datapatch on the specified PDBs (pdb1 in cdb1 and pdb2 in cdb2).

Example 12: Skipping Datapatch on Closed PDBs

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --skipClosedPDBs
```

Patches Oracle home to version 19.12.0.0.0 while skipping running datapatch on any closed PDBs.

Example 13: Rolling Back Oracle Home

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --rollback
```

Rolls back the last applied patch on the Oracle home named `DB_HOME_NAME`.

Example 14: Combination of Patching with Prerequisite Checks and Specific Nodes

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --executePrereqs --nodes node1,node2
```

Patches Oracle home to version 19.12.0.0.0, runs prerequisite checks, and applies the patch only on node1 and node2.

Example 15: Skipping Unreachable Nodes and Specific PDBs

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0
--skipUnreachableNodes --skipPDBs cdb1:pdb1
```

Patches Oracle home `DB_HOME_NAME` to version 19.12.0.0.0 while skipping unreachable nodes and avoiding running datapatch on pdb1 within cdb1.

Example 16: Checking Oracle Home Version Post-Patch

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0

opatch lsinventory
```

This example shows how to check the Oracle home version after a successful patch by running `opatch lsinventory`.

Example 17: Rolling Back Patching with Specific Session ID

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
rollback --sessionID 67890
```

Rolls back the Oracle home patching for a session ID of 67890.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 18: Patching with Skipping Prerequisite Checks

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --skipPrereqs
```

Patches the Oracle home but skips the prerequisite checks before applying the patch.

Example 19: Applying a Patch to a Custom Oracle Home Image

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0
--imageLocation /custom/location/image.zip
```

Patches Oracle home using a custom image file located at /custom/location/image.zip.

Example 20: Skipping Specific Nodes and Running Prereqs

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --skipUnreachableNodes --executePrereqs
```

Skips patching unreachable nodes and runs prerequisite checks before applying the patch.

Example 21: Skipping Datapatch on All PDBs in Multiple CDBs

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --skipPDBs cdb1:pdb1,cdb2:pdb2,cdb3:pdb3
```

Patches Oracle home but skips datapatch on the specified PDBs in multiple CDBs.

Example 22: Continuing Patching with Downtime on Multiple Nodes

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.12.0.0.0
--continueWithDbDowntime --nodes node3,node4
```

Continues patching on node3 and node4 with database downtime allowed.

Example 23: Skipping Datapatch on PDBs and Closed PDBs

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0 --skipDatapatch --skipClosedPDBs
```

Patches Oracle home while skipping both the datapatch and closed PDBs.

Example 24: Rolling Back and Reapplying Patch

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
rollback

dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.12.0.0.0
```

Rolls back the current patch and then reapplies the patch to the Oracle home.

Example 25: Skipping Datapatch and Allowing Downtime on a Specific Node

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.13.0.0.0 --skipDatapatch --continueWithDbDowntime --nodes node1
```

Patches Oracle home to version 19.13.0.0.0 on node1, skipping the datapatch step and allowing downtime.

Example 26: Specifying Drain Timeout During Database Shutdown

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --targetVersion 19.13.0.0.0
--drainTimeoutInSeconds 300
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Patches the Oracle home `DB_HOME_NAME` to version 19.13.0.0.0 and allows a 5-minute timeout (300 seconds) for draining resources during shutdown.

Example 27: Running Patching in the Background

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.13.0.0.0 --waitForCompletion false
```

Patches Oracle home to version 19.13.0.0.0 and runs the patching process in the background without waiting for completion.

Example 28: Rolling Back Patch on a Subset of Nodes

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
rollback --nodes node1,node2
```

Rolls back the last applied patch on node1 and node2 only for the specified Oracle home.

Example 29: Skipping Prerequisites and Patching Multiple Nodes

```
dbaascli dbHome patch --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
targetVersion 19.13.0.0.0 --skipPrereqs --nodes node3,node4
```

Patches Oracle home to version 19.13.0.0.0 on node3 and node4 without running prerequisite checks.

Example 30: Rolling Back Patch and Skipping Unreachable Nodes

```
dbaascli dbHome patch --oracleHomeName DB_HOME_NAME --rollback --
skipUnreachableNodes
```

Rolls back the last patch on Oracle home `DB_HOME_NAME` and skips unreachable nodes during the rollback process.

dbaascli database move

To move the database from one home to another, use the `dbaascli database move` command.

Prerequisites

- Before performing a move operation, ensure that all of the database instances associated with the database are up and running.
- Run the command as the `root` user.

Syntax

```
dbaascli database move
{
  --oracleHome <value> | --oracleHomeName <value>
}
--dbname <value>
[--executePrereqs]
[--resume [--sessionID <value>]]
[--rollback [--sessionID <value>]]
[--skipDatapatch]
[--skipPDBs <value>]
[--skipClosedPDBs]
[--continueWithDbDowntime]
[--allowParallelDBMove]
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
[--waitForCompletion <value>]
[--nodeList <value>]
```

Where:

- `--oracleHome` specifies Oracle home path
- `--oracleHomeName` specifies the name of Oracle home
- `--dbname` specifies the name of the database
- `--executePrereqs` runs the prerequisite checks and report the results
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--rollback` rolls the database back to previous home
  - `--sessionID` specifies to resume a specific session ID
- `--skipDatapatch` skips running the datapatch on the databases
- `--skipPdbs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: *pdb1*,*pdb2*...
- `--skipClosedPDBs` skips patching closed PDBs
- `--continueWithDbDowntime` continues patching with database downtime. This option can be used in environments wherein there is only one active instance up and the patching operation can be continued even with a downtime.
- `--allowParallelDBMove` allows database move in parallel.
- `--waitForCompletion` specifies `false` to run the operation in the background. **Valid values:** `true|false`
- `--nodeList` specifies a comma-delimited list of nodes if operation has to be performed on a subset of nodes

Frequently Asked Questions

Q: What is the dbaascli database move command used for?

A: The `dbaascli database move` command is used to move a database from one Oracle home to another.

Q: What are the prerequisites for using the dbaascli database move command?

A: Before performing a move operation, ensure that all database instances associated with the database are up and running. Additionally, the command must be run as the `root` user.

Q: What does the --oracleHome parameter specify?

A: The `--oracleHome` parameter specifies the path of the Oracle home to which the database will be moved.

Q: What does the --oracleHomeName parameter specify?

A: The `--oracleHomeName` parameter specifies the name of the Oracle home to which the database will be moved.

Q: What is the purpose of the --dbname parameter?

A: The `--dbname` parameter specifies the name of the database that you want to move.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What does the --executePrereqs option do?

A: The `--executePrereqs` option runs the prerequisite checks and reports the results.

Q: What is the --resume option used for?

A: The `--resume` option resumes a previously interrupted or incomplete move operation.

Q: How is --sessionID used in the command?

A: The `--sessionID` specifies the session ID to resume a previous run or rollback.

Q: What does the --rollback option do?

A: The `--rollback` option rolls the database back to its previous Oracle home.

Q: What does the --skipDatapatch option do?

A: The `--skipDatapatch` option skips running the datapatch on the databases during the move operation.

Q: What is the function of the --skipPDBs option?

A: The `--skipPDBs` option skips running the datapatch on a specified comma-delimited list of PDBs (e.g., pdb1, pdb2).

Q: What does the --skipClosedPDBs option do?

A: The `--skipClosedPDBs` option skips patching of closed PDBs.

Q: What does --continueWithDbDowntime mean?

A: The `--continueWithDbDowntime` option allows the move operation to proceed even if there is only one active instance up, allowing for downtime during the process.

Q: What is the purpose of the --allowParallelDBMove option?

A: The `--allowParallelDBMove` option allows the database move to be performed in parallel, potentially speeding up the process.

Q: What does --waitForCompletion specify?

A: The `--waitForCompletion` option specifies whether to wait for the operation to complete. Setting it to false runs the operation in the background.

Q: How do I use the --nodeList parameter?

A: The `--nodeList` parameter specifies a comma-delimited list of nodes on which the move operation will be performed, if it is not to be applied to all nodes.

Q: What should I do if I encounter issues with the dbaascli database move command?

A: Ensure all database instances are running and verify that you are running the command as the `root` user. If issues persist, consult the detailed command documentation or open a support ticket with Oracle.

Q: Can I perform a move operation if one of the database instances is down?

A: No, all associated database instances must be up and running before performing the move operation.

Q: What happens if the move operation is interrupted?

A: You can use the `--resume` option to continue the move operation from where it left off by using the same session or specifying the `--sessionID`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What does the --allowParallelDBMove option do?

A: It enables the database move to be performed in parallel, which can reduce the time it takes to complete the operation, especially on larger environments.

Q: How do I monitor the progress of a move operation that is running in the background?

A: When using `--waitForCompletion false`, the command does not wait for the operation to complete. You can check the status of the operation manually using appropriate Oracle logs or status commands.

Q: What is the significance of the --skipClosedPDBs option?

A: It skips patching for PDBs that are closed, reducing the operation time if there are PDBs that don't need to be patched.

Q: Can the database move be rolled back at any time?

A: Yes, the move operation can be rolled back using the `--rollback` option, either by specifying the session ID or simply rolling back to the previous Oracle home.

Q: What is the role of --nodeList in a multi-node environment?

A: In a multi-node environment, you can restrict the move operation to specific nodes by providing a comma-delimited list of node names with `--nodeList`.

Q: Can I move the database to a new Oracle home while skipping specific nodes in a multi-node environment?

A: Yes, you can use the `--nodeList` option to specify which nodes to include in the move operation. Any nodes not listed will be skipped.

Q: What is the maximum number of nodes I can specify with the --nodeList parameter?

A: The `--nodeList` parameter allows you to specify a comma-delimited list of as many nodes as needed, limited only by your environment configuration. Ensure all nodes are valid and reachable.

Q: How do I know which PDBs are closed before using the --skipClosedPDBs option?

A: You can query the `v$pdbs` view to check the status of your PDBs. Any PDBs with a status of "`MOUNTED`" or "`CLOSED`" will be skipped when using `--skipClosedPDBs`.

Q: How do I verify if a rollback has completed successfully?

A: After running the rollback command, you can review the database logs or use the Oracle alert logs to verify that the database has been successfully rolled back to its previous Oracle home.

Q: Is there a way to force the move operation if some prerequisites fail?

A: The move command enforces prerequisite checks for system stability. You cannot bypass critical prerequisite failures. Address any issues reported by the `--executePrereqs` option before proceeding with the move.

Example Use Cases

Example 1: Basic Database Move by Oracle Home Path

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Moves the database *ORCL* to the Oracle home located at `/u01/app/oracle/product/19.0.0/dbhome_1`.

Example 2: Database Move by Oracle Home Name

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL
```

Moves the database *ORCL* to the Oracle home named `DB_HOME_NAME`.

Example 3: Running Prerequisite Checks Before Moving

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --executePrereqs
```

Moves the database *ORCL* to the Oracle home while running the prerequisite checks beforehand.

Example 4: Resuming a Previous Move Operation

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --resume
```

Resumes a previous move operation for the *ORCL* database.

Example 5: Resuming a Move Operation with a Specific Session ID

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL --resume --
sessionID 12345
```

Resumes the move operation for the ORCL database using session ID 12345.

Example 6: Rolling Back a Move Operation

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --rollback
```

Rolls back the move operation for the ORCL database, restoring it to the previous Oracle home.

Example 7: Rolling Back a Move Operation with a Session ID

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL --rollback --
sessionID 67890
```

Rolls back the move operation for ORCL using session ID 67890.

Example 8: Skipping Datapatch

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --skipDatapatch
```

Moves the database *ORCL* without running datapatch on the databases.

Example 9: Skipping Specific PDBs During Datapatch

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --skipPDBs pdb1,pdb2
```

Moves the *ORCL* database to a new Oracle home but skips running datapatch on the specified PDBs (pdb1 and pdb2).

Example 10: Skipping Datapatch on Closed PDBs

Copyright © 2008, 2026, Oracle and/or its affiliates.

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --skipClosedPDBs
```

Moves the *ORCL* database and skips running datapatch on any closed PDBs.

Example 11: Allowing Database Downtime During Move

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL --
continueWithDbDowntime
```

Moves the *ORCL* database to the specified Oracle home while allowing downtime of the database during the move process.

Example 12: Moving Database in Parallel

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --allowParallelDBMove
```

Moves the database *ORCL* to the specified Oracle home with the option to run the move in parallel for better performance.

Example 13: Running the Operation in the Background

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --waitForCompletion false
```

Moves the database *ORCL* to a new Oracle home but runs the operation in the background.

Example 14: Specifying Nodes for the Move

dbaascli database move `--oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --dbname ORCL --nodeList` *node1,node2*

Moves the database *ORCL* to the specified Oracle home but performs the operation only on node1 and node2.

Example 15: Combination of Move with Prerequisite Checks, Skipping Specific PDBs, and Allowing Downtime

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname ORCL --
executePrereqs --skipPDBs pdb1 --continueWithDbDowntime
```

Moves the *ORCL* database to the specified Oracle home, runs prerequisite checks, skips running datapatch on pdb1, and allows database downtime during the operation.

Example 16: Combination of Move in Parallel and Running in the Background

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --allowParallelDBMove --waitForCompletion false
```

Moves the *ORCL* database to a new Oracle home, runs the move in parallel, and executes the operation in the background.

Example 17: Combining Move with Parallel Execution and Skipping Closed PDBs

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname TESTDB --allowParallelDBMove --skipClosedPDBs
```

Moves the *TESTDB* database to the new Oracle home `/u02/app/oracle/product/19.0.0/dbhome_2`, while running the operation in parallel and skipping datapatch on closed PDBs.

Example 18: Running Prerequisites Check Only

Copyright © 2008, 2026, Oracle and/or its affiliates.

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname PRODDB --executePrereqs
```

Checks the prerequisites for moving the *PRODDB* database to the Oracle home located at `/u02/app/oracle/product/19.0.0/dbhome_2` without actually performing the move.

Example 19: Skipping Datapatch for Specific PDBs

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname HRDB --skipPDBs pdb1,pdb3
```

Moves the *HRDB* database to the new Oracle home, but skips running datapatch for pdb1 and pdb3.

Example 20: Running the Move on Specific Nodes

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname FINDB --nodeList node1,node3
```

Moves the *FINDB* database to the new Oracle home only on node1 and node3.

Example 21: Database Move with Downtime Allowed

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname ORCL --continueWithDbDowntime
```

Moves the *ORCL* database to the specified Oracle home while allowing downtime during the move operation.

Example 22: Combination of Parallel Move and Skipping Datapatch

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname CRMDB --allowParallelDBMove --skipDatapatch
```

Moves the *CRMDB* database in parallel, skipping the datapatch process.

Example 23: Move Operation in Background with a Node List

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname SALESDB --waitForCompletion false --nodeList node2,node3
```

Moves the *SALESDB* database to the specified Oracle home in the background, and the operation is applied only on node2 and node3.

Example 24: Database Move with Prerequisite Check and Allowing Parallel Move

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_2 --
dbname ORCL --executePrereqs --allowParallelDBMove
```

Moves the *ORCL* database to the new Oracle home after performing the prerequisite checks and running the move operation in parallel.

Example 25: Rolling Back a Move Operation and Skipping Closed PDBs

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname DEVDB --rollback --skipClosedPDBs
```

Rolls back the move operation for the *DEVDB* database, skipping any closed PDBs.

Example 26: Moving the Database with Specific Downtime and Parallel Execution

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname FINDB --allowParallelDBMove --continueWithDbDowntime
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Moves the *FINDB* database to the specified Oracle home while allowing database downtime and enabling parallel execution to speed up the process.

Example 27: Checking Database Move Prerequisites without Executing the Move

```
dbaascli database move --oracleHome /u01/app/oracle/product/19.0.0/dbhome_1 --
dbname HRDB --executePrereqs
```

Runs prerequisite checks to validate that the *HRDB* database can be moved to the specified Oracle home without executing the move itself.

Example 28: Moving Database and Running the Command in the Background on Specific Nodes

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_3 --
dbname PRODDB --waitForCompletion false --nodeList node1,node4
```

Moves the *PRODDB* database to a new Oracle home, executing the operation in the background, and applying it only on node1 and node4.

Example 29: Combining Prerequisite Checks, Skipping Closed PDBs, and Allowing Parallel Execution

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_2 --
dbname CRMDB --executePrereqs --skipClosedPDBs --allowParallelDBMove
```

Performs prerequisite checks before moving the *CRMDB* database to the new Oracle home, skips patching closed PDBs, and allows the operation to run in parallel for faster execution.

Example 30: Database Move with Rollback on Specific Session ID and Skipping Datapatch

```
dbaascli database move --oracleHomeName DB_HOME_NAME --dbname DEVDB --rollback --
sessionID 45678 --skipDatapatch
```

Rolls back a previously executed move operation for the *DEVDB* database to its prior Oracle home using session ID 45678, skipping the datapatch process during the rollback.

Example 31: Moving Database with Allow Parallel Execution and Specifying Datapatch Skipping for PDBs

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0/dbhome_3 --
dbname ANALYTICDB --allowParallelDBMove --skipPDBs pdb2,pdb4
```

Moves the *ANALYTICDB* database in parallel to the specified Oracle home and skips the datapatch process for pdb2 and pdb4.

Miscellaneous

Q: How can I skip the catbundle run while patching Oracle Database 11.2.0.4.0?

A: To skip the `catbundle` run during the Oracle Database patching process, use the `--skipDatapatch` option with the `dbaascli database move` or `dbaascli dbHome patch` command.

Q: What are the best practices to follow during Oracle Database patching?

A: Oracle recommends performing out-of-place patching using the `dbaascli database move` command to minimize the patching window.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Oracle advises using the `--allowParallelDBMove` option to enable parallel patching, which can expedite the process.

Q: Can warnings reported during Oracle Database patch prerequisites be ignored?

A: It is advisable to address and resolve any warnings reported during the prerequisites check before proceeding with the patching process. Ignoring warnings may lead to issues during the actual patching.

Q: How can I continue with Oracle Database patching if only one database instance is up and running?

A: It is recommended to have at least two instances running to avoid database downtime. If running two instances is not possible, you can use the `--continueWithDbDowntime` option with the `dbaascli database move` or `dbaascli dbHome patch` command to proceed with patching despite the downtime.

Q: In a Data Guard environment, does datapatch get executed on both the primary and standby databases?

A: No, in a Data Guard environment, datapatch is executed only as part of the primary database patching process.

Q: Can Interim Software Updates (one-off patches) or individual patches be applied manually on Oracle Homes in Exadata Cloud@Customer (ExaDB-C@C) environments?

A: Yes, one-off patches or individual patches can be applied manually to Oracle Homes in ExaDB-C@C environments. However, it is recommended to use the Oracle Database Software Image option for a more streamlined and supported patching process.

Q: How do I patch multiple Oracle databases running from the same Oracle Home when each database is running on just one node?

A: Use the `dbaascli dbHome patch` command to patch the specified Oracle Home, which will apply the patch to all databases running from that home. It is recommended to have multiple instances running to avoid downtime. If running multiple instances is not possible, you can use the `--continueWithDbDowntime` option to proceed with patching despite the downtime.

Pluggable Database (PDB) Management

This section covers the management of Pluggable Databases (PDBs) within a container database (CDB). It includes commands for creating (`dbaascli pdb create`), deleting (`dbaascli pdb delete`), and cloning PDBs (`dbaascli pdb localClone`, `dbaascli pdb remoteClone`). You can manage PDB states with commands to open, close, or bounce PDBs, and retrieve connection details (`dbaascli pdb getConnectString`). Additional commands support backup, recovery, and relocating PDBs, ensuring comprehensive control over PDB lifecycle and operations.

- 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 create
  To create a new pluggable database (PDB), use the `dbaascli pdb create` command.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- dbaascli pdb delete
  To delete a pluggable database (PDB) run the `dbaascli pdb delete` command.
- dbaascli pdb getConnectString
  To display Oracle Net connect string information for a pluggable database (PDB) run the `dbaascli pdb getConnectString` command.
- dbaascli pdb getDetails
  To view details of a pluggable database (PDB), use the `dbaascli pdb getDetails` command.
- dbaascli pdb list
  To view the list of pluggable databases (PDB) in a container database, use the `dbaascli pdb list` command.
- dbaascli pdb localClone
  To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the `dbaascli pdb localClone` command.
- dbaascli pdb open
  To open a pluggable database (PDB), use the `dbaascli pdb open` command.
- dbaascli pdb recover
  To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.
- dbaascli pdb refresh
  To refresh a specified pluggable database (PDB), use the `dbaascli pdb refresh` command.
- dbaascli pdb 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 pdb relocate
  To relocate the specified PDB from the remote database into local database, use the `dbaascli pdb relocate` command.

dbaascli pdb backup

To backup a pluggable database (PDB), query PDB backups, and delete a PDB backup, use the `dbaascli pdb backup` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli pdb backup --pdbName <value> --dbname <value>
        {
            --start
                {
                    [--level1]
                    | [--archival --tag <value>]
                }
            | --delete --backupTag <value>
            | --status --uuid <value>
            | --getBackupReport --json <value> --tag <value>
            | --list [--json <value>]
        }
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Where:

```
--pdbName: PDB name.
--dbname: Oracle Database name.
--start | --delete | --status | --getBackupReport | --list
--start: Begins PDB backup.
     [--level1 | --archival]
     [--level1: Creates a Level-1 (incremental) backup.]
     [--archival: Creates an archival full backup.]
         --tag: Specify backup tag.
--delete: Deletes archival backup.
         --backupTag: Specify backup tag to delete.
--status
         --uuid <value>
--getBackupReport: Returns backup report.
        --json: Specify the file name for JSON output.
        --tag: Specify backup tag.
--list: Returns PDB backup information.
        [--json: Specify the file name for JSON output.]
```

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb backup command?

A: The `dbaascli pdb backup` command is used to create backups for a pluggable database (PDB), query backup status, generate backup reports, and delete PDB backups in an Exadata Cloud@Customer environment.

Q: What is the prerequisite for using the dbaascli pdb backup command?

A: The command must be run as the `root` user, and you need to be connected to an Exadata Cloud@Customer virtual machine.

Q: How do I start a PDB backup using the dbaascli pdb backup command?

A: You can start a PDB backup using the `--start` option. For example:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --start
```

Q: What options can be used with the --start flag?

A: With the `--start` flag, you can specify:

`--level1` for a Level-1 incremental backup

`--archival` for a full archival backup (which also requires a `--tag` to specify the backup tag)

Q: How do I create a Level-1 incremental PDB backup?

A: Use the `--level1` flag with the `--start` option to create a Level-1 incremental backup:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --start --level1
```

Q: How do I create an archival PDB backup?

A: Use the `--archival` flag with the `--start` option and specify a backup tag using `--tag`:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --start --archival --
tag <backup_tag>
```

Q: How do I delete a specific PDB backup?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: To delete a specific backup, use the `--delete` flag and specify the backup tag using `--backupTag`:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --delete --backupTag
<backup_tag>
```

Q: How can I check the status of a PDB backup?

A: Use the `--status` flag along with the backup `--uuid` to check the status of a specific backup:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --status --uuid
<backup_uuid>
```

Q: How do I retrieve a PDB backup report in JSON format?

A: To get a backup report in JSON format, use the `--getBackupReport` option, specify the file name with `--json`, and provide the backup tag with `--tag`:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --getBackupReport --
json <file_name> --tag <backup_tag>
```

Q: How can I list all PDB backups for a specific PDB?

A: Use the `--list` option to get a list of all backups for a given PDB:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --list
```

Optionally, you can output the list in JSON format using the `--json` flag:

```
dbaascli pdb backup --pdbName <PDB_Name> --dbname <DB_Name> --list --json
<file_name>
```

Q: What does the --pdbName option do?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to back up, query, or delete backups for.

Q: What is the purpose of the --dbname option?

A: The `--dbname` option specifies the name of the Oracle Database to which the PDB belongs.

Q: How do I specify a backup tag for a PDB backup?

A: You specify a backup tag using the `--tag` option when starting an archival backup or when retrieving a backup report:

```
--tag <backup_tag>
```

Q: Can I run PDB backups in JSON mode?

A: Yes, both the backup report (`--getBackupReport`) and backup listing (`--list`) options support output in JSON format. You specify a JSON file name using the `--json` option.

Example 7-30 Examples

- To take level1 backup for a PDB *pdb1* in a CDB *myTestDb*:

  ```
  dbaascli pdb backup --dbname myTestDb --pdbName pdb1 --start --level1
  ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

- 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 Cloud@Customer 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`.

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb bounce command?

A: The `dbaascli pdb bounce` command is used to bounce (restart) a pluggable database (PDB) in an Exadata Cloud@Customer environment.

Q: What are the prerequisites for using the dbaascli pdb bounce command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

Q: How do I bounce a PDB by specifying its name?

A: To bounce a PDB by specifying its name, use the following syntax:

Copyright © 2008, 2026, Oracle and/or its affiliates.

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbName <PDB_Name>
```

Q: How do I bounce a PDB by using its unique identifier (UID)?

A: To bounce a PDB using its unique identifier (UID), use the following syntax:

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbUID <PDB_UID>
```

Q: What is the --dbname option used for?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the pluggable database (PDB) being bounced.

Q: What is the --pdbName option used for?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to bounce.

Q: What is the --pdbUID option used for?

A: The `--pdbUID` option specifies the unique identifier (UID) of the pluggable database (PDB) that you want to bounce.

Q: How do I specify the target open mode for the PDB when bouncing it?

A: You can use the `--openMode` option to specify the desired open mode for the PDB after bouncing. The valid values are `READ_WRITE` and `READ_ONLY`. For example:

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbName <PDB_Name> --openMode
READ_WRITE
```

Q: Can I open the PDB in read-only mode after bouncing it?

A: Yes, you can use the `--openMode READ_ONLY` option to open the PDB in read-only mode after bouncing:

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbName <PDB_Name> --openMode READ_ONLY
```

Q: What is the default open mode if --openMode is not specified?

A: If `--openMode` is not specified, the PDB will be opened in its default open mode, which is typically `READ_WRITE`.

Q: Can I use both --pdbName and --pdbUID in the same command?

A: No, you should specify either `--pdbName` or `--pdbUID`, but not both in the same command.

Q: How can I restart a PDB and ensure it opens in read-write mode?

A: To restart a PDB and ensure it opens in read-write mode, use the `--openMode READ_WRITE` option:

```
dbaascli pdb bounce --dbname <CDB_Name> --pdbName <PDB_Name> --openMode
READ_WRITE
```

Q: Is it mandatory to specify the open mode when using the dbaascli pdb bounce command?

A: No, specifying the `--openMode` is optional. If not provided, the PDB will open in its default mode.

Q: What happens if I don't specify the --openMode flag?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: If the `--openMode` flag is not specified, the PDB will be opened in its default mode, which is typically `READ_WRITE`.

Example 7-31 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.

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb close command?

A: The `dbaascli pdb close` command is used to close a pluggable database (PDB) in an Exadata Cloud@Customer environment.

Q: What are the prerequisites for using the dbaascli pdb close command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: How do I close a PDB by specifying its name?

A: To close a PDB by specifying its name, use the following syntax:

```
dbaascli pdb close --dbname <CDB_Name> --pdbName <PDB_Name>
```

Q: How do I close a PDB by specifying its unique identifier (UID)?

A: To close a PDB by using its unique identifier (UID), use the following syntax:

```
dbaascli pdb close --dbname <CDB_Name> --pdbUID <PDB_UID>
```

Q: What does the --dbname option do in the dbaascli pdb close command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the pluggable database (PDB) you want to close.

Q: What does the --pdbName option do in the dbaascli pdb close command?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to close.

Q: What is the purpose of the --pdbUID option in the dbaascli pdb close command?

A: The `--pdbUID` option allows you to specify the unique identifier (UID) of the pluggable database (PDB) that you want to close.

Q: Can I close the PDB on a specific instance of the CDB?

A: No, upon successful completion, the PDB is closed on all instances of the container database (CDB).

Q: Is it possible to specify both --pdbName and --pdbUID in the same command?

A: No, you can specify either `--pdbName` or `--pdbUID`, but not both in the same command.

Q: What happens when the dbaascli pdb close command completes successfully?

A: When the command completes successfully, the pluggable database (PDB) is closed on all instances of the container database (CDB).

Q: How do I close a specific PDB within a CDB using its UID?

A: You can close a specific PDB using its UID by running:

```
dbaascli pdb close --dbname <CDB_Name> --pdbUID <PDB_UID>
```

Q: What happens if I forget to specify either --pdbName or --pdbUID?

A: You must specify either `--pdbName` or `--pdbUID` in the command. If neither is provided, the command will not run.

Q: Can I use the dbaascli pdb close command for a CDB directly?

A: No, the command is designed to close a pluggable database (PDB) within a container database (CDB), not the CDB itself.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-32 dbaascli pdb close

```
dbaascli pdb close --dbname cdb name --pdbName pdb name associated with the
CDB


dbaascli pdb close --dbname cdb name --pdbUID con_uid of that pdb
```

dbaascli pdb create

To create a new pluggable database (PDB), use the `dbaascli pdb create` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb create --pdbName <value> --dbName <value> [--maxCPU <value>] [--
maxSize <value>] [--pdbAdminUserName <value>] [--lockPDBAdminAccount <value>]
[--resume [--sessionID <value>]] [--executePrereqs] [--waitForCompletion
<value>]
        {
            [--blobLocation <value>]
            | [--standbyBlobFromPrimary <value>]
        }
        [--pdbTdeKeyVersionOCID <value>]
```

Where:

- `--pdbName` specifies the name of the new PDB that you want to create
- `--dbName` specifies the name of the container database that hosts the new PDB
- `--maxCPU` optionally specifies the maximum number of CPUs that are available to the PDB. Setting this option is effectively the same as setting the `CPU_COUNT` parameter in the PDB
- `--maxSize` optionally specifies the maximum total size of data files and temporary files for tablespaces belonging to the PDB. Setting this option is effectively the same as setting the `MAXSIZE PDB` storage clause in the `CREATE PLUGGABLE DATABASE` SQL command. You can impose a limit by specifying an integer followed by a size unit (`K`, `M`, `G`, or `T`), or you can specify `UNLIMITED` to explicitly enforce no limit
- `--pdbAdminUserName` specifies the new PDB admin user name
- `--lockPDBAdminAccount` specifies `true` or `false` to lock the PDB admin user account. Default value is `true`.
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--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.

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb create command?

A: The `dbaascli pdb create` command is used to create a new pluggable database (PDB) in a container database (CDB) in an Exadata Cloud@Customer environment.

Q: What are the prerequisites for using the dbaascli pdb create command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

Q: What does the --pdbName option do in the dbaascli pdb create command?

A: The `--pdbName` option specifies the name of the new pluggable database (PDB) you want to create.

Q: What does the --dbName option do in the dbaascli pdb create command?

A: The `--dbName` option specifies the name of the container database (CDB) that will host the new pluggable database (PDB).

Q: Can I limit the CPU resources for the new PDB?

A: Yes, you can use the `--maxCPU` option to specify the maximum number of CPUs that the PDB can use. This is equivalent to setting the `CPU_COUNT` parameter in the PDB.

Q: How can I limit the storage size of a PDB?

A: You can use the `--maxSize` option to specify the maximum total size of data files and temporary files for the PDB. You can either set a size limit (in K, M, G, or T) or specify `UNLIMITED` for no limit.

Q: What is the --pdbAdminUserName option used for?

A: The `--pdbAdminUserName` option specifies the name of the admin user for the new PDB. This user will have administrative privileges within the PDB.

Q: Is it possible to lock the admin user account during PDB creation?

A: Yes, you can use the `--lockPDBAdminAccount` option to specify whether the PDB admin account should be locked. The default value is `true` (locked).

Q: What does the --resume option do in the dbaascli pdb create command?

A: The `--resume` option allows you to resume a previously failed PDB creation process.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: How do I specify a session ID for resuming a previous run?

A: You can specify a session ID using the `--sessionID` option to resume a specific session of the PDB creation process.

Q: What is the purpose of the --executePrereqs option?

A: The `--executePrereqs` option specifies whether to run only the prerequisite checks for PDB creation. You can set this option to `yes` or `no`.

Q: Can I run the PDB creation process in the background?

A: Yes, you can use the `--waitForCompletion` option and set it to false to run the operation in the background.

Q: What is the --standbyBlobFromPrimary option used for?

A: The `--standbyBlobFromPrimary` option specifies the location of the standby blob file, which is prepared from the primary database. This is required for standby database PDB operations.

Q: How will I be prompted for the PDB admin password during the creation process?

A: During the PDB creation process, you will be prompted to specify the administration password for the new PDB.

Q: Can I create a standby PDB using the dbaascli pdb create command?

A: Yes, if you are creating a standby PDB, you can use the `--standbyBlobFromPrimary` option to specify the location of the standby blob file from the primary database.

Q: What happens if I don't use the --maxSize option?

A: If you do not specify the `--maxSize` option, the PDB will not have a storage size limit unless otherwise defined by the CDB policies.

Q: What happens if I do not provide the --pdbAdminUserName option?

A: If you do not provide the `--pdbAdminUserName` option, the PDB will be created without a specified admin user, and you will need to manually configure the admin user after creation.

Q: Can I resume a failed PDB creation at any point in the process?

A: Yes, as long as the session has not been terminated, you can resume a failed PDB creation using the `--resume` and `--sessionID` options.

Example 7-33 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 © 2008, 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

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb delete command?

A: The `dbaascli pdb delete` command is used to delete a pluggable database (PDB) from a container database (CDB) in an Exadata Cloud@Customer environment.

Q: What are the prerequisites for running the dbaascli pdb delete command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

Q: What does the --dbName option specify in the dbaascli pdb delete command?

A: The `--dbName` option specifies the name of the container database (CDB) that hosts the PDB you want to delete.

Q: How can I specify which PDB to delete using the dbaascli pdb delete command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: You can specify the PDB to delete using either the `--pdbName` option (specifies the PDB name) or the `--pdbUID` option (specifies the PDB UID).

Q: Can I run the prerequisite checks without actually deleting the PDB?

A: Yes, you can use the `--executePrereqs` option and set it to `yes` to run only the prerequisite checks for the PDB deletion operation.

Q: How can I run the PDB deletion process in the background?

A: Use the `--waitForCompletion` option and set it to `false` to run the deletion process in the background.

Q: What does the --resume option do in the dbaascli pdb delete command?

A: The `--resume` option allows you to resume a previously failed PDB deletion process.

Q: How do I resume a specific session for a PDB deletion?

A: You can specify a session ID using the `--sessionID` option to resume a specific session for the PDB deletion process.

Q: What does the --allStandbyPrepared option do?

A: The `--allStandbyPrepared` option is used to confirm that the deletion operation has been successfully run on all standby databases before proceeding with the primary PDB deletion.

Q: What is the purpose of the --cleanupRelocatedPDB option?

A: The `--cleanupRelocatedPDB` option cleans up the source database after a PDB has been relocated, ensuring no residuals are left after the relocation.

Q: Can I delete a PDB that has already been relocated?

A: Yes, you can use the `--cleanupRelocatedPDB` option to delete a PDB that has already been relocated to a new CDB.

Q: How do I ensure that the delete operation runs successfully on standby databases?

A: Use the `--allStandbyPrepared` option to confirm that the operation has run successfully on all standby databases before proceeding.

Q: What happens if the delete process fails and needs to be resumed?

A: You can resume the delete process by using the `--resume` option, and if needed, specify the session ID with `--sessionID`.

Q: What does setting --waitForCompletion to false do?

A: Setting `--waitForCompletion` to `false` allows the delete process to run in the background, letting you continue working without waiting for the operation to finish.

Example: dbaascli pdb delete
To delete a PDB from a standard database in a non-Data Guard environment or from Standby database in Data Guard environment.

```
dbaascli pdb delete --dbName db721 --pdbName pdb1
```

To create PDB from Primary database in Data Guard environment:

```
dbaascli pdb create --dbName db721 --pdbName pdb1 --allStandbyPrepared
```


Copyright © 2008, 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

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb getConnectString command?

A: The `dbaascli pdb getConnectString` command is used to display the Oracle Net connect string information for a pluggable database (PDB) in an Exadata Cloud@Customer environment.

Q: What are the prerequisites for using the dbaascli pdb getConnectString command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

Q: How do I retrieve the connect string of a PDB by specifying its name?

A: To retrieve the connect string by specifying the PDB name, use the following syntax:

`dbaascli pdb getConnectString --dbname <CDB_Name> --pdbName <PDB_Name>`

Q: How do I retrieve the connect string of a PDB by specifying its unique identifier (UID)?

A: To retrieve the connect string using the PDB's unique identifier (UID), use the following syntax:

`dbaascli pdb getConnectString --dbname <CDB_Name> --pdbUID <PDB_UID>`

Q: What does the --dbname option do in the dbaascli pdb getConnectString command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the pluggable database (PDB) for which you want to display the connect string information.

Q: What does the --pdbName option do in the dbaascli pdb getConnectString command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `--pdbName` option specifies the name of the pluggable database (PDB) for which you want to retrieve the Oracle Net connect string information.

Q: What is the purpose of the --pdbUID option in the dbaascli pdb getConnectString command?

A: The `--pdbUID` option allows you to specify the unique identifier (UID) of the pluggable database (PDB) for which you want to display the connect string.

Q: Can I use both --pdbName and --pdbUID in the same command?

A: No, you can use either `--pdbName` or `--pdbUID`, but not both in the same command.

Q: What type of information is returned by the dbaascli pdb getConnectString command?

A: The command returns the Oracle Net connect string information for the specified pluggable database (PDB).

Q: Can I retrieve the connect string for a PDB on a specific container database instance?

A: No, the connect string is for the PDB as a whole, not for a specific instance of the container database.

Q: How can I get the connect string information if I only know the PDB's unique identifier (UID)?

A: You can retrieve the connect string using the PDB's UID by running:

```
dbaascli pdb getConnectString --dbname <CDB_Name> --pdbUID <PDB_UID>
```

Q: What happens if I don't provide either --pdbName or --pdbUID?

A: You must specify either `--pdbName` or `--pdbUID` to retrieve the connect string. The command will not run without one of these options.

Q: Is the connect string information for the PDB always the same across all instances of the CDB?

A: Yes, the connect string information is consistent for the PDB across all instances of the container database (CDB).

Example 7-34 dbaascli pdb getConnectString

```
dbaascli pdb getConnectString --dbname dbname --pdbName pdbName
```

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>
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        | --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

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb getDetails command?

A: The `dbaascli pdb getDetails` command is used to view details of a pluggable database (PDB) hosted in a container database (CDB) in an Exadata Cloud@Customer environment.

Q: What are the prerequisites for running the dbaascli pdb getDetails command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine.

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

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the PDB for which you want to view details.

Q: How do you specify the PDB for which you want to view details?

A: You can specify the PDB using either the `--pdbName` option (to provide the PDB name) or the `--pdbUID` option (to provide the PDB UID).

Q: What is the difference between --pdbName and --pdbUID?

A: The `--pdbName` option uses the name of the PDB to fetch details, whereas the `--pdbUID` option uses the unique identifier (UID) of the PDB to fetch its details.

Q: Can I use both --pdbName and --pdbUID together in the dbaascli pdb getDetails command?

A: No, you can specify either the `--pdbName` or the `--pdbUID` option to get details of the PDB, but not both at the same time.

Q: What are some use cases for the dbaascli pdb getDetails command?

A: You can use the `dbaascli pdb getDetails` command to:

- Retrieve details about a specific PDB in a CDB.
- Verify the configuration of a PDB.
- Check the status of a PDB within a CDB.

Q: How can I view details of a PDB based on its name?

A: To view details of a PDB based on its name, use the following syntax:

`dbaascli pdb getDetails --dbname <CDB_Name> --pdbName <PDB_Name>`

Q: How can I view details of a PDB based on its UID?

A: To view details of a PDB based on its UID, use the following syntax:

`dbaascli pdb getDetails --dbname <CDB_Name> --pdbUID <PDB_UID>`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: Can this command be used for multiple PDBs in one execution?

A: No, the command can be used to fetch details of one PDB at a time by specifying either its name or UID.

Example 7-35 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.

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-36 dbaascli pdb list

```
dbaascli pdb list --dbname cdb name
```

dbaascli pdb localClone

To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the `dbaascli pdb localClone` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb localClone --pdbName <value> --dbName <value> [--targetPDBName
<value>] [--powerLimit <value>] [--maxCPU <value>] [--maxSize <value>] [--
resume [--sessionID <value>]] [--executePrereqs] [--waitForCompletion <value>]
    {
      [--blobLocation <value>]
      | [--standbyBlobFromPrimary <value>]
    }
    [--excludeUserTablespaces <value>] [--excludePDBData <value>] [--
pdbAdminUserName <value>] [--lockPDBAdminAccount <value>] [--
sourcePDBServiceConvertList <value>]
    {
      [--createFromSnapshot
        {
          --snapshotName <value>
          | --snapshotUID <value>
        }
        [--copyDataFiles]]
      | [--snapshot]
    }
```

Where:

- `--pdbName` specifies the name of the new PDB that you want to clone

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--dbName` specifies the name of the database
- `--targetPDBName` specifies the name for the target PDB (new cloned PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the clone operation. Valid value is between 1 and 128
- `--maxCPU` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specifies the location of the standby blob file which is prepared from the primary database. This is required only for standby database PDB operations.

  > **Note**
  >
  > The parameters `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive.

- `--excludeUserTablespaces` option to skip user table spaces, example t1,t2,t3.
- `--excludePDBData` specify true/yes to skip user data from source PDB.
- `--pdbAdminUserName` specify new PDB admin user name.
- `--lockPDBAdminAccount` specify `true` or `false` to lock the PDB admin user account. Default value is `true`.
- `--sourcePDBServiceConvertList` specify a comma-delimited list of source to target service names which need to be converted. Syntax is `source_srv1:new_srv1,source_srv2:new_srv2`.
- `--createFromSnapshot` | `--snapshot`
  - `--createFromSnapshot` specify to create PDB from PDB snapshot `--snapshotName` | `--snapshotUID`
    - `--snapshotName` specify the snapshot name to create PDB from PDB.
    - `--snapshotUID` specify the snapshot UID to create PDB from PDB.

    `--copyDataFiles` specify this option to skip snapshot copy.
  - `--snapshot` specify to create PDB with snapshot copy.

The newly cloned PDB inherits administration passwords from the source PDB.

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb localClone command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `dbaascli pdb localClone` command is used to create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB) in an Exadata Cloud@Customer environment.

Q: What are the prerequisites for running the dbaascli pdb localClone command?

A: The command must be run as the `oracle` user, and you must be connected to an Exadata Cloud@Customer virtual machine. Additionally, the source PDB must already exist in the specified CDB.

Q: What does the --dbName option specify in the dbaascli pdb localClone command?

A: The `--dbName` option specifies the name of the container database (CDB) that hosts the source PDB from which the new PDB will be cloned.

Q: What does the --pdbName option specify in the dbaascli pdb localClone command?

A: The `--pdbName` option specifies the name of the new PDB that you want to create as a clone of the existing PDB in the same CDB.

Q: Can I clone a PDB with a different name using the dbaascli pdb localClone command?

A: Yes, you can specify a different name for the cloned PDB using the `--targetPDBName` option. If this option is not provided, the cloned PDB will inherit the name of the source PDB.

Q: What does the --resume option do in the dbaascli pdb localClone command?

A: The `--resume` option allows you to resume a previously interrupted PDB cloning operation.

Q: How do I limit the CPU resources available to the cloned PDB?

A: You can limit the CPU resources for the cloned PDB using the `--maxCPU` option, which specifies the maximum number of CPUs that will be allocated to the new PDB.

Q: Can I run the PDB cloning operation in the background?

A: Yes, you can run the operation in the background by setting the `--waitForCompletion` option to `false`. If you set it to `true`, the operation will run in the foreground and wait for completion.

Q: What is the purpose of the --maxSize option in the dbaascli pdb localClone command?

A: The `--maxSize` option specifies the maximum storage size (in GB) for the newly cloned PDB. If no size is specified, the cloned PDB inherits the same storage limits as the source PDB.

Q: Can I control the parallelism of the PDB clone operation?

A: Yes, you can control the degree of parallelism for the cloning operation using the `--powerLimit` option. This option accepts values between 1 and 128 to define the degree of parallelism.

Q: What is the --primaryDBWalletTar option used for?

A: The `--primaryDBWalletTar` option specifies the location of the primary database wallet tar file. This option is only required if the cloning operation involves standby database PDB operations.

Q: Can I run only the prerequisite checks for the cloning operation?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Yes, you can run only the prerequisite checks by using the `--executePrereqs` option and setting it to `yes`. The valid values are `yes` and `no`.

Q: What happens if the PDB cloning operation fails or is interrupted?

A: If the cloning operation fails or is interrupted, you can resume it by using the `--resume` option to continue from where the operation stopped.

Example 7-37 dbaascli pdb localClone

```
dbaascli pdb localClone --dbName db35 --pdbName PDB35 --targetPDBName
local_clone1 --maxCPU 2 --maxSize 15
```

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

Upon successful completion, the PDB is opened on all of the container database instances.

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb open command?

A: The `dbaascli pdb open` command is used to open a pluggable database (PDB) in an Oracle container database (CDB) in an Exadata Cloud@Customer environment.

Q: Who can run the dbaascli pdb open command?

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

Q: What does the --pdbName option specify in the dbaascli pdb open command?

A: The `--pdbName` option specifies the name of the PDB that you want to open.

Q: What does the --pdbUID option specify in the dbaascli pdb open command?

A: The `--pdbUID` option specifies the unique identifier (UID) of the PDB that you want to open.

Q: What does the --dbname option specify in the dbaascli pdb open command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the PDB.

Q: What is the purpose of the --openMode option?

A: The `--openMode` option specifies the mode in which the PDB will be opened. Valid values are `READ_WRITE` and `READ_ONLY`.

Q: Can I start services when opening the PDB?

A: Yes, you can use the `--startServices` option to either start all services associated with the PDB by specifying all or provide a comma-delimited list of specific services to start.

Q: What happens if I set the --waitForCompletion option to false?

A: If `--waitForCompletion` is set to `false`, the command will run in the background, and the user does not need to wait for the operation to complete. If set to `true`, the command will wait for completion before exiting.

Q: What does the --setPDBRefreshModeNone option do?

A: The `--setPDBRefreshModeNone` option converts a refreshable PDB (one that is regularly updated from a primary database) into a non-refreshable PDB.

Q: What is the function of the --skipPDBRefresh option?

A: The `--skipPDBRefresh` option allows you to skip the refresh operation when opening a refreshable PDB, preventing the PDB from syncing with the primary database at that time.

Q: What does the --pdbAdminUserName option do in the dbaascli pdb open command?

A: The `--pdbAdminUserName` option allows you to specify a new PDB admin username when opening the PDB.

Q: What happens if the dbaascli pdb open command is successful?

A: Upon successful completion, the specified PDB will be opened on all instances of the container database (CDB).

Q: Is it possible to run the dbaascli pdb open command for a refreshable PDB?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Yes, the command can be used for refreshable PDBs. The `--setPDBRefreshModeNone` option converts the PDB to non-refreshable, and the `--skipPDBRefresh` option skips the refresh operation during the opening process.

Q: What is the default open mode for a PDB if no --openMode is specified?

A: If no `--openMode` is specified, the PDB is typically opened in `READ_WRITE` mode by default.

Example 7-38 dbaascli pdb open

```
dbaascli pdb open --dbname cdb name --pdbName pdb name associated with the CDB


dbaascli pdb open --dbname cdb name --pdbUID con_uid of that pdb
```

Optional: `--openMode READ_WRITE/READ_ONLY`

dbaascli pdb recover

To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.

Prerequisite

- Run the command as the `root` user.
- Database must be configured with backup storage destination details where backups are stored.

Syntax

```
dbaascli pdb recover --dbname <value> --pdbName <value>
        {
            --start
                {
                    --untilTime <value> [--nonUTC]
                    | --untilSCN <value>
                    | --latest
                    | --tag <value>
                }
            | --status --uuid <value>
        }
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbName` specifies the name of the PDB that you want to recover
- `--start` begins the recovery of a PDB.
  - `--untilTime` recovers PDB until time. Input format: `DD-MON-YYYY HH24:MI:SS`
  - `--untilSCN` recovers PDB until SCN
  - `--latest` recovers PDB to the last known state
  - `--tag` recovers PDB to archival tag
- `--status` displays the details about a PDB recovery job process

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--uuid` unique identifier of the PDB recovery operation. Input format: `xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx`

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb recover command?

A: The `dbaascli pdb recover` command is used to recover a pluggable database (PDB) to a previous state using backups stored in a configured backup storage destination.

Q: Who can run the dbaascli pdb recover command?

A: The command must be run as the `root` user.

Q: What is required before running the dbaascli pdb recover command?

A: Before running the command, the database must be configured with backup storage destination details where backups are stored.

Q: What does the --pdbName option specify in the dbaascli pdb recover command?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to recover.

Q: What does the --dbname option specify in the dbaascli pdb recover command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the PDB.

Q: What are the possible options for starting a PDB recovery using the --start option?

A: You can recover the PDB using one of the following options:

- `--untilTime <value>`: Recovers the PDB until a specified time (format: `DD-MON-YYYY HH24:MI`).
- `--untilSCN <value>`: Recovers the PDB until a specified System Change Number (SCN).
- `--latest`: Recovers the PDB to the latest known state.
- `--tag <value>`: Recovers the PDB to a specific archival tag.

Q: What is the format required for specifying the time in the --untilTime option?

A: The time must be in the format `DD-MON-YYYY HH24:MI:SS`.

Q: How can I recover a PDB to the latest state using dbaascli pdb recover?

A: To recover the PDB to the latest known state, use the `--latest` option:

```
dbaascli pdb recover --pdbName <value> --dbname <value> --start --latest
```

Q: How do I recover a PDB to a specific archival tag?

A: You can recover the PDB to a specific tag using the `--tag` option:

```
dbaascli pdb recover --pdbName <value> --dbname <value> --start --tag <tag_value>
```

Q: Can I recover a PDB using a specific SCN?

A: Yes, you can recover the PDB to a specific SCN using the `--untilSCN` option:

```
dbaascli pdb recover --pdbName <value> --dbname <value> --start --untilSCN
<SCN_value>
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What does the --status option do in the dbaascli pdb recover command?

A: The `--status` option is used to check the status of a recovery operation. You must provide the `--uuid` to specify the recovery session.

Q: How can I check the status of a PDB recovery?

A: To check the status of a recovery operation, use the `--status` option with the `--uuid` of the recovery session:

```
dbaascli pdb recover --pdbName <value> --dbname <value> --status --uuid
<uuid_value>
```

Q: What happens if I specify the --latest option in the recovery command?

A: If you specify the `--latest` option, the PDB will be recovered to the most recent state available in the backup.

Example 7-39 Examples

- To recover a PDB *pdb1* in a CDB *myTestDb* to latest:

  ```
  dbaascli pdb recover --dbname myTestDb --pdbName pdb1 --start --latest
  ```

- To query the status of PDB recovery request submitted with `uuid` *81a17352362011ecbc3000163e8e4fac*:

  ```
  dbaascli pdb recover --dbname myTestDb --pdbName pdb1 --status --uuid
  81a17352362011ecbc3000163e8e4fac
  ```

Related Topics

- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer 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

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--waitForCompletion`: specify `false` to run the operation in the background. Valid values: `true|false`

Frequently Asked Questions

Q: What is the purpose of the dbaascli pdb refresh command?

A: The `dbaascli pdb refresh` command is used to refresh a specified pluggable database (PDB) in a container database (CDB).

Q: Who can run the dbaascli pdb refresh command?

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

Q: What does the --dbname option specify in the dbaascli pdb refresh command?

A: The `--dbname` option specifies the name of the container database (CDB) that hosts the pluggable database (PDB) to be refreshed.

Q: What does the --pdbName option specify in the dbaascli pdb refresh command?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) that you want to refresh.

Q: What does the --pdbUID option specify in the dbaascli pdb refresh command?

A: The `--pdbUID` option specifies the unique identifier (UID) of the pluggable database (PDB) that you want to refresh.

Q: What does the --waitForCompletion option do in the dbaascli pdb refresh command?

A: The `--waitForCompletion` option specifies whether the operation should be run in the foreground or background. If set to true, the operation will run in the foreground and wait for completion. If set to false, the operation will run in the background.

Q: How can I refresh a PDB and run the operation in the background?

A: To refresh a PDB and run the operation in the background, use the `--waitForCompletion false` option:

```
dbaascli pdb refresh --dbname <value> --pdbName <value> --waitForCompletion false
```

Q: How do I refresh a PDB using its unique identifier (UID)?

A: You can refresh the PDB using the `--pdbUID` option:

```
dbaascli pdb refresh --dbname <value> --pdbUID <value>
```

Q: Can I specify both --pdbName and --pdbUID together in the dbaascli pdb refresh command?

A: No, you must specify either `--pdbName` or `--pdbUID`, but not both, when refreshing a PDB.

Q: What happens if I don't include the --waitForCompletion option in the command?

A: If you don't specify the `--waitForCompletion` option, the default behavior will be to wait for the operation to complete before returning control to the user.

Q: Can I refresh a PDB while the database is running?

A: Yes, you can refresh a PDB while the database is running, as long as the command is executed by a user with appropriate privileges.

Copyright © 2008, 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 Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli pdb remoteClone

To create a new pluggable database (PDB) as a clone of an existing PDB in another container database (CDB), use the `dbaascli pdb remoteClone` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli pdb remoteClone --pdbName <value> --dbName <value> --
sourceDBConnectionString <value> [--targetPDBName <value>] [--powerLimit
<value>] [--maxCPU <value>] [--maxSize <value>] [--resume [--sessionID
<value>]] [--executePrereqs] [--waitForCompletion <value>] [--
sourcePDBExportedTDEKeyFile <value>]
        {
            [--blobLocation <value>]
            | [--standbyBlobFromPrimary <value>]
        }
        [--excludeUserTablespaces <value>] [--excludePDBData <value>] [--
pdbAdminUserName <value>] [--lockPDBAdminAccount <value>] [--
sourcePDBServiceConvertList <value>] [--upgradePDB]
        {
            [--createFromSnapshot --snapshotName <value> [--copyDataFiles]]
            | [--snapshot]
        }
        [--refreshablePDB --refreshMode <value> [--refreshIntervalInMinutes
<value>]] [--updateDBBlockCacheSize]
        {
            [--dblinkUsername <value> [--honorCaseSensitiveUserName]]
            | [--dbLinkName <value>]
        }
        [--keepSourceDBServices <value>]
```

Where:

- `--pdbName` specifies the name of the source PDB that you want to clone
- `--dbname` specifies the name (`DB_NAME`) of the CDB that hosts the newly cloned PDB
- `--sourceDBConnectionString` specifies the source database connection string in the format *`scan_name`*`:`*`scan_port/database_service_name`*
- `--targetPDBName` specifies the name for the target PDB (new cloned PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the clone operation. Valid value is between 1 and 128
- `--maxCPU` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--sourcePDBExportedTDEKeyFile` specifies the source PDB exported key file. This variable is applicable to only 12.1 database.
- `--blobLocation` specifies the custom path where the standby blob file will be generated in a Data Guard environment
- `--standbyBlobFromPrimary` specify the location of the standby blob file, which is prepared from the primary database. This is required only for standby database PDB operations

> **Note**
>
> The parameters `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive.

- `--excludeUserTablespaces` option to skip user table spaces, example *t1*,*t2*,*t3*.
- `--excludePDBData` specify `true/yes` to skip user data from source PDB.
- `--pdbAdminUserName` specifies new PDB admin user name
- `--lockPDBAdminAccount` specify `true` or `false` to lock the PDB admin user account. Default value is `true`.
- `--sourcePDBServiceConvertList` specify a comma-delimited list of source to target service names, which need to be converted. Syntax is `source_srv1:new_srv1`, `source_srv2:new_srv2`.
- `--upgradePDB` specify to upgrade the PDB as part of this operation
- `--createFromSnapshot | --snapshot`
  - `--createFromSnapshot` specify to create PDB from PDB snapshot
    - `--snapshotName` specify the snapshot name to create PDB from PDB.

    `--copyDataFiles` specify this option to skip snapshot copy.
  - `--snapshot` specify to create PDB with snapshot copy.
- `--refreshablePDB` specifies to create refreshable PDB
  - `--refreshMode` specifies refresh mode for refreshable PDB. Valid values: `AUTO|MANUAL`
    - `--refreshIntervalInMinutes` specifies refresh interval for `refreshablePDB` in minutes
  - `--dblinkUsername` specifies common user of a remote database used for database link to connect to the remote database
    - `--honorCaseSensitiveUserName` indicates specified username is case sensitive
- `--dbLinkName` specifies an existing database link name in the target database that points to the remote database

When promoted, you must supply the SYS user password for the source PDB. The newly cloned PDB inherits administration passwords from the source PDB. The cloned PDB is named using the following format: `dbname_sourcepdbname`. This command is supported only for databases that are not in a Data Guard configuration and use Oracle Database version 12.2.0.1, or later.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Frequently Asked Questions

Q: What is the dbaascli pdb remoteClone command used for?

A: The `dbaascli pdb remoteClone` command is used to create a new Pluggable Database (PDB) as a clone of an existing PDB in another container database (CDB).

Q: What user should execute the dbaascli pdb remoteClone command?

A: The command should be executed as either the `root` or `oracle` user.

Q: What is required when prompted during the dbaascli pdb remoteClone operation?

A: You must supply the SYS user password for the source PDB.

Q: What does the --pdbName parameter specify?

A: The `--pdbName` parameter specifies the name of the source PDB that you want to clone.

Q: What does the --dbName parameter represent?

A: The `--dbName` parameter represents the name (`DB_NAME`) of the CDB that will host the newly cloned PDB.

Q: How should the --sourceDBConnectionString be formatted?

A: The `--sourceDBConnectionString` should be formatted as <*scan_name*>:<*scan_port*>/<*database_service_name*>.

Q: What is the purpose of the --targetPDBName parameter?

A: The `--targetPDBName` parameter specifies the name for the newly cloned PDB.

Q: What does --powerLimit control?

A: The `--powerLimit` parameter controls the degree of parallelism used for the cloning operation. The valid value is between 1 and 128.

Q: What does the --maxCPU parameter define?

A: The `--maxCPU` parameter defines the maximum number of CPUs to be allocated for the PDB cloning process.

Q: What is the function of --maxSize?

A: The `--maxSize` parameter specifies the maximum storage size in GB for the new PDB.

Q: What does the --resume parameter do?

A: The `--resume` parameter resumes the previous cloning operation.

Q: What should you provide with the --resume option?

A: You can specify a `--sessionID` to resume a specific session if you are resuming a previous operation.

Q: What does --executePrereqs control?

A: The `--executePrereqs` parameter determines if only the prerequisites for the cloning operation should be run. Valid values are `yes` or `no`.

Q: How does --waitForCompletion affect the operation?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `--waitForCompletion` parameter specifies whether to wait for the operation to complete or run it in the background. Valid values are `true` or `false`.

Q: What is specified by the --sourcePDBExportedTDEKeyFile parameter?

A: The `--sourcePDBExportedTDEKeyFile` parameter specifies the exported key file from the source PDB. This parameter is applicable only for Oracle Database version 12.1.

Q: What does the --blobLocation parameter define?

A: The `--blobLocation` parameter specifies the custom path where the standby BLOB file will be generated in a Data Guard environment.

Q: When is --standbyBlobFromPrimary used?

A: The `--standbyBlobFromPrimary` parameter specifies the location of the standby BLOB file prepared from the primary database. This is required only for standby database PDB operations.

Q: Can --blobLocation and --standbyBlobFromPrimary be used together?

A: No, `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive and cannot be used together.

Q: What does the --excludeUserTablespaces option do?

A: The `--excludeUserTablespaces` option allows you to skip specific user tablespaces from being cloned. For example, t1,t2,t3.

Q: What is the effect of --excludePDBData?

A: The `--excludePDBData` option specifies whether to skip user data from the source PDB during cloning. Valid values are `true` or `yes`.

Q: What is specified by --pdbAdminUserName?

A: The `--pdbAdminUserName` parameter specifies the new admin user name for the cloned PDB.

Q: What does the --lockPDBAdminAccount option control?

A: The `--lockPDBAdminAccount` option specifies whether to lock the PDB admin user account. The default value is `true`.

Q: What does --sourcePDBServiceConvertList specify?

A: The `--sourcePDBServiceConvertList` parameter specifies a comma-delimited list of source to target service name conversions. For example, *source_srv1*:*new_srv1*,*source_srv2*:*new_srv2*.

Q: What is the purpose of --refreshablePDB?

A: The `--refreshablePDB` parameter specifies whether to create a refreshable PDB.

Q: What does --refreshMode control?

A: The `--refreshMode` parameter controls the refresh mode for a refreshable PDB. Valid values are `AUTO` or `MANUAL`.

Q: How does --refreshIntervalInMinutes work?

A: The `--refreshIntervalInMinutes` parameter specifies the interval in minutes for refreshing the refreshable PDB.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What is --dblinkUsername used for?

A: The `--dblinkUsername` parameter specifies a common user of a remote database used for the database link to connect to the remote database.

Q: What does the --honorCaseSensitiveUserName option indicate?

A: The `--honorCaseSensitiveUserName` option indicates that the specified username is case sensitive.

Q: What is the effect of --updateDBBlockCacheSize?

A: The `--updateDBBlockCacheSize` option enables the application to set the DB block cache size initialization parameters to support data copy with a different block size.

Q: What should I do if I encounter an error with the dbaascli pdb remoteClone command?

A: Review the error message for details, ensure all parameters are correctly specified, and verify that you have the necessary permissions and credentials. Additionally, check that the source and target databases meet all the requirements.

Q: What if I forget the SYS user password for the source PDB?

A: You will need to reset or recover the SYS user password for the source PDB. Without it, the cloning operation cannot be completed.

Example 7-40 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 pdb relocate

To relocate the specified PDB from the remote database into local database, use the `dbaascli pdb relocate` command.

Prerequisite

Run the command as the `oracle` user. When prompted, you must supply the SYS user password for the source database.

Syntax

```
dbaascli pdb relocate --pdbName <value> --dbName <value> --
sourceDBConnectionString <value> [--targetPDBName <value>] [--powerLimit
<value>] [--maxCpu <value>] [--maxSize <value>] [--resume [--sessionID
<value>]] [--executePrereqs] [--sourcePDBServices <value>] [--
sourcePDBReadOnlyServices <value>] [--relocateAvailabilityType <value>] [--
waitForCompletion <value>]
        {
            [--blobLocation <value>]
            | [--standbyBlobFromPrimary <value>]
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        }
        [--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`.
- `--updateDBBlockCachesize` specifies to enable application to set `db block cache size` initialization parameters to support data copy with different block size

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--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.

Frequently Asked Questions

Q: What is the dbaascli pdb relocate command used for?

A: The `dbaascli pdb relocate` command is used to relocate a Pluggable Database (PDB) from a remote database to a local database.

Q: What user should run the dbaascli pdb relocate command?

A: The command should be run as the `Oracle` user.

Q: What is required when prompted during the dbaascli pdb relocate operation?

A: You must supply the SYS user password for the source database.

Q: What does the --pdbName parameter specify?

A: The `--pdbName` parameter specifies the name of the source PDB to be relocated.

Q: What is the purpose of the --dbName parameter?

A: The `--dbName` parameter specifies the target database name where the PDB will be relocated.

Q: How should the --sourceDBConnectionString be formatted?

A: The `--sourceDBConnectionString` should be formatted as <*scan_name*>:<*scan_port*>/<*database_service_name*>.

Q: What does the --targetPDBName parameter do?

A: The `--targetPDBName` parameter specifies a new name for the relocated PDB.

Q: What is the use of --powerLimit?

A: The `--powerLimit` parameter specifies the degree of parallelism to be used during the relocate operation.

Q: How does --maxCpu affect the relocation process?

A: The `--maxCpu` parameter specifies the maximum number of CPUs to be allocated for the PDB relocation process.

Q: What does the --maxSize parameter define?

A: The `--maxSize` parameter defines the maximum storage size in GB for the new PDB.

Q: What is the function of --resume?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `--resume` parameter indicates that the relocation operation should resume from where it left off.

Q: What should I provide with the --resume option?

A: You can specify a `--sessionID` to resume a specific session if you are resuming a previous operation.

Q: What does the --executePrereqs parameter do?

A: The `--executePrereqs` parameter determines if only the prerequisites for the operation should be run. Valid values are yes or no.

Q: What is specified by the --sourcePDBServices parameter?

A: The `--sourcePDBServices` parameter specifies a list of comma-delimited source PDB services.

Q: What does the --sourcePDBReadOnlyServices parameter list?

A: The `--sourcePDBReadOnlyServices` parameter lists a comma-delimited list of source PDB read-only services.

Q: What is the effect of --waitForCompletion?

A: The `--waitForCompletion` parameter specifies whether to run the operation in the background. Valid values are `true` or `false`.

Q: What does the --blobLocation parameter specify?

A: The `--blobLocation` parameter specifies the location of a custom directory where the standby BLOB file will be generated in a Data Guard environment.

Q: When should I use --standbyBlobFromPrimary?

A: Use `--standbyBlobFromPrimary` to specify the location of the standby BLOB file, which is prepared from the primary database. This is required only for standby operations.

Q: Can I use --blobLocation and --standbyBlobFromPrimary together?

A: No, the `--blobLocation` and `--standbyBlobFromPrimary` parameters are mutually exclusive and cannot be used together.

Q: What does --upgradePDB do?

A: The `--upgradePDB` parameter specifies whether to upgrade the PDB as part of the relocation operation. Valid values are `true` or `false`.

Q: What is the purpose of --updateDBBlockCacheSize?

A: The `--updateDBBlockCacheSize` option allows the application to set the DB block cache size initialization parameters to support data copy with a different block size.

Q: What does the --skipOpenPDB option do?

A: The `--skipOpenPDB` option indicates that the PDB should not be opened at the end of the relocation operation.

Q: When should I use --completePDBRelocate?

A: Use `--completePDBRelocate` to complete the PDB relocation if it is done as a two-step operation.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What should I do if I encounter an error while using the dbaascli pdb relocate command?

A: Check the error message for details, ensure all parameters are correctly specified, and verify that you have the necessary permissions and credentials. You might also need to review the prerequisites and configurations.

Q: What if I forget the SYS user password for the source database?

A: You will need to reset or recover the SYS user password for the source database. Without it, you cannot complete the relocation operation.

Example 7-41 dbaascli pdb relocate

```
dbaascli pdb relocate --sourceDBConnectionString test-
scan.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521/
source_cdb_service_name --pdbName source_pdb --dbName target_db
```

System Management

This section focuses on overseeing and managing Oracle homes within the system. It includes commands like `dbaascli system getDBHomes` to view details about all Oracle Database homes and `dbaascli system getGridHomes` to list the details of all Grid Infrastructure homes. These commands provide essential information for maintaining and troubleshooting the overall system environment.

- dbaascli system getDBHomes
  To view information about all the Oracle homes, use the `dbaascli system getDBHomes` command.
- dbaascli system getGridHomes
  To list the details of all Grid homes, use the `dbaascli system getGridHomes` command.

dbaascli 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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: No, the `dbaascli system getDBHomes` command does not have any parameters.

Q: What kind of information does the dbaascli system getDBHomes command provide?

A: The command provides details about all Oracle homes on the system, including their paths and other relevant information.

Q: How can I interpret the output from the dbaascli system getDBHomes command?

A: The output will list all Oracle homes with information such as the location of each Oracle home. This information can help in managing and configuring Oracle environments.

Q: What should I do if the dbaascli system getDBHomes command does not return any output?

A: Ensure that you are running the command as the `root` or `oracle` user and verify that Oracle homes are properly installed on the system. You may also want to check system permissions and configurations.

Q: What if I receive an error message while executing the dbaascli system getDBHomes command?

A: Check the error message for specific details, verify that you have the appropriate permissions, and ensure that the `dbaascli` tool is correctly installed and configured.

Q: Can I run dbaascli system getDBHomes on a non-Oracle system?

A: No, the `dbaascli system getDBHomes` command is specific to Oracle systems and requires Oracle software to be installed.

Example 7-42 dbaascli system getDBHomes

```
dbaascli system getDBHomes
```

dbaascli system getGridHomes

To list the details of all Grid homes, use the `dbaascli system getGridHomes` command.

Prerequisite

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli system getGridHomes
```

Frequently Asked Questions

Q: What is the dbaascli system getGridHomes command used for?

A: The `dbaascli system getGridHomes` command is used to list the details of all Grid homes on a system.

Q: What user should execute the dbaascli system getGridHomes command?

A: The command should be executed as either the `root` or `oracle` user.

Q: Are there any parameters for the dbaascli system getGridHomes command?

A: No, the `dbaascli system getGridHomes` command does not have any parameters.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What kind of information does the dbaascli system getGridHomes command provide?

A: The command provides details about all Grid homes on the system, including their locations and other relevant information.

Q: How can I interpret the output from the dbaascli system getGridHomes command?

A: The output will list all Grid homes with information such as their paths and configurations. This helps in managing and configuring Oracle Grid infrastructure.

Q: What should I do if the dbaascli system getGridHomes command does not return any output?

A: Ensure you are running the command as the `root` or `oracle` user and verify that Oracle Grid homes are properly installed on the system. Check system permissions and configurations if needed.

Q: What if I receive an error message while executing the dbaascli system getGridHomes command?

A: Review the error message for specific details, ensure you have the appropriate permissions, and verify that the `dbaascli` tool is properly installed and configured.

Q: Can I run dbaascli system getGridHomes on a system without Oracle Grid infrastructure?

A: No, the `dbaascli system getGridHomes` command is specific to systems with Oracle Grid infrastructure installed. If no Grid homes are present, the command may not return any results.

Transparent Data Encryption (TDE) Management

This section encompasses the management of Transparent Data Encryption (TDE) for securing database data. It includes commands for handling encryption keys and keystores, such as `dbaascli tde addSecondaryHsmKey` to add secondary HSM keys, `dbaascli tde rotateMasterKey` to rotate the master key, and `dbaascli tde encryptTablespacesInPDB` to encrypt tablespaces within a pluggable database. You can also convert between FILE and HSM-based TDE (`dbaascli tde fileToHsm`, `dbaascli tde hsmToFile`), manage key versions, and retrieve key details with various commands. These tools ensure effective encryption management and data security.

- dbaascli tde addSecondaryHsmKey
  To add a secondary HSM (KMS) key to the existing HSM (KMS) configuration, use the `dbaascli tde addSecondaryHsmKey` 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 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- dbaascli tde getHsmKeys
  To get TDE active key details, use the `dbaascli tde getHsmKeys` command.
- dbaascli tde getMkidForKeyVersionOCID
  To get Master Key ID associated with the KMS key version OCID, use the `dbaascli tde getMkidForKeyVersionOCID` command.
- dbaascli tde getPrimaryHsmKey
  To get primary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde getPrimaryHsmKey` command.
- dbaascli tde hsmToFile
  To convert HSM (KMS/OKV) based TDE to FILE based TDE, use the `dbaascli tde hsmToFile` command.
- dbaascli tde listKeys
  To list TDE master keys, use the `dbaascli tde listKeys` command.
- dbaascli tde removeSecondaryHsmKey
  To remove secondary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde removeSecondaryHsmKey` command.
- dbaascli tde rotateMasterKey
  To rotate the master key for database encryption, use the `dbaascli tde rotateMasterKey` command.
- dbaascli tde setKeyVersion
  To set the version of the primary key to be used in DB/CDB or PDB, use the `dbaascli tde setKeyVersion` command.
- dbaascli tde setPrimaryHsmKey
  To change the primary HSM (KMS) key for the existing HSM (KMS) configuration, use the `dbaascli tde setPrimaryHsmKey` command.
- dbaascli tde status
  To display information about the keystore for the specified database, use the `dbaascli tde status` command.

dbaascli 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Frequently Asked Questions

Q: What does the dbaascli tde addSecondaryHsmKey command do?

A: The `dbaascli tde addSecondaryHsmKey` command adds a secondary HSM (KMS) key to the existing HSM (KMS) configuration for an Exadata Cloud@Customer database.

Q: Who should run the dbaascli tde addSecondaryHsmKey command?

A: The command must be run as the `root` user.

Q: On which machine should I run the dbaascli tde addSecondaryHsmKey command?

A: You need to connect to an Exadata Cloud@Customer virtual machine using SSH to run this command.

Q: Where can I find more details about connecting to a virtual machine to run this command?

A: You can refer to the guide "Connecting to a Virtual Machine with SSH" for instructions on how to connect.

Q: What does the --secondaryKmsKeyOCID option specify?

A: The `--secondaryKmsKeyOCID` option specifies the OCID (Oracle Cloud Identifier) of the secondary KMS key to be added to the existing HSM (KMS) configuration.

Q: What does the --dbname option do?

A: The `--dbname` option allows you to specify the name of the database for which the secondary KMS key should be added. It is optional.

Q: What does the --precheckOnly option do?

A: The `--precheckOnly` option, when set to `yes`, runs a precheck of the operation without making any actual changes. The valid values are `yes` or `no`.

Q: Can I run the precheck only without making changes?

A: Yes, you can use the `--precheckOnly yes` option to run just the precheck without making changes.

Q: Can you give an example of how to run this command to add a secondary HSM key?

A: Here's an example:

```
dbaascli tde addSecondaryHsmKey --secondaryKmsKeyOCID ocid1.kms.key.oc1..example
```

Q: How do I run the command for a specific database?

A: You can specify the database name like this:

```
dbaascli tde addSecondaryHsmKey --secondaryKmsKeyOCID ocid1.kms.key.oc1..example
--dbname mydatabase
```

Q: How do I run the command with a precheck only?

A: To run the precheck, use the following syntax:

```
dbaascli tde addSecondaryHsmKey --secondaryKmsKeyOCID ocid1.kms.key.oc1..example
--precheckOnly yes
```

Q: What should I do if the command fails?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Ensure that you are running the command as the `root` user and that you have connected to the correct Exadata Cloud@Customer virtual machine. Also, verify the OCID of the KMS key and check if the required permissions are granted.

Q: How can I check if I have the correct OCID for the secondary KMS key?

A: You can retrieve the OCID of the KMS key from the Oracle Cloud Infrastructure console, under the Key Management Service (KMS) section.

Q: What permissions are required to add a secondary KMS key?

A: You need appropriate permissions in Oracle Cloud Infrastructure for KMS operations, including the ability to manage KMS keys for the relevant compartment.

Q: Can I use the dbaascli tde addSecondaryHsmKey command without specifying the --dbname option?

A: Yes, the `--dbname` option is optional. If omitted, the command applies to all databases using the existing HSM (KMS) configuration.

Q: What happens if I add a secondary KMS key?

A: The secondary KMS key will be added to the existing configuration, providing an additional layer of encryption key management redundancy.

Q: Can I remove a secondary KMS key once it is added?

A: No, once a secondary KMS key is added, it cannot be removed. You can only rotate or update keys in the future.

Example 7-43 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 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>]]
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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?

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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.

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>
   ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli tde enableWalletRoot

To enable `wallet_root` spfile parameter for the existing database, use the `dbaascli tde enableWalletRoot` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde enableWalletRoot --dbname <value>
[--dbRestart <value>]
[--executePrereqs]
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the Oracle Database.
- `--dbrestart` specifies the database restart option. Valid values are: `rolling` or `full`. Default value: `rolling`
  If you do not pass the `dbrestart` argument, then the database restarts in a `rolling` manner.
- `--precheckOnly` runs only the precheck for this operation. Valid values are: `yes` or `no`
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.

Frequently Asked Questions

Q: What does the dbaascli tde enableWalletRoot command do?

A: The `dbaascli tde enableWalletRoot` command enables the `wallet_root` parameter in the `spfile` for an existing Oracle database on Exadata Cloud@Customer.

Q: Who should run the dbaascli tde enableWalletRoot command?

A: The command must be run as the `root` user.

Q: On which machine should I run the dbaascli tde enableWalletRoot command?

A: You must connect to an Exadata Cloud@Customer virtual machine using SSH to run this command.

Q: Where can I find instructions to connect to the virtual machine?

A: You can refer to the guide "Connecting to a Virtual Machine with SSH" for instructions on connecting.

Q: What does the --dbRestart option do?

A: The `--dbRestart` option specifies how the database should be restarted after enabling `wallet_root`. The valid values are:

- `rolling`: Restarts the database in a rolling manner (default behavior).
- `full`: Performs a full database restart.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What does the --dbname option do?

A: The `--dbname` option allows you to specify the name of the Oracle Database for which the `wallet_root` parameter should be enabled.

Q: What does the --precheckOnly option do?

A: The `--precheckOnly` option runs a precheck of the operation without making actual changes. The valid values are `yes` or `no`.

Q: What happens if I do not specify the --dbRestart option?

A: If you do not specify the `--dbRestart` option, the database will restart in a rolling manner by default.

Q: Can you give an example of how to enable wallet_root for a specific database?

A: Here's an example to enable `wallet_root` for a database named *`mydatabase`*:

```
dbaascli tde enableWalletRoot --dbname mydatabase
```

Q: How do I enable wallet_root and specify a full database restart?

A: You can enable `wallet_root` with a full database restart using the following command:

```
dbaascli tde enableWalletRoot --dbname mydatabase --dbRestart full
```

Q: How do I run the command with a precheck only?

A: To perform a precheck without making changes, use the following syntax:

```
dbaascli tde enableWalletRoot --dbname mydatabase --precheckOnly yes
```

Q: What are the prerequisites for running the dbaascli tde enableWalletRoot command?

A: You must run the command as the `root` user and be connected to the correct Exadata Cloud@Customer virtual machine.

Q: Do I need to restart the database to enable wallet_root?

A: Yes, the database will need to restart either in a rolling manner (default) or fully, depending on the option you choose.

Q: What should I do if the command fails?

A: Ensure that you are running the command as the root user, and verify that the database name (`--dbname`) is correct. Check for any precheck errors if you are running with `--precheckOnly`.

Q: What if the database fails to restart after running the command?

A: Verify that the correct restart option was used (`rolling` or `full`) and check the database logs for any errors. You may need to manually restart the database if the automatic restart fails.

Q: How can I check if wallet_root was enabled successfully?

A: You can verify the change by checking the database's `spfile` or using Oracle SQL queries to confirm that the `wallet_root` parameter is enabled.

Q: Can I enable wallet_root without restarting the database?

A: No, the database needs to restart for the change to take effect. You can choose between a rolling restart or a full restart.

Q: What is the difference between a rolling and full database restart?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-44 dbaascli tde enableWalletRoot

```
dbaascli tde enableWalletRoot --dbname db name --dbrestart rolling|full


dbaascli tde enableWalletRoot --dbname orcl


dbaascli tde enableWalletRoot --dbname orcl--dbrestart full
```

dbaascli tde encryptTablespacesInPDB

To encrypt all the tablespaces in the specified PDB, use the `dbaascli tde encryptTablespacesInPDB` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde encryptTablespacesInPDB --dbname <value> --pdbName <value>
[--executePrereqs]
```

Where:

- `--pdbName` specifies the name of the PDB to encrypt all the tablespaces.
- `--dbname` specifies the name of the Oracle Database.
- `--executePrereqs` execute the prerequisites checks and report the results.

Frequently Asked Questions

Q: What does the dbaascli tde encryptTablespacesInPDB command do?

A: The `dbaascli tde encryptTablespacesInPDB` command encrypts all the tablespaces in the specified pluggable database (PDB) for an Oracle Database on Exadata Cloud@Customer.

Q: Who should run the dbaascli tde encryptTablespacesInPDB command?

A: The command must be run as the `root` user.

Q: On which machine should I run the dbaascli tde encryptTablespacesInPDB command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: You need to connect to an Exadata Cloud@Customer virtual machine using SSH to run this command.

Q: Where can I find instructions for connecting to the virtual machine?

A: Refer to the guide "Connecting to a Virtual Machine with SSH" for connection instructions.

Q: What does the --pdbName option specify?

A: The `--pdbName` option specifies the name of the pluggable database (PDB) whose tablespaces need to be encrypted.

Q: What does the --dbname option do?

A: The `--dbname` option allows you to specify the name of the Oracle Database to which the PDB belongs.

Q: What does the --precheckOnly option do?

A: The `--precheckOnly` option runs a precheck of the encryption operation without making any actual changes. Valid values are `yes` or `no`.

Q: What does the --useSysdbaCredential option do?

A: The `--useSysdbaCredential` option specifies whether SYSDBA credentials should be used for the operation. Valid values are true or false.

Q: Can you give an example of how to encrypt tablespaces in a specific PDB?

A: Here's an example to encrypt all tablespaces in a PDB named *mypdb*:

```
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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: You may need SYSDBA credentials for this operation if specified using the `--useSysdbaCredential` option.

Q: What should I do if the command fails?

A: Ensure you are running the command as the root user, and verify that the PDB name (`--pdbName`) and database name (`--dbname`) are correct. You can also run the command with `--precheckOnly yes` to check for issues before running the full encryption.

Q: What should I do if encryption of the tablespaces fails?

A: Check the database logs and ensure that you have the necessary privileges and resources to perform the encryption. You may also need to verify that there is enough space to handle the encryption process.

Q: How can I check if the tablespaces in a PDB are encrypted?

A: You can query the database views related to encryption, such as `V$ENCRYPTED_TABLESPACES`, to verify if the tablespaces have been successfully encrypted.

Q: How do I verify if the precheck was successful?

A: If you ran the command with `--precheckOnly yes`, you can check the output for any warnings or errors indicating potential issues with the encryption process.

Q: Can I encrypt the tablespaces for multiple PDBs simultaneously?

A: No, you need to run the `dbaascli tde encryptTablespacesInPDB` command separately for each PDB.

Q: Can I partially encrypt some tablespaces in a PDB?

A: No, this command encrypts all tablespaces within the specified PDB. For partial encryption, you would need to use different database management commands.

Q: Does encrypting tablespaces impact database performance?

A: Encrypting tablespaces can have a temporary performance impact during the encryption process. However, the impact should be minimal once the encryption is complete.

Q: Can I undo the encryption of tablespaces?

A: No, once the tablespaces are encrypted, the encryption cannot be undone. You can only rotate or re-encrypt the keys as needed.

Q: What happens if the operation is interrupted during the encryption process?

A: If the operation is interrupted, you may need to rerun the command. The system will resume encryption from where it left off, and you can verify the status using database views.

Example 7-45 dbaascli tde encryptTablespacesInPDB

```
dbaascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb


dbaascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb --
executePrereqs
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli tde fileToHsm

To convert FILE based TDE to HSM (KMS/OKV) based TDE, use the `dbaascli tde fileToHsm` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde fileToHsm --kmsKeyOCID <value> --dbname <value>
[--skipPatchCheck <value>]
[--executePrereqs ]
[--primarySuc <value>]
{
    [--resume [--sessionID <value>]] | [--revert [--sessionID <value>]]
}
[--waitForCompletion <value>]
```

Where:

- `--kmsKeyOCID` specifies the KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE
- `--dbname` specifies the name of the database
- `--skipPatchCheck` skips validation check for required patches if the value passed for this argument is `true`. Valid values: `true` or `false`
- `--executePrereqs` 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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `--kmsKeyOCID` parameter specifies the KMS key OCID that will be used for TDE encryption when transitioning from file-based to HSM-based TDE.

Q: What does the --dbname parameter do?

A: The `--dbname` parameter specifies the name of the database for which you are converting the TDE from file-based to HSM-based.

Q: Can I skip the patch validation check while converting TDE?

A: Yes, by using the `--skipPatchCheck` parameter with the value `true`, you can skip the validation check for required patches.

Q: What is the --executePrereqs parameter used for?

A: The `--executePrereqs` parameter allows you to run only the prechecks for the TDE conversion process without performing the actual conversion. Valid values are `yes` or `no`.

Q: What does the --primarySuc parameter do in a Data Guard setup?

A: The `--primarySuc` parameter is used in a Data Guard environment to indicate that the command has been successfully run on the primary database. It should be specified on the standby database after the primary conversion is complete.

Q: How do I resume a previous TDE conversion?

A: You can resume a previously incomplete TDE conversion by using the `--resume` parameter. Optionally, you can specify a specific session ID with `--sessionID`.

Q: How do I revert a TDE conversion?

A: To revert a previous TDE conversion, use the `--revert` parameter. You can also provide the specific session ID you want to revert using `--sessionID`.

Q: How do I specify a session ID when resuming or reverting a TDE conversion?

A: You can use the `--sessionID` parameter to specify the ID of the session you want to resume or revert. Example: `--resume --sessionID <ID>` or `--revert --sessionID <ID>`.

Q: What happens if I set --waitForCompletion to false?

A: If you set `--waitForCompletion` to `false`, the TDE conversion process will run in the background, and the command prompt will return immediately. If set to `true`, the command will wait for the process to finish before returning control to the user.

Q: What are the valid values for the --waitForCompletion parameter?

A: Valid values are `true` or `false`. Setting it to true makes the command wait until the process is complete; setting it to `false` runs the process in the background.

Q: Can I run 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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-46 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.

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Refer to the guide "Connecting to a Virtual Machine with SSH" for instructions on connecting.

Q: What does the --dbname option do?

A: The `--dbname` option allows you to specify the name of the Oracle Database for which you want to retrieve TDE key details.

Q: What does the --infoFile option do?

A: The `--infoFile` option specifies the file path where the list of key OCIDs (Oracle Cloud Identifiers) will be saved. The output is in JSON format.

Q: Can you give an example of how to retrieve TDE key details for a specific database?

A: Here's an example to get the TDE key details for a database named *mydatabase*:

```
dbaascli tde getHsmKeys --dbname mydatabase
```

Q: How do I save the TDE key details to a file?

A: You can specify a file path using the `--infoFile` option to save the output in JSON format:

```
dbaascli tde getHsmKeys --dbname mydatabase --infoFile /path/to/output.json
```

Q: What are the prerequisites for running the dbaascli tde getHsmKeys command?

A: You must run the command as the `root` user and be connected to the Exadata Cloud@Customer virtual machine.

Q: Do I need SYSDBA credentials to retrieve TDE key details?

A: No, SYSDBA credentials are not required to run the `dbaascli tde getHsmKeys` command.

Q: In what format is the TDE key information saved when using the --infoFile option?

A: The output is saved in JSON format.

Q: What information is included in the TDE key details?

A: The details include key OCIDs and other metadata about the active encryption keys stored in the HSM for the specified database.

Q: What should I do if the command fails to retrieve the key details?

A: Ensure that you are running the command as the root user and that the database name (`--dbname`) is correct. Check your connection to the Exadata Cloud@Customer virtual machine.

Q: How do I check if the output file was created successfully?

A: You can check the specified file path for the output JSON file. If the file is missing, verify that the file path is correct and that you have write permissions to the directory.

Q: What should I do if the output file is empty?

A: Ensure that the database specified contains active TDE keys and that the `--dbname` parameter is correct. You may also need to check if there are any errors in the database logs.

Q: Can I retrieve TDE key details for multiple databases at once?

A: No, you must run the `dbaascli tde getHsmKeys` command separately for each database.

Q: How can I use the output file from the --infoFile option in other operations?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-47 dbaascli tde getHsmKeys

```
dbaascli tde getHsmkeys --dbname dbname


dbaascli tde getHsmkeys --dbname dbname --infoFile infoFilePath
```

dbaascli tde getMkidForKeyVersionOCID

To get Master Key ID associated with the KMS key version OCID, use the `dbaascli tde getMkidForKeyVersionOCID` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde getMkidForKeyVersionOCID --kmsKeyVersionOCID <value>
[--dbname <value>]
[--waitForCompletion <value>]
```

Where:

- `--kmsKeyVersionOCID` specifies the KMS key version OCID to set
- `--dbname` specifies the name of the database
- `--waitForCompletion` specify `false` to run the operation in background. Valid values : `true|false`.

Frequently Asked Questions

Q: What is the purpose of the dbaascli tde getMkidForKeyVersionOCID command?

A: The `dbaascli tde getMkidForKeyVersionOCID` command retrieves the Master Key ID (MKID) associated with a specific KMS key version OCID in Oracle Database Cloud Service environments.

Q: What are the prerequisites for running the dbaascli tde getMkidForKeyVersionOCID command?

A: You must:

- Run the command as the `root` user.
- Be connected to an Exadata Cloud@Customer virtual machine via SSH.

Q: Who can run the dbaascli tde getMkidForKeyVersionOCID command?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Only the `root` user can run this command.

Q: What does the --kmsKeyVersionOCID parameter specify?

A: The `--kmsKeyVersionOCID` parameter specifies the KMS key version OCID for which you want to retrieve the associated Master Key ID (MKID).

Q: What does the --dbname parameter specify?

A: The `--dbname` parameter specifies the name of the database for which the KMS key version OCID is being queried.

Q: Is the --dbname parameter mandatory?

A: No, the `--dbname` parameter is optional. If you don't specify a database name, the command will retrieve the MKID for the default database on the system.

Q: What should I do if I don't know the KMS key version OCID?

A: You must retrieve the KMS key version OCID from your KMS management console or service provider before using this command. Without it, the command cannot retrieve the Master Key ID (MKID).

Q: Can I run this command on a non-Exadata Cloud@Customer environment?

A: No, this command is specifically for use in an Exadata Cloud@Customer environment, and you need to connect to a virtual machine using SSH to execute it.

Q: What happens if I run the command without specifying a database name using --dbname?

A: If the `--dbname` parameter is not provided, the command will attempt to retrieve the MKID for the default database configured on the system.

Q: What should I do if I encounter an error while retrieving the MKID?

A: Ensure that:

- You are running the command as the `root` user.
- You are correctly connected to the Exadata Cloud@Customer virtual machine.
- The KMS key version OCID you provided is valid. If the error persists, check the system logs for more details.

Q: How do I connect to the Exadata Cloud@Customer virtual machine?

A: You can connect to the virtual machine via SSH. Refer to the Exadata Cloud@Customer documentation for steps on how to securely connect.

Example 7-48 dbaascli tde getMkidForKeyVersionOCID

```
dbaascli tde getMkidForKeyVersionOCID --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli tde getPrimaryHsmKey

To get primary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde getPrimaryHsmKey` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde getPrimaryHsmKey
[--dbname]
```

Where:

- `--dbname` specifies the name of the database

Frequently Asked Questions

Q: What is the purpose of the dbaascli tde getPrimaryHsmKey command?

A: The `dbaascli tde getPrimaryHsmKey` command retrieves the primary Hardware Security Module (HSM) key from the existing HSM (KMS) configuration in an Oracle Database environment.

Q: What are the prerequisites for running the dbaascli tde getPrimaryHsmKey command?

A: You must:

- Run the command as the `root` user.
- Be connected to an Exadata Cloud@Customer virtual machine via SSH.

Q: Who can execute the 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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-49 dbaascli tde getPrimaryHsmKey

```
dbaascli tde getPrimaryHsmKey --dbname dbname
```

dbaascli tde hsmToFile

To convert HSM (KMS/OKV) based TDE to FILE based TDE, use the `dbaascli tde hsmToFile` command.

Run the command as the `root` user.

Syntax

```
dbaascli tde hsmToFile
[--dbname <value>]
{
    [--prepareStandbyBlob <value> [--blobLocation <value>]
   | [--standbyBlobFromPrimary <value>]
}
]
[--skipPatchCheck <value>]
[--executePrereqs ]
[--primarySuc <value>]
{
     [--resume [--sessionID <value>]] |
     [--revert [--sessionID <value>]]
}
[--waitForCompletion <value>]
```

Where:

- `--dbname` specifies the name of the database
- `--prepareStandbyBlob` specify `true` to generate a blob file containing the artifacts needed to perform the operation in a DG environment.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specify the location of the standby blob file which is prepared from the primary database. This is required only for standby operations.]
- `--skipPatchCheck` skips validation check for required patches if the value passed for this argument is `true`. Valid values: `true` or `false`
- `--executePrereqs` execute the prerequisites checks and report the results.
- `--primarySuc` specify this property in the standby database of the Data Guard environment once the command is successfully run on the primary database
- `--resume` resumes the previous run
  - `--sessionID` specifies to resume a specific session ID
- `--revert` specifies to roll back the previous run
  - `--sessionID` specifies to rollback a specific session ID
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true|false`

Frequently Asked Questions

Q: What is the purpose of the dbaascli tde hsmToFile command?

A: The `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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: The `--primarySuc` parameter is used in a Data Guard setup to indicate that the conversion has successfully run on the primary database. It should be used when running the conversion on the standby database.

Q: How can I resume a previous hsmToFile conversion?

A: You can resume a previous conversion by using the `--resume` parameter. Optionally, you can specify the session ID of the previous run with `--sessionID`.

Q: What is the purpose of the --revert parameter?

A: The `--revert` parameter is used to roll back a previously initiated conversion process in case of failure or if you need to undo the operation.

Q: What happens if I set --waitForCompletion to false?

A: If you set `--waitForCompletion` to `false`, the operation will run in the background, allowing you to continue other tasks. If set to `true`, the command will wait for the process to complete before returning control to the user.

Q: What should I do if I need to convert the TDE in a standby database in a Data Guard setup?

A: In a Data Guard setup, after converting TDE on the primary database, you must run the command on the standby database using the `--primaryDBWalletTar` parameter, specifying the wallet tar file from the primary database, and include `--primarySuc`.

Q: What should I do if I want to skip checking for required patches during the conversion?

A: You can skip the patch check by using the `--skipPatchCheck` parameter and setting it to `true`.

Q: How do I check if the system is ready for the hsmToFile conversion without making changes?

A: You can perform only the prechecks by using the `--executePrereqs` parameter and setting it to `yes`.

Q: What should I do if the conversion process is interrupted?

A: You can use the `--resume` parameter to restart the process from where it left off. Optionally, you can specify a particular session ID with `--sessionID`.

Q: What should I do if the conversion process fails?

A: If the conversion fails, you can roll back the process using the `--revert` parameter. Additionally, review any error messages and check system logs for more details.

Q: Can I run the 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-50 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.

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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: Is the --dbname parameter mandatory?

A: No, the `--dbname` parameter is optional. If not specified, the command will list the TDE master keys for the default database configured on the system.

Q: What should I do if I want to save the list of keys to a file?

A: You should provide the `--file` parameter along with the desired file path. For example:

```
dbaascli tde listKeys --file /path/to/output.txt
```

Q: What happens if I don't provide a database name with --dbname?

A: If the `--dbname` parameter is not provided, the command will list the TDE master keys for the default database on the system.

Q: Can I use this command in environments other than Exadata Cloud@Customer?

A: This command is designed specifically for Exadata Cloud@Customer environments. Ensure you are connected to the appropriate virtual machine to run it.

Q: What should I do if the command fails to list keys?

A: Ensure that:

- You are running the command as the `root` user.
- You are connected to the Exadata Cloud@Customer virtual machine.
- The database name (if specified) is correct. Check the error messages and logs for more details on the failure.

Q: Can I run the dbaascli tde listKeys command while the database is running?

A: Yes, the command can be executed while the database is running. It simply lists the TDE master keys and does not alter the state of the database.

Q: Do I need special permissions to run this command?

A: You must run this command as the `root` user. Without root permissions, you will not be able to execute the command.

Q: What is the purpose of listing TDE master keys?

A: Listing TDE master keys allows you to review the encryption keys being used for protecting your database's data. It is essential for monitoring and managing encryption settings.

Q: How do I connect to the Exadata Cloud@Customer virtual machine to run the command?

A: You can connect to the virtual machine using SSH. Refer to the Exadata Cloud@Customer documentation for instructions on how to establish a secure connection.

Example 7-51 dbaascli tde listKeys

```
dbaascli tde listKeys --dbname dbname


dbaascli tde listKeys --dbname dbname --infoFilePath infoFilePath
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli tde removeSecondaryHsmKey

To remove secondary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde removeSecondaryHsmKey` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde removeSecondaryHsmKey --dbname <value>
[--confirmDeletion]
[--secondaryKmsKeyOCID]
[--executePrereqs]
```

Where:

- `--dbname` specifies the name of the database
- `--confirmDeletion` if not specified the user will be prompted while deleting all existing HSM(KMS) keys.
- `--secondaryKmsKeyOCID` secondary KMS key to be removed from existing HSM(KMS) configuration. If not specified all secondary KMS keys will be removed.
- `--executePrereqs` execute the prerequisites checks and report the results.

Frequently Asked Questions

Q: What is the purpose of the dbaascli tde removeSecondaryHsmKey command?

A: The `dbaascli tde removeSecondaryHsmKey` command is used to remove a secondary Hardware Security Module (HSM) key from the existing HSM (KMS) configuration in an Oracle Database environment.

Q: What are the prerequisites for running the dbaascli tde removeSecondaryHsmKey command?

A: You must:

- Run the command as the `root` user.
- Be connected to an Exadata Cloud@Customer virtual machine using SSH.

Q: What does the --force parameter do in the dbaascli tde removeSecondaryHsmKey command?

A: The `--force` parameter allows the removal of the secondary HSM key without prompting the user for confirmation. If not specified, the command will prompt the user before deleting any keys.

Q: What does the --secondaryKmsKeyOCID parameter specify?

A: The `--secondaryKmsKeyOCID` parameter specifies the OCID (Oracle Cloud Identifier) of the secondary KMS key you want to remove from the existing HSM configuration.

Q: What does the --dbname parameter do?

A: The `--dbname` parameter specifies the name of the database for which the secondary HSM key is being removed.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What is the purpose of the --precheckOnly parameter?

A: The `--precheckOnly` parameter, if set to `yes`, will only run the prechecks to validate the readiness for the removal operation without actually removing the secondary HSM key. If set to `no`, the full removal operation is performed.

Q: Is the --force parameter mandatory?

A: No, the `--force` parameter is optional. If it's not specified, the system will prompt the user for confirmation before proceeding with the key removal.

Q: Is the --secondaryKmsKeyOCID parameter mandatory?

A: Yes, you must provide the `--secondaryKmsKeyOCID` to identify the specific secondary HSM key that you want to remove from the configuration.

Q: Is the --dbname parameter mandatory?

A: No, the `--dbname` parameter is optional. If not specified, the command will attempt to remove the secondary HSM key from the default database on the system.

Q: What should I do if I want to remove the secondary HSM key without any user prompts?

A: You should use the `--force` parameter to bypass the confirmation prompt and remove the secondary HSM key directly:

```
dbaascli tde removeSecondaryHsmKey --force --secondaryKmsKeyOCID <value>
```

Q: How can I test whether the system is ready to remove the secondary HSM key without actually removing it?

A: You can use the `--precheckOnly` parameter set to yes to perform a precheck:

```
dbaascli tde removeSecondaryHsmKey --precheckOnly yes --secondaryKmsKeyOCID
<value>
```

Q: What happens if I don't provide a database name with --dbname?

A: If the `--dbname` parameter is not specified, the command will attempt to remove the secondary HSM key from the default database configured on the system.

Q: What should I check if the command fails to remove the secondary HSM key?

A: Ensure that:

- You are running the command as the `root` user.
- You are connected to the Exadata Cloud@Customer virtual machine.
- The correct `--secondaryKmsKeyOCID` and `--dbname` values are provided. Check the error messages and logs for more details on the failure.

Q: What should I do if the removal operation fails partway through?

A: If the operation fails, review the error logs and try running the command with `--precheckOnly` to ensure the system is ready for the operation. If necessary, correct any issues before retrying.

Q: Can I run the dbaascli tde removeSecondaryHsmKey command while the database is running?

A: Yes, the command can be executed while the database is running, as it does not require the database to be stopped.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What is the purpose of removing a secondary HSM key?

A: Removing a secondary HSM key is typically done when the key is no longer needed or when you want to manage the encryption keys used in your TDE (Transparent Data Encryption) configuration.

Q: How do I connect to the Exadata Cloud@Customer virtual machine to run the command?

A: You can connect to the virtual machine using SSH. Refer to the Exadata Cloud@Customer documentation for instructions on establishing a secure connection.

Example 7-52 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
```

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

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--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.

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`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Q: What does the --standbyBlobFromPrimary parameter specify?

A: The `--standbyBlobFromPrimary` parameter specifies the location of the standby BLOB file that was generated from the primary database. This parameter is used when performing the master key rotation on a standby database in a Data Guard environment.

Q: Is the --rotateMasterKeyOnAllPDBs parameter mandatory?

A: No, the `--rotateMasterKeyOnAllPDBs` parameter is optional. If it is not specified, the master key will only be rotated for the database (or specific PDB) provided in the `--dbname` or `--pdbName` parameters.

Q: Is the --pdbName parameter required if I'm rotating keys for a CDB?

A: No, the `--pdbName` parameter is only required if you want to rotate the master key for a specific Pluggable Database (PDB). It is optional when rotating the key for the entire CDB.

Q: Do I need to use the --prepareStandbyBlob and --standbyBlobFromPrimary parameters for standalone databases?

A: No, these parameters are only relevant in a Data Guard environment where a standby database is involved.

Q: How can I rotate the master key for all PDBs in a CDB?

A: You should use the `--rotateMasterKeyOnAllPDBs` parameter set to true to rotate the master key for all PDBs in the CDB. For example:

```
dbaascli tde rotateMasterKey --dbname CDB_NAME --rotateMasterKeyOnAllPDBs true
```

Q: How do I run a check to validate that the system is ready for master key rotation without performing the actual operation?

A: You can use the `--executePrereqs` parameter to run the prerequisite checks. This will report any issues that might prevent the master key rotation:

```
dbaascli tde rotateMasterKey --dbname DB_NAME --executePrereqs
```

Q: What should I do if the operation was interrupted, and I want to resume it?

A: You can use the `--resume` parameter to resume the previously interrupted operation. If you have a session ID, provide it with the `--sessionID` parameter:

```
dbaascli tde rotateMasterKey --dbname DB_NAME --resume --sessionID <value>
```

Q: How can I prepare for key rotation in a Data Guard environment?

A: You should use the `--prepareStandbyBlob` parameter to generate a BLOB file that contains the required artifacts for rotating the master key in a standby environment:

```
dbaascli tde rotateMasterKey --dbname DB_NAME --prepareStandbyBlob true --
blobLocation /path/to/blob
```

Q: How do I apply the standby BLOB file from the primary database when rotating keys on a standby database?

A: Use the `--standbyBlobFromPrimary` parameter to specify the location of the BLOB file that was prepared on the primary database:

```
dbaascli tde rotateMasterKey --dbname DB_NAME --standbyBlobFromPrimary /path/to/
blob
```

Q: What should I check if the master key rotation fails?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Ensure that:

- You are running the command as the `root` user.
- The database name (`--dbname`) is correct.
- Any prerequisite checks were run using `--executePrereqs` to ensure readiness. Review the error logs for more detailed information on the failure.

Q: What should I do if the operation does not complete successfully in a Data Guard environment?

A: Ensure that the BLOB file from the primary database was prepared correctly using `--prepareStandbyBlob`, and then use `--standbyBlobFromPrimary` to apply it on the standby database.

Q: Can I run the dbaascli tde rotateMasterKey command while the database is running?

A: Yes, the command can be executed while the database is running. However, it is recommended to run prerequisite checks beforehand using the `--executePrereqs` option.

Q: Why is rotating the master key important?

A: Rotating the master key improves database security by ensuring that the encryption keys used for data protection are periodically updated, reducing the risk of key compromise.

Q: Do I need to restart the database after rotating the master key?

A: No, restarting the database is not required after rotating the master key. The key rotation will take effect immediately without any service disruption.

dbaascli tde setKeyVersion

To set the version of the primary key to be used in DB/CDB or PDB, use the `dbaascli tde setKeyVersion` command.

Run the command as the `root` user.

Syntax

```
dbaascli tde setKeyVersion --kmsKeyVersionOCID <value> --dbname <value>
[--pdbName <value>]
[--masterKeyID <value>]
[--standbySuc]
[--executePrereqs]
[--waitForCompletion <value>]
```

Where:

- `--kmsKeyVersionOCID` specifies the KMS key version OCID to set.
- `--dbname` specifies the name of the database.
- `--pdbName` name of the PDB to use the key version OCID.
- `--masterKeyID` specifies the master key ID of the given key version OCID. This is applicable to the Data Guard environment.
- `--standbySuc` specify this property in the primary database of the Data Guard environment once the command is successfully run on the standby database
- `--executePrereqs` execute the prerequisites checks and report the results.

Copyright © 2008, 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 tde setKeyVersion command?

A: The `dbaascli tde setKeyVersion` command is used to set the version of the primary encryption key that should be used for Transparent Data Encryption (TDE) in a database or Pluggable Database (PDB). This allows for the specific version of a KMS key to be assigned to the database.

Q: What are the prerequisites for using the dbaascli tde setKeyVersion command?

A: You must run the command as the `root` user and ensure that you are connected to an Exadata Cloud@Customer virtual machine.

Q: What does the --kmsKeyVersionOCID parameter specify?

A: The `--kmsKeyVersionOCID` parameter specifies the KMS key version OCID (Oracle Cloud Identifier) that you want to set for the database or PDB.

Q: What does the --dbname parameter specify?

A: The `--dbname` parameter specifies the name of the Oracle Database for which the key version will be set.

Q: What is the purpose of the --pdbName parameter?

A: The `--pdbName` parameter specifies the name of the Pluggable Database (PDB) within a Container Database (CDB) where you want to set the specific KMS key version.

Q: What is the --masterKeyID parameter used for?

A: The `--masterKeyID` parameter specifies the master key ID that is associated with the given KMS key version OCID. This is particularly important in a Data Guard environment.

Q: What is the role of the --standbySuc parameter?

A: The `--standbySuc` parameter is used in a Data Guard environment. It specifies that this property should be set on the primary database after successfully running the command on the standby database.

Q: What does the --executePrereqs parameter do?

A: The `--executePrereqs` parameter specifies whether to run prerequisite checks before executing the operation. Valid values are `yes` or `no`.

Q: What does the --waitForCompletion parameter control?

A: The `--waitForCompletion` parameter determines whether the operation will run synchronously (waiting for completion) or asynchronously (in the background). Valid values are `true` or `false`.

Q: Is the --pdbName parameter required if setting the key version for a CDB?

A: No, the `--pdbName` parameter is only required if you are setting the key version for a specific Pluggable Database (PDB). It is optional if you are setting the key version for the entire Container Database (CDB).

Q: Is the --masterKeyID parameter necessary for non-Data Guard environments?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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>
  ```
- 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-53 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-
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

Copyright © 2008, 2026, Oracle and/or its affiliates.

- `--allStandbyPrepared` specify to confirm that the operation has been successfully run on all the standby databases.
- `--bounceDatabase` specify this flag to do rolling database bounce for this operation
- `--executePrereqs` execute the prerequisites checks and report the results.
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.

Frequently Asked Questions

Q: What is the purpose of the dbaascli tde setPrimaryHsmKey command?

A: The `dbaascli tde setPrimaryHsmKey` command is used to change the primary HSM (Hardware Security Module) or KMS (Key Management Service) key in an existing HSM/KMS configuration for Transparent Data Encryption (TDE).

Q: What are the prerequisites for running the dbaascli tde setPrimaryHsmKey command?

A: The command must be executed as the `root` user, and the environment should be an Exadata Cloud@Customer virtual machine.

Q: What does the --primaryKmsKeyOCID parameter specify?

A: The `--primaryKmsKeyOCID` parameter specifies the OCID (Oracle Cloud Identifier) of the primary KMS key to be set for the TDE environment.

Q: What is the function of the --dbname parameter?

A: The `--dbname` parameter specifies the name of the Oracle Database for which the primary HSM/KMS key will be set.

Q: What does the --standbySuc parameter do?

A: The `--standbySuc` parameter is used in a Data Guard environment. It specifies that the command should be run on the primary database after successfully executing it on the standby database.

Q: What is the purpose of the --precheckOnly parameter?

A: The `--precheckOnly` parameter allows you to run only the prechecks for this operation. It validates the environment without making any actual changes. Valid values are `yes` or `no`.

Q: What does the --bounceDatabase parameter control?

A: The `--bounceDatabase` parameter specifies whether the database should be bounced (restarted) in a rolling manner as part of the operation. This ensures minimal downtime by restarting parts of the database one by one.

Q: How do I set the primary KMS key for my database?

A: To set the primary KMS key, run the following command:

```
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
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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?

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 7-54 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-
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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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>
```

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)?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: Yes, by specifying the appropriate database name using the `--dbname` parameter, you can view the keystore status for both CDBs and PDBs.

Q: What does it mean if the command returns an error about database connectivity?

A: This could indicate an issue with the connection to the database or a problem with your environment. Ensure that the database is running and accessible, and verify your SSH connection to the Exadata Cloud@Customer virtual machine.

Q: What happens if the database name is incorrect?

A: If the `--dbname` parameter specifies an incorrect or non-existent database, the command will fail, and you'll receive an error message indicating the problem.

Q: How do I troubleshoot if the keystore status indicates an unexpected state?

A: If the keystore status indicates an unexpected state, review the database logs for more details and check the configuration of the keystore to ensure it is properly set up.

Q: Can I automate the check of keystore status for monitoring purposes?

A: Yes, you can script the `dbaascli tde status` command to check keystore status periodically or integrate it into your database monitoring tools.

Q: How do I verify that Transparent Data Encryption (TDE) is correctly enabled?

A: You can verify TDE is correctly enabled by checking the status of the keystore using the `dbaascli tde status` command. A valid and open keystore indicates that TDE is properly configured.

Example 7-55 dbaascli tde status

```
dbaascli tde status --dbname dbname
```

Deprecated dbaascli commands

The `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.

- dbaascli patch db apply
- dbaascli patch db prereq
- dbaascli tde status
  To display information about the keystore for the specified database, use the `dbaascli tde status` command.

Copyright © 2008, 2026, Oracle and/or its affiliates.

dbaascli patch db apply

> **Note**
>
> `dbaascli patch db prereq` and `dbaascli patch db apply` commands have been deprecated in `dbaascli` release 21.2.1.2.0, and replaced with `dbaascli grid patch`, `dbaascli dbhome patch`, and `dbaascli database move` commands.
> For more information, see:
>
> - `dbaascli grid patch`
> - `dbaascli dbhome patch`
> - `dbaascli database move`
> - *Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli*

Related Topics

- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.
- Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli
  Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud@Customer system.

dbaascli patch db prereq

> **Note**
>
> `dbaascli patch db prereq` and `dbaascli patch db apply` commands have been deprecated in `dbaascli` release 21.2.1.2.0, and replaced with `dbaascli grid patch`, `dbaascli dbhome patch`, and `dbaascli database move` commands.
> For more information, see:
>
> - `dbaascli grid patch`
> - `dbaascli dbhome patch`
> - `dbaascli database move`
> - *Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- dbaascli grid patch
  To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.
- dbaascli dbHome patch
  To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- dbaascli database move
  To move the database from one home to another, use the `dbaascli database move` command.
- Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli
  Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud@Customer system.

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

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>
```

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?

Copyright © 2008, 2026, Oracle and/or its affiliates.

A: If the keystore status indicates an unexpected state, review the database logs for more details and check the configuration of the keystore to ensure it is properly set up.

Q: Can I automate the check of keystore status for monitoring purposes?

A: Yes, you can script the `dbaascli tde status` command to check keystore status periodically or integrate it into your database monitoring tools.

Q: How do I verify that Transparent Data Encryption (TDE) is correctly enabled?

A: You can verify TDE is correctly enabled by checking the status of the keystore using the `dbaascli tde status` command. A valid and open keystore indicates that TDE is properly configured.

Example 7-56 dbaascli tde status

```
dbaascli tde status --dbname dbname
```

Monitoring and Managing Exadata Storage Servers with ExaCLI

Learn to use the ExaCLI command-line utility to perform monitoring and management functions on Exadata storage servers in the Exadata Cloud Service.

- About the ExaCLI Command
  The ExaCLI command provides a subset of the commands found in the on-premises Exadata command line utility.
- Exadata Storage Server Username and Password
  You need a username and password to connect to the Exadata Storage Server.
- ExaCLI Command
  Use ExaCLI (`exacli`) to configure cell, database node configuration, and objects in the remote node environment, and to monitor your Oracle Exadata Database Service on Cloud@Customer service and objects.
- Connecting to a Storage Server with ExaCLI
  To use ExaCLI on storage servers, you will need to know your target storage server's IP address.

About the ExaCLI Command

The ExaCLI command provides a subset of the commands found in the on-premises Exadata command line utility.

ExaCLI offers a subset of the commands found in the on-premises Exadata command line utility. The utility runs on the database compute nodes in the Exadata Cloud Service.

Related Topics

- Using the CellCLI Utility

Exadata Storage Server Username and Password

You need a username and password to connect to the Exadata Storage Server.

On Exadata Database Service on Cloud@Customer, the preconfigured user for Exadata Storage Server is `cloud_user_`*`clustername`*, where *`clustername`* is the name of the virtual machine (VM) cluster that is being used.

Copyright © 2008, 2026, Oracle and/or its affiliates.

You can determine the name of the VM cluster by running the following `crsctl` command as the `grid` user on any cluster node:

```
crsctl get cluster name
```

The password for `cloud_user_clustername` is initially set to a random value, which you can view by running the following command as the `opc` user on any cluster node:

```
/opt/exacloud/get_cs_data.py
```

ExaCLI Command

Use ExaCLI (`exacli`) to configure cell, database node configuration, and objects in the remote node environment, and to monitor your Oracle Exadata Database Service on Cloud@Customer service and objects.

Purpose

ExaCLI (`exacli`) enables you configure your Oracle Exadata Database Service on Cloud@Customer system, and to obtain real-time information about your Exadata Cloud Service. To obtain information about the services and options on your system, run ExaCLI using the monitoring command parameter that you require.

To obtain a list of the system monitoring parameters you can use with ExaCLI, run the `LIST` parameter.

Syntax

```
exacli -c [username@]remotehost[:port]
[-l username]
[--xml]
[--cookie-jar filename]
[-e {command | 'command; command' | @batchfile}]
```

Options

Option	Description
`-c [username@]remotehost` or `--connect [username@]remotehost[:port]`	Specifies the remote node to which you want to connect. ExaCLI prompts for the user name if not specified.
`-l username` or `--login-name username`	Specifies the user name to log into the remote node. The preconfigured user is `cloud_user_clustername`.
`--xml`	Displays the output in XML format.
`--cookie-jar [filename]`	Specifies the file name of the cookie jar to use. If you do not specify a file name, then the cookie is stored in a default cookie jar located at *HOME*`/.exacli/cookiejar`, where *HOME* is the home directory of the operating system user running the `exacli` command. The presence of a valid cookie allows the ExaCLI user to run commands without requiring the user to log in during subsequent ExaCLI sessions.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Option	Description
`-e command` or `-e 'command[; command]'` or `-e @batchFile`	Specifies either the ExaCLI commands to run, or a batch file. After running the commands, ExaCLI quits. If you are specifying multiple commands to run, then enclose the commands in single quotes to prevent the shell from interpreting the semicolon. To start an interactive ExaCLI session, omit this command.
`--cert-proxy proxy[:port]`	Specifies the proxy server that you want to use when downloading certificates. If `port` is omitted, then port 80 is used by default.
`-n` or `--no-prompt`	Suppresses prompting for user input.

Command Parameters

To obtain information about objects and services on your system, use these ExaCLI command parameters.

Table 7-1 Command

Command Parameter	Description
`ACTIVEREQUEST`	Lists all active requests that are currently being served by the storage servers.
`ALERTDEFINITION`	Lists all possible alerts and their sources for storage servers.
`ALERTHISTORY`	Lists all alerts that have been issues for the storage servers.
`CELL`	Used to list the details of a specific attribute of the storage servers or storage cells. The syntax is as follows: `LIST CELL ATTRIBUTES A,B,C`, with A, B, and C being attributes. To see all cell attributes, use the `LIST CELL ATTRIBUTES ALL` command.
`CELLDISK`	Lists the attributes of the cell disks in the storage servers. Use the following syntax to list the cell disk details: `LIST CELLDISK cell_disk_name DETAIL`.
`DATABASE`	Lists details of the databases. Uses the regular `LIST` command syntax: `LIST DATABASE` and `LIST DATABASE DETAIL`. You can also use this command to show an individual attribute with the following syntax: `LIST DATABASE ATTRIBUTES NAME`.
`FLASHCACHE`	Lists the details of the Exadata system's flash cache. For this object, you can use the following syntax patterns: `LIST FLASHCACHE DETAIL` or `LIST FLASHCACHE ATTRIBUTES attribute_name`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-1 (Cont.) Command

Command Parameter	Description
FLASHCACHECONTENT	Lists the details of all objects in the flash cache, or the details of a specified object ID. To list all the details of all objects, use LIST FLASHCACHECONTENT DETAIL. To list details for a specific object, use a where clause as follows: LIST FLASHCACHECONTENT WHERE objectNumber=12345 DETAIL. Example query: finding the object_id value of an object select object_name, data_object_id from user_objects where object_name = 'BIG_CENSUS'; OBJECT_NAME DATA_OBJECT_ID ---------- --- BIG_CENSUS 29152
FLASHLOG	Lists the attributes for the Oracle Exadata Smart Flash Log.
GRIDDISK	Lists the details of a particular grid disk. The syntax is similar to the CELLDISK command syntax. To view all attributes: LIST GRIDDISK grid_disk_name DETAIL. To view specified attributes of the grid disk: LIST GRIDDISK grid_disk_name ATTRIBUTES size, name.
IBPORT	Lists details of the InfiniBand ports. Syntax is LIST IBPORT DETAIL.
IORMPLAN	Use the ExaCLI CREATE, ALTER, DROP, and LIST commands with IORMPLAN. To see the details of all IORM plans, use LIST IORMPLAN DETAIL. You can also use the command to create and alter IORM plans, and to apply plans to storage servers.
IORMPROFILE	Lists any IORM profiles that have been set on the storage servers. You can also refer back to the profile attribute on the DATABASE object if a database has an IORM profile on it. Syntax is LIST IORMPROFILE.
LIST	Lists the command parameter options available with ExaCLI for the Exadata Database Service on Cloud@Customer services and objects.
LUN	The LUN (logical unit number) object returns the number and the detail of the physical disks in the storage servers. List the LUNs of the disks with LIST LUN. List the details of each LUN with LIST LUN lun_number DETAIL.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-1 (Cont.) Command

Command Parameter	Description
METRICCURRRENT	Lists the current metrics for a particular object type. Syntax is `LIST METRICCURRENT WHERE objectType = 'CELLDISK'`. This command also allows for sorting and results limits as seen in the following example: `LIST METRICCURRENT attributes name, metricObjectName ORDER BY metricObjectName asc, name desc LIMIT 5`
METRICDEFINITION	Lists metric definitions for the object that you can then get details for. With the command `LIST metricDefinition WHERE objectType=cell`, you can get all the metrics for that object type. You can then use the metric definition object again to get details for one of those specific metrics just listed: `LIST metricDefinition WHERE name= IORM_MODE DETAIL`
METRICHISTORY	List metrics over a specified period of time. For example, with the command `LIST METRICHISTORY WHERE ageInMinutes < 30`, you can list all the metrics collected over the past 30 minutes. You can also use the predicate `collectionTime` to set a range from a specific time. Use `collectionTime` as shown in the follow example: `LIST METRICHISTORY WHERE collectionTime > '2018-04-01T21:12:00-10:00'`. The metric history object can also be used to see a specific metric using the object's name (for example, `LIST METRICHISTORY CT_FD_IO_RQ_SM`) or with a "where" clause to get objects with similar attributes like name (for example, `LIST METRICHISTORY WHERE name like 'CT_.*'`).
OFFLOADGROUP	Lists the attributes for the offload group that are running on your storage servers. You can list all details for all groups with `LIST OFFLOADGROUP DETAIL`, or list the attributes for a specific group, as shown in the following example: `LIST OFFLOADGROUP offloadgroup4`. List specific attributes with `LIST OFFLOADGROUP ATTRIBUTES name`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-1 (Cont.) Command

Command Parameter	Description
PHYSICALDISK	Lists all physical disks. Use the results of LIST PHYSICALDISK to identify a specific disk for further investigation, then list the details of that disk using the command as follows: LIST PHYSICALDISK 20:10 DETAIL. To list the details of flash disks, use the command as follows: LIST PHYSICALDISK FLASH_1_0 DETAIL).
PLUGGABLEDATABASE	Lists all PDBs. View the details of a specific PDB with LIST PLUGGABLEDATABASE pdb_name.
QUARANTINE	Lists all SQL statements that you prevented from using Smart Scans. The syntax is LIST QUARANTINE DETAIL. You can also use a "where" clause on any of the available attributes.
DIAGPACK	Use the ExaCLI CREATE, ALTER, DROP, and LIST commands with DIAGPACK to list the diagnostic packages and their status in your Exadata system. The syntax is LIST DIAGPACK [DETAIL], with DETAIL being an optional attribute. Use CREATE DIAGPACK with the packStartTime attribute to gather logs and trace files into a single compressed file for downloading, as in the following example: CREATE DIAGPACK packStartTime=2019_12_15T00_00_00. You can also use the value now with packStartTime: CREATE DIAGPACK packStartTime=now. To download a diagnostic package, use DOWNLOAD DIAGPACK *package_name local_directory*. For example, the following command downloads a diagnostic package to the /tmp directory: DOWNLOAD DIAGPACK cfclcx2647_diag_2018_06_03T00_44_24_1 /tmp.

Usage Notes

- Notes for the --cookie-jar option:
  - The user name and password are sent to the remote node for authentication. On successful authentication, the remote node issues a cookie (the login credentials) that is stored in the specified filename on the database node. If filename is not specified, the cookie is stored in a default cookie jar located at HOME/.exacli/cookiejar, where HOME is the home directory of the operating system user running the ExaCLI command. For the opc user, the home is /home/opc.
  - The operating system user running the ExaCLI command is the owner of the cookie jar file.
  - A cookie jar can contain multiple cookies from multiple users on multiple nodes in parallel sessions.
  - Cookies are invalidated after 24 hours.

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - If the cookie is not found or is no longer valid, ExaCLI prompts for the password. The new cookie is stored in the cookie jar identified by `filename`, or the default cookie jar if `filename` is not specified.
  - Even without the `--cookie-jar` option, ExaCLI still checks for cookies from the default cookie jar. However, if the cookie does not exist or is no longer valid, the new cookie will not be stored in the default cookie jar if the `--cookie-jar` option is not specified.
- Notes for the `-e` option:
  - ExaCLI exits after running the commands.
  - If specifying multiple commands to run, be sure to enclose the commands in single quotes to prevent the shell from interpreting the semi-colon.
  - The batch file is a text file that contains one or more ExaCLI commands to run.
- Notes for the `-n` (`--no-prompt`) option:
  - If ExaCLI needs additional information from the user, for example, if ExaCLI needs to prompt the user for a password (possibly because there were no valid cookies in the cookie-jar) or to prompt the user to confirm the remote node's identity, then ExaCLI prints an error message and exits.

Examples

Example 7-57 Starting an Interactive ExaCLI Session on a Storage Server

This example shows the user on an Exadata compute node issuing the command to log in to ExaCLI start an interactive ExaCLI session on a storage server:

```
exacli -l cloud_user_clustername -c 192.168.136.7
```

See "Finding the IP addresses of storage cells using the cellip.ora file" for information about how to determine your storage server IP address.

After you are logged in, run additional commands as follows:

```
exacli cloud_user_clustername@192.168.136.7> LIST DATABASE
ASM
HRCDB
```

Example 7-58 Issuing a Single Command on a Compute Node

This example shows a single command issued on a compute node that does the following:

- Connects to a storage server
- Performs a `LIST` action
- Exits the session (specified with the `-e` option)

```
exacli -l cloud_user_clustername -c 192.168.136.7 --xml --cookie-jar -e list
griddisk detail
```

Related Topics

- Connecting to a Storage Server with ExaCLI
  To use ExaCLI on storage servers, you will need to know your target storage server's IP address.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Connecting to a Storage Server with ExaCLI

To use ExaCLI on storage servers, you will need to know your target storage server's IP address.

If you do not know the IP address of the node you want to connect to, you can find it by viewing the contents of the `cellip.ora` file.

The following example illustrates how to do so on the UNIX command line for a quarter rack system. (Note that a quarter rack has three storage cells, and each cell has two connections, so a total of six IP addresses are shown.)

```
cat /etc/oracle/cell/
network-config/cellip.oracle
cell="192.168.136.5;cell="192.168.136.6"
cell="192.168.136.7;cell="192.168.136.8"
cell="192.168.136.9;cell="192.168.136.10"
```

If you are connecting to a storage cell for the first time using ExaCLI, you may be prompted to accept an SSL certificate. The ExaCLI output in this case will look like the following:

```
exacli -l cloud_user_clustername -c 192.168.136.7 --cookie-jar
No cookies found for cloud_user_clustername@192.168.136.7
Password: *********
EXA-30016: This connection is not secure. You have asked ExaCLI to connect to
cell 192.168.136.7 securely. The identity of 192.168.136.7 cannot be verified.
Got certificate from server:
C=US,ST=California,L=Redwood City,O=Oracle Corporation,OU=Oracle
Exadata,CN=ed1cl03clu01-priv2.usdc2.oraclecloud.com
Do you want to accept and store this certificate? (Press y/n)
```

Accept the self-signed Oracle certificate by pressing "y" to continue using ExaCLI.

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). For more information on policies, see *Getting Started with Policies*.

- Prerequisites for Using Metrics
  The following prerequisites must be met for the metrics to flow out of the VM Cluster.
- View Metrics for a VM Cluster
  To view the metrics for Guest VMs using the console, use this procedure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- View Metrics for a Database
  To view the metrics for a database using the console, use this procedure.
- View Metrics for VM Clusters in a Compartment
  To view the metrics for databases in a compartment using the console, use this procedure.
- View Metrics for Databases in a Compartment
  To view the metrics for databases in a compartment using the console, use this procedure.
- Manage Oracle Trace File Analyzer

Related Topics

- Getting Started with Policies

Prerequisites for Using Metrics

The following prerequisites must be met for the metrics to flow out of the VM Cluster.

- Metrics on the VM Clusters depend on the Oracle Trace File Analyzer (TFA) agent. Ensure that this component is up and running. AHF version 23.2.0 or higher is required for capturing metrics from the VM Clusters. To start, stop, or check the status of TFA, see *Manage Oracle Trace File Analyzer*.
- To view the metrics on the Oracle Cloud Infrastructure Console, the TFA flag `defaultocimonitoring` must be set to ON. This flag is set to ON by default and you need not perform any action to set this. If you are not seeing metrics on the Console, then as `root` user on the guest VM, check if the flag is set to ON.

```
tfactl get defaultocimonitoring
.-------------------------------------------------------------------.
|                            <host name>                            |
+-------------------------------------------------------------+-------+
| Configuration Parameter                                     | Value |
+-------------------------------------------------------------+-------+
| Send CEF metrics to OCI Monitoring ( defaultOciMonitoring ) | ON    |
'-------------------------------------------------------------+-------'
```

If the `defaultocimonitoring` flag is set to OFF, then run the `tfactl set defaultocimonitoring=on` or `tfactl set defaultocimonitoring=ON` command to turn it on:

```
tfactl set defaultocimonitoring=on
Successfully set defaultOciMonitoring=ON
.-------------------------------------------------------------------.
|                            <host name>                            |
+-------------------------------------------------------------+-------+
| Configuration Parameter                                     | Value |
+-------------------------------------------------------------+-------+
| Send CEF metrics to OCI Monitoring ( defaultOciMonitoring ) | ON    |
'-------------------------------------------------------------+-------'
```

- The network communication uses the existing path that the Control Plane Server uses to reach the OCI Services. For more information, see *Network Requirements for Oracle Exadata Database Service on Cloud@Customer*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Network Requirements for Oracle Exadata Database Service on Cloud@Customer
  Review the network requirements for provisioning Oracle Exadata Database Service on Cloud@Customer at your site.
- Manage Oracle Trace File Analyzer

View Metrics for a VM Cluster

To view the metrics for Guest VMs using the console, use this procedure.

> **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 **Exadata Database Service on Cloud@Customer**.
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. Under **Resources**, click **Metrics**.

   A chart for each metric is displayed. By default, the metrics for the last 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*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Viewing Default Metric Charts

View Metrics for a Database

To view the metrics for a database using the console, use this procedure.

> **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 **Exadata Database Service on Cloud@Customer**.
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 list of databases, click the database for which you want to view the metrics.
5. Under **Resources**, click **Metrics**.

   A chart for each metric is displayed. By default, the metrics for the last hour are displayed.

   You can only select the **oci_database** namespace from the **Metric namespace** drop-down.
6. 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.
7. 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.
8. 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 © 2008, 2026, Oracle and/or its affiliates.

Related Topics

- Viewing Default Metric Charts

View Metrics for VM Clusters in a Compartment

To view the metrics for databases in a compartment using the console, use this procedure.

> **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, then you can show metrics aggregated across the clusters by selecting **Aggregate Metric Streams**.
7. If you want to limit the metrics you see, next to **Dimensions**, click **Add** (click **Edit** if you have already added dimensions).
8. In the **Dimension Name** field, select a dimension.
9. In the **Dimension Value** field, select a value.
10. Click **Done**.
11. In the Edit dimensions dialog click **+Additional Dimension** to add an additional dimension. Click **X** to remove a dimension.
12. To create an alarm on a specific metric, click **Options** and select **Create an Alarm on this Query**. See *Managing Alarms* for information on setting and using alarms.

> **Note**
>
> If you don't see any metrics, check the network settings and AHF version listed in the prerequisites section.

Related Topics

- Managing Alarms

Copyright © 2008, 2026, Oracle and/or its affiliates.

View Metrics for Databases in a Compartment

To view the metrics for databases in a compartment using the console, use this procedure.

> **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, then you can show metrics aggregated across the databases by selecting **Aggregate Metric Streams**.
7. If you want to limit the metrics you see, next to **Dimensions**, click **Add** (click **Edit** if you have already added dimensions).
8. In the **Dimension Name** field, select a dimension.
9. In the **Dimension Value** field, select a value.
10. Click **Done**.
11. In the Edit dimensions dialog click **+Additional Dimension** to add an additional dimension. Click **X** to remove a dimension.
12. To create an alarm on a specific metric, click **Options** and select **Create an Alarm on this Query**. See *Managing Alarms* for information on setting and using alarms.

> **Note**
>
> If you don't see any metrics, check the network settings and AHF version listed in the prerequisites section.

Related Topics

- Managing 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- To check the run status of Oracle Trace File Analyzer, run the `tfactl status` command as `root` or a non-root user:

```
# tfactl status
.-----------------------------------------------------------------------------
----------------------.
| Host           | Status of TFA | PID    | Port | Version    | Build
ID             | Inventory Status|
+---------------+--------------+--------+------+------------
+---------------------+------------+
| node1          | RUNNING       |  41312 | 5000 | 22.1.0.0.0 |
22100020220310214615| COMPLETE    |
| node2          | RUNNING       | 272300 | 5000 | 22.1.0.0.0 |
22100020220310214615| COMPLETE    |
'---------------+--------------+--------+------+------------
+---------------------+------------'
```

- To start the Oracle Trace File Analyzer daemon on the local node, run the `tfactl start` command as `root`:

```
# tfactl start
Starting TFA..
Waiting up to 100 seconds for TFA to be started..
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
Successfully started TFA Process..
. . . . .
TFA Started and listening for commands
```

- To stop the Oracle Trace File Analyzer daemon on the local node, run the `tfactl stop` command as `root`:

```
# tfactl stop
Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..
```

Metrics for Oracle Exadata Database Service on Cloud@Customer in the Monitoring Service

Review the metrics emitted by the Oracle Exadata Database Service on Cloud@Customer service in the `oci_database_cluster` and `oci_database` namespaces for Oracle Databases.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Dimensions

All the metrics discussed in this topic include the following dimensions.

- **resourceId**: The OCID of the VM Cluster.
- **resourceName**: The name of the VM Cluster.

NOT_SUPPORTED

The metrics listed in the following table are automatically available for the VM cluster.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
`ASMDiskgroupUtilization`	**ASM Diskgroup Utilization**	percentage	Percentage of usable space used in a Disk Group. Usable space is the space available for growth. DATA disk group stores our Oracle database files. RECO disk group contains database files for recovery such as archives and flashback logs.	10 minutes	`deploymentType` `diskgroupName`

`CpuUtilization`	**CPU Utilization**	percentage	Percent CPU utilization.	1 minute	`hostName`
					`deploymentType`
					`resourceId_dbnode`
					`resourceName_dbnode`

`FilesystemUtilization`	**Filesystem Utilization**	percentage	Percent utilization of provisioned filesystem.	1 minute	`hostName`
					`deploymentType`
					`filesystemName`
					`resourceId_dbnode`
					`resourceName_dbnode`

`LoadAverage`	**Load Average**	number	System load average over 5 minutes.	1 minute	`hostName`
					`deploymentType`
					`resourceId_dbnode`
					`resourceName_dbnode`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
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
OcpusAllocated	**OCPU Allocated**	integer	The number of OCPUs allocated.	1 minute	deploymentType
SwapUtilization	**Swap Utilization**	percentage	Percent utilization of total swap space.	1 minute	hostName deploymentType resourceId_dbnode resourceName_dbnode

NOT_SUPPORTED

The metrics listed in the following table are automatically available for the database.

Copyright © 2008, 2026, Oracle and/or its affiliates.

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, which is two times the number of OCPUs.	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 © 2008, 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 © 2008, 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 © 2008, 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 Cloud@Customer Events

Oracle Exadata Database Service on Cloud@Customer resources emit events, which are structured messages that indicate changes in resources.

Copyright © 2008, 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 Cloud@Customer
  Learn about the event types available for Oracle Exadata Database Service on Cloud@Customer resources.
- Exadata Infrastructure Event Types
  Review the list of event types that Exadata Infrastructure instances emit.
- Exascale DB Storage Vaults Event Types
  Review the list of event types that exascale-db-storage-vaults emit.
- VM Cluster Network Event Types
  Review the list of event types that VM cluster networks emit.
- VM Cluster Event Types
  Review the list of event types that VM clusters emit.
- Backup Destination Event Types
  Review the list of event types that backup destinations emit.
- Database Node Event Types (Oracle Exadata Database Service on Cloud@Customer)
  Review the list of event types that database nodes emit.
- Database Home Event Types (Oracle Exadata Database Service on Cloud@Customer)
  Review the list of event types that Database Homes emit.
- Database Event Types (Oracle Exadata Database Service on Cloud@Customer)
  Review the list of event types that databases emit.
- Database and Grid Infrastructure Patching Event Types
  Review the list of event types that Database and Grid Infrastructure Patching emit.
- Data Guard Event Types
  Review the list of event types that Data Gurad groups and Data Guard Associations emit.
- Key Store Event Types
  Review the list of event types that Key Store emits.
- Exadata Infrastructure Maintenance Event Types
  Review the list of event types that Exadata Infrastructure Maintenance emits.
- Storage Expansion Event Types
  Review the list of event types that storage expansion emits.
- Database Software Images Event Types
  Review the list of event types that Database Software Image emits.
- Database Upgrade Event Types
  Review the list of event types that Database Upgrade emit.
- Pluggable Database Event Types
  Review the list of event types that Pluggable Databases emit.
- VM Node Subsetting Event Types
  Review the list of event types that VM Node Subsetting emits.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Database Service Events
  The Database Service emits events, which are structured messages that indicate changes in resources.
- Compute Expansion Event Types
  Review the list of event types that compute expansion emits.
- Serial Console Connection Event Types
  Review the list of event types that serial console connection emits.
- Serial Console History Event Types
  Review the list of new event types that serial console history emits.
- Interim Software Updates Event Types

About Event Types on Oracle Exadata Database Service on Cloud@Customer

Learn about the event types available for Oracle Exadata Database Service on Cloud@Customer resources.

Oracle Exadata Database Service on Cloud@Customer 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

Exadata Infrastructure Event Types

Review the list of event types that Exadata Infrastructure instances emit.

Table 7-2 Exadata Infrastructure Event Types

Friendly Name	Event Type
Activate Begin	com.oraclecloud.databaseservice.activateexadatainfrastructure.begin
Activate End	com.oraclecloud.databaseservice.activateexadatainfrastructure.end
Change Compartment	com.oraclecloud.databaseservice.changeexadatainfrastructurecompartment
Configuration File Download	com.oraclecloud.databaseservice.downloadexadatainfrastructureconfigfile
Create Begin	com.oraclecloud.databaseservice.createexadatainfrastructure.begin
Create End	com.oraclecloud.databaseservice.createexadatainfrastructure.end
Delete Begin	com.oraclecloud.databaseservice.deleteexadatainfrastructure.begin
Delete End	com.oraclecloud.databaseservice.deleteexadatainfrastructure.end

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-2 (Cont.) Exadata Infrastructure Event Types

Friendly Name	Event Type
Update Begin	com.oraclecloud.databaseservice.updateexadatainfrastructure.begin
Update End	com.oraclecloud.databaseservice.updateexadatainfrastructure.end
Exadata Infrastructure - Connectivity Status	com.oraclecloud.databaseservice.exadatainfrastructureconnectstatus
Exadata Infrastructure - Maintenance Custom action time Begin (ROLLING)	com.oraclecloud.databaseservice.exacinfrastructuremaintenancecustomactiontime.begin
Exadata Infrastructure - Maintenance Custom action time End (ROLLING)	com.oraclecloud.databaseservice.exacinfrastructuremaintenancecustomactiontime.end
Exadata Infrastructure - IB Switch maintenance Begin	com.oraclecloud.databaseservice.exacinfrastructuremaintenanceibswitch.begin
Exadata Infrastructure - IB Switch maintenance End	com.oraclecloud.databaseservice.exacinfrastructuremaintenanceibswitch.end
Exadata Infrastructure - Storage server maintenance Begin	com.oraclecloud.databaseservice.exacinfrastructuremaintenancestorageservers.begin
Exadata Infrastructure - Storage server maintenance End	com.oraclecloud.databaseservice.exacinfrastructuremaintenancestorageservers.end
Exadata Infrastructure - Configure Exascale Begin	com.oraclecloud.databaseservice.configureexascalecloudexadatainfrastructure.begin
Exadata Infrastructure - Configure Exascale End	com.oraclecloud.databaseservice.configureexascalecloudexadatainfrastructure.end

Example 7-59 Exadata Infrastructure Example

This is a reference event for Exadata Infrastructure instances:

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.createexadatainfrastructure.begin",
    "source": "databaseservice",
    "eventTypeVersion": "version",
    "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_name",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "resourceName": "my_exadata_infra",
      "resourceId": "ExadataInfra-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "shape": "ExadataCC.Base3.48",
        "timeZone": "US/Pacific",
        "displayName": "testDisplayName"
      }
    }
  }
```

This is a reference event for Exadata Infrastructure - Connectivity Status:

```
{
  "eventType" :
"com.oraclecloud.databaseservice.exadatainfrastructureconnectstatus",
  "cloudEventsVersion" : "0.1",
  "eventTypeVersion" : "2.0",
  "source" : "DatabaseService",
  "eventTime" : "2020-06-02T06:07:40.141Z",
  "contentType" : "application/json",
  "data" : {
    "compartmentId" :
"ocid1.compartment.oc1..aaaaaaaayrygl4guhplo5rtiumx3eh4mk7grrkrqspzaltmvbxecnb
vhkrga",
    "compartmentName" : "DBaaSInteg20160918ExaccTest",
    "resourceName" : "MVM_HR",
    "resourceId" : "ocid1.exadatainfrastructure.oc1.ap-
hyderabad-1.abuhsljrp2vzzenmqmctqciwro6euhhsmlrewiiemiktov5xyfsu5hiufjsq",
    "availabilityDomain" : "",
    "additionalDetails" : {
      "timeCreated" : "2020-05-28T00:23:18Z",
      "timeUpdated" : "2020-06-02T06:07:40Z",
      "lifecycleState" : "DISCONNECTED",
      "lifecycleDetails" : "Exadata Infrastructure is not reachable. Please
lodge a Service Request (SR) with Oracle Support and provide Infrastructure-
id: ocid1.exadatainfrastructure.oc1.ap-
hyderabad-1.abuhsljrp2vzzenmqmctqciwro6euhhsmlrewiiemiktov5xyfsu5hiufjsq.",
      "shape" : "ExadataCC.Half3.200",
      "timeZone" : "UTC"
    },
    "definedTags" : {
      "Oracle-Tags" : {
        "CreatedBy" : "test-user@example.com",
        "CreatedOn" : "2020-05-28T00:23:18.291Z"
      }
    }
  },
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "eventID" : "cde92d45-a06b-4b69-a125-6005dd8c2f0c",
  "extensions" : {
    "compartmentId" :
"ocid1.compartment.oc1..aaaaaaaayrygl4guhplo5rtiumx3eh4mk7grrkrqspzaltmvbxecnb
vhkrga"
  }
}
```

This is a reference event for Exadata Infrastructure - Custom action time Begin:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqlasasjrs5agc5dbs2dupjpkv56dyepsuqkdbyklzaei7kk
ut442gd7taksq5q",
  "serviceName": "Database",
  "displayName": "Exadata Infrastructure - Custom action time Begin",
  "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancecustomactiontim
e.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",
        "string"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeScheduled",
      "type": "string"
    },
    {
      "name": "timeStarted",
      "type": "string"
    },
    {
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "name": "availabilityDomain",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "lifecycleState",
        "type": "string"
      },
      {
        "name": "message",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "description",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "shape",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "displayName",
        "type": [
          "null",
          "int"
        ]
      }
    ],
    "exampleEvent": {
      "cloudEventsVersion": "0.1",
      "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
      "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancecustomactiontim
e.begin",
      "source": "databaseservice",
      "eventTypeVersion": "2.0",
      "eventTime": "2019-06-27T21:16:04.000Z",
      "contentType": "application/json",
      "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "exaccinfrastructuremaintenancecustomactiontime",
        "compartmentId": "ocid1.compartment.oc1.......unique_id",
        "compartmentName": "example_compartment",
        "resourceName": "my_exacc_infrastructure",
        "resourceId": "ocid1.exadatainfrastructure.oc1.....unique_id",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbmaintenancerun.oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeScheduled": "2019-08-29T12:30:00.000Z",
        "timeStarted": "2019-08-29T12:30:00.000Z",
        "description": "ExaCC Infrastructure Maintenance Notifications",
        "message": "detail message",
        "shape": "ExadataCC.Base3.48",
        "displayName": "testDisplayName"
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


This is a reference event for Exadata Infrastructure - Custom action time End:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrs5agc5dbs2dupjpkv56dyepsuqdfdfkdbyklzaei7kk
ut442gd7taksq5q",
  "serviceName": "Database",
  "displayName": "Exadata Infrastructure - Custom action time End",
  "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancecustomactiontim
e.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",
        "string"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "timeCreated",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "type": "string"
      },
      {
        "name": "timeScheduled",
        "type": "string"
      },
      {
        "name": "timeStarted",
        "type": "string"
      },
      {
        "name": "availabilityDomain",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "lifecycleState",
        "type": "string"
      },
      {
        "name": "message",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "description",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "shape",
        "type": [
          "null",
          "string"
        ]
      },
      {
        "name": "displayName",
        "type": [
          "null",
          "int"
        ]
      }
    ],
    "exampleEvent": {
      "cloudEventsVersion": "0.1",
      "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
      "eventType":
  "com.oraclecloud.databaseservice.exaccinfrastructuremaintenancecustomactiontim
  e.end",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "exaccinfrastructuremaintenancecustomactiontime",
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "example_compartment",
      "resourceName": "my_exacc_infrastructure",
      "resourceId": "ocid1.exadatainfrastructure.oc1.....unique_id",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbmaintenancerun.oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeScheduled": "2019-08-29T12:30:00.000Z",
        "timeStarted": "2019-08-29T12:30:00.000Z",
        "description": "ExaCC Infrastructure Maintenance Notifications",
        "message": "detail message",
        "shape": "ExadataCC.Base3.48",
        "displayName": "testDisplayName"
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


This is a reference event for Exadata Infrastructure - IB Switch maintenance Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrs5agc5dbs2dupjpkv56asasdyepsuqkdbyklzaei7kk
ut442gd7taksq5q",
  "serviceName": "Database",
  "displayName": "Exadata Infrastructure - IB Switch maintenance Begin",
  "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenanceibswitch.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",
        "string"
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    ]
  },
  {
    "name": "freeFormTags",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "timeCreated",
    "type": "string"
  },
  {
    "name": "timeScheduled",
    "type": "string"
  },
  {
    "name": "timeStarted",
    "type": "string"
  },
  {
    "name": "availabilityDomain",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "lifecycleState",
    "type": "string"
  },
  {
    "name": "message",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "description",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "shape",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "displayName",
    "type": [
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "null",
        "int"
      ]
    }
  ],
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b98",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenanceibswitch.begin"
,
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "ExaccInfrastructureMaintenanceIBSwitch",
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "example_compartment",
      "resourceName": "my_exacc_infrastructure",
      "resourceId": "ocid1.exadatainfrastructure.oc1.....unique_id",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbmaintenancerun.oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeScheduled": "2019-08-29T12:30:00.000Z",
        "timeStarted": "2019-08-29T12:30:00.000Z",
        "description": "ExaCC Infrastructure Maintenance Notifications",
        "message": "detail message",
        "shape": "ExadataCC.Base3.48",
        "displayName": "testDisplayName"
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


This is a reference event for Exadata Infrastructure - IB Switch maintenance End:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrs5agc5dbs2dupjpkv56dyepsusdsdqkdbyklzaei7kk
ut442gd7taksq5q",
  "serviceName": "Database",
  "displayName": "Exadata Infrastructure - IB Switch maintenance End",
  "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenanceibswitch.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
"additionalDetails": [
  {
    "name": "id",
    "type": "string"
  },
  {
    "name": "defineTags",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "freeFormTags",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "timeCreated",
    "type": "string"
  },
  {
    "name": "timeScheduled",
    "type": "string"
  },
  {
    "name": "timeStarted",
    "type": "string"
  },
  {
    "name": "availabilityDomain",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "lifecycleState",
    "type": "string"
  },
  {
    "name": "message",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "description",
    "type": [
      "null",
      "string"
    ]
  },
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    {
      "name": "shape",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "displayName",
      "type": [
        "null",
        "int"
      ]
    }
  ],
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b98",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenanceibswitch.end",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "ExaccInfrastructureMaintenanceIBSwitch",
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "example_compartment",
      "resourceName": "my_exacc_infrastructure",
      "resourceId": "ocid1.exadatainfrastructure.oc1.....unique_id",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbmaintenancerun.oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeScheduled": "2019-08-29T12:30:00.000Z",
        "timeStarted": "2019-08-29T12:30:00.000Z",
        "description": "ExaCC Infrastructure Maintenance Notifications",
        "message": "detail message",
        "shape": "ExadataCC.Base3.48",
        "displayName": "testDisplayName"
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


This is a reference event for Exadata Infrastructure - Storage server maintenance Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrsasas5agc5dbs2dupjpkv56dyepsuqkdbyklzaei7kk
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
ut442gd7taksq5q",
  "serviceName": "Database",
  "displayName": "Exadata Infrastructure - Storage server maintenance Begin",
  "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancestorageservers.
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",
        "string"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeScheduled",
      "type": "string"
    },
    {
      "name": "timeStarted",
      "type": "string"
    },
    {
      "name": "availabilityDomain",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "message",
      "type": [
        "null",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
            "string"
          ]
        },
        {
          "name": "description",
          "type": [
            "null",
            "string"
          ]
        },
        {
          "name": "shape",
          "type": [
            "null",
            "string"
          ]
        },
        {
          "name": "displayName",
          "type": [
            "null",
            "int"
          ]
        }
      ],
      "exampleEvent": {
        "cloudEventsVersion": "0.1",
        "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b55",
        "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancestorageservers.
begin",
        "source": "databaseservice",
        "eventTypeVersion": "2.0",
        "eventTime": "2019-06-27T21:16:04.000Z",
        "contentType": "application/json",
        "data": {
          "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
          "eventName": "ExaccInfrastructureMaintenanceStorageServers",
          "compartmentId": "ocid1.compartment.oc1.......unique_id",
          "compartmentName": "example_compartment",
          "resourceName": "my_exacc_infrastructure",
          "resourceId": "ocid1.exadatainfrastructure.oc1.....unique_id",
          "availabilityDomain": "all",
          "freeFormTags": {},
          "definedTags": {},
          "additionalDetails": {
            "id": "ocid1.dbmaintenancerun.oc1...unique_ID",
            "lifecycleState": "AVAILABLE",
            "timeCreated": "2019-08-29T12:00:00.000Z",
            "timeScheduled": "2019-08-29T12:30:00.000Z",
            "timeStarted": "2019-08-29T12:30:00.000Z",
            "description": "ExaCC Infrastructure Maintenance Notifications",
            "message": "detail message",
            "shape": "ExadataCC.Base3.48",
            "displayName": "testDisplayName"
          }
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```

This is a reference event for Exadata Infrastructure - Storage server maintenance End:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljasrs5agddc5dbs2dupjpkv56dyepsuqkdbyklzaei7kk
ut442gd7taksq5q",
  "serviceName": "Database",
  "displayName": "Exadata Infrastructure - Storage server maintenance End",
  "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancestorageservers.
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",
        "string"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeScheduled",
      "type": "string"
    },
    {
      "name": "timeStarted",
      "type": "string"
    },
    {
      "name": "availabilityDomain",
      "type": [
        "null",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
            "string"
          ]
        },
        {
          "name": "lifecycleState",
          "type": "string"
        },
        {
          "name": "message",
          "type": [
            "null",
            "string"
          ]
        },
        {
          "name": "description",
          "type": [
            "null",
            "string"
          ]
        },
        {
          "name": "shape",
          "type": [
            "null",
            "string"
          ]
        },
        {
          "name": "displayName",
          "type": [
            "null",
            "int"
          ]
        }
      ],
      "exampleEvent": {
        "cloudEventsVersion": "0.1",
        "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b55",
        "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancestorageservers.
end",
        "source": "databaseservice",
        "eventTypeVersion": "2.0",
        "eventTime": "2019-06-27T21:16:04.000Z",
        "contentType": "application/json",
        "data": {
          "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
          "eventName": "ExaccInfrastructureMaintenanceStorageServers",
          "compartmentId": "ocid1.compartment.oc1.......unique_id",
          "compartmentName": "example_compartment",
          "resourceName": "my_exacc_infrastructure",
          "resourceId": "ocid1.exadatainfrastructure.oc1.....unique_id",
          "availabilityDomain": "all",
          "freeFormTags": {},
          "definedTags": {},
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "additionalDetails": {
        "id": "ocid1.dbmaintenancerun.oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeScheduled": "2019-08-29T12:30:00.000Z",
        "timeStarted": "2019-08-29T12:30:00.000Z",
        "description": "ExaCC Infrastructure Maintenance Notifications",
        "message": "detail message",
        "shape": "ExadataCC.Base3.48",
        "displayName": "testDisplayName"
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}
```


Exascale DB Storage Vaults Event Types

Review the list of event types that exascale-db-storage-vaults emit.

Table 7-3 Exascale DB Storage Vaults Event Types

Friendly Name	Event Type
Exascale Database Storage Vault - Create Begin	com.oraclecloud.databaseservice.createexascaledbstoragevault.begin
Exascale Database Storage Vault - Create End	com.oraclecloud.databaseservice.createexascaledbstoragevault.end
Exascale Database Storage Vault - Update Begin	com.oraclecloud.databaseservice.updateexascaledbstoragevault.begin
Exascale Database Storage Vault - Update End	com.oraclecloud.databaseservice.updateexascaledbstoragevault.end
Exascale Database Storage Vault - Delete Begin	com.oraclecloud.databaseservice.deleteexascaledbstoragevault.begin
Exascale Database Storage Vault - Delete End	com.oraclecloud.databaseservice.deleteexascaledbstoragevault.end
Exascale Database Storage Vault - Change Compartment Begin	com.oraclecloud.databaseservice.changeexascaledbstoragevaultcompartment.begin
Exascale Database Storage Vault - Change Compartment End	com.oraclecloud.databaseservice.changeexascaledbstoragevaultcompartment.end

VM Cluster Network Event Types

Review the list of event types that VM cluster networks emit.

Table 7-4 VM Cluster Network Event Types

Friendly Name	Event Type
Create Begin	com.oraclecloud.databaseservice.createvmclusternetwork.begin

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-4 (Cont.) VM Cluster Network Event Types

Friendly Name	Event Type
Create End	com.oraclecloud.databaseservice.createvmclusternetwork.end
Network Validation File Download	com.oraclecloud.databaseservice.downloadvmclusternetworkconfigfile
Terminate Begin	com.oraclecloud.databaseservice.deletevmclusternetwork.begin
Terminate End	com.oraclecloud.databaseservice.deletevmclusternetwork.end
Update Begin	com.oraclecloud.databaseservice.updatevmclusternetwork.begin
Update End	com.oraclecloud.databaseservice.updatevmclusternetwork.end
Validate Begin	com.oraclecloud.databaseservice.validatevmclusternetwork.begin
Validate End	com.oraclecloud.databaseservice.validatevmclusternetwork.end

Example 7-60 VM Cluster Network Example

This is a reference event for VM cluster networks:

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.createvmclusternetwork.begin",
    "source": "databaseservice",
    "eventTypeVersion": "version",
    "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_name",
      "resourceName": "my_vmcluster_network",
      "resourceId": "VmClusterNetwork-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "exadataInfrastructureId": "ExadataInfra-unique_ID",
        "displayName": "testDisplayName"
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      }
    }
  }
```

VM Cluster Event Types

Review the list of event types that VM clusters emit.

Table 7-5 VM Cluster Event Types

Friendly Name	Event Type
Change Compartment	com.oraclecloud.databaseservice.changevmclustercompartment
Create Begin	com.oraclecloud.databaseservice.createvmcluster.begin
Create End	com.oraclecloud.databaseservice.createvmcluster.end
Terminate Begin	com.oraclecloud.databaseservice.deletevmcluster.begin
Terminate End	com.oraclecloud.databaseservice.deletevmcluster.end
Update Begin	com.oraclecloud.databaseservice.updatevmcluster.begin
Update End	com.oraclecloud.databaseservice.updatevmcluster.end

Example 7-61 VM Cluster Example

This is a reference event for VM clusters:

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.createvmclusternetwork.begin",
    "source": "databaseservice",
    "eventTypeVersion": "version",
    "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_name",
      "resourceName": "my_vmcluster_network",
      "resourceId": "VmClusterNetwork-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "exadataInfrastructureId": "ExadataInfra-unique_ID",
        "displayName": "testDisplayName"
      }
    }
  }
```

This is a reference event for VM Cluster Create Begin:

```
{
  "cloudEventsVersion": "0.1",
  "eventId": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType": "com.oraclecloud.databaseservice.createvmcluster.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,
      "dbVersion": "19.0.0.0",
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "giVersion": "19.0.0.0",
      "dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
      "timeZone": "US/Pacific"
    }
  }
}
```

Backup Destination Event Types

Review the list of event types that backup destinations emit.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-6 Backup Destination Event Types

Friendly Name	Event Type
Change Compartment	com.oraclecloud.databaseservice.changebackupdestinationcompartment
Create	com.oraclecloud.databaseservice.createbackupdestination
Terminate	com.oraclecloud.databaseservice.deletebackupdestination
Update	com.oraclecloud.databaseservice.updatebackupdestination

Example 7-62 Backup Destination Example

This is a reference event for backup destinations:

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.createbackupdestination",
    "source": "databaseservice",
    "eventTypeVersion": "version",
    "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_name",
      "resourceName": "my_backupdestination",
      "resourceId": "BackupDestination-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {}
    }
  }
```

Database Node Event Types (Oracle Exadata Database Service on Cloud@Customer)

Review the list of event types that database nodes emit.

Table 7-7 Database Node Event Types

Friendly Name	Event Type
Update Begin	com.oraclecloud.databaseservice.dbnodeaction.begin
Update End	com.oraclecloud.databaseservice.dbnodeaction.end

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-63 Database Node Example

This is a reference event for database nodes:

```
{
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType": "com.oraclecloud.databaseservice.dbnodeaction.begin",
  "source": "databaseservice",
  "eventTypeVersion": "version",
  "eventTime": "2019-06-27T21:16:04Z",
  "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"
    }
  }
}
```

Database Home Event Types (Oracle Exadata Database Service on Cloud@Customer)

Review the list of event types that Database Homes emit.

Table 7-8 Database Home Event Types

Friendly Name	Event Type
Create Begin	com.oraclecloud.databaseservice.created bhome.begin
Create End	com.oraclecloud.databaseservice.created bhome.end
Terminate Begin	com.oraclecloud.databaseservice.deleted bhome.begin

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-8 (Cont.) Database Home Event Types

Friendly Name	Event Type
Terminate End	com.oraclecloud.databaseservice.deletedbhome.end
Update Begin	com.oraclecloud.databaseservice.updatedbhome.begin
Update End	com.oraclecloud.databaseservice.updatedbhome.end

Example 7-64 Database Home Example

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": "version",
    "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_name",
      "resourceName": "my_dbhome",
      "resourceId": "DbHome-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "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": "testDisplayName"
      }
    }
  }
```

Database Event Types (Oracle Exadata Database Service on Cloud@Customer)

Review the list of event types that databases emit.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-9 Database Event Types

Friendly Name	Event Type
Create Begin	com.oraclecloud.databaseservice.createdatabase.begin
Create End	com.oraclecloud.databaseservice.createdatabase.end
Delete Backup Begin	com.oraclecloud.databaseservice.deletebackup.begin
Delete Backup End	com.oraclecloud.databaseservice.deletebackup.end
Restore Begin	com.oraclecloud.databaseservice.restoredatabase.begin
Restore End	com.oraclecloud.databaseservice.restoredatabase.end
Terminate Begin	com.oraclecloud.databaseservice.deletedatabase.begin
Terminate End	com.oraclecloud.databaseservice.deletedatabase.end
Update Begin	com.oraclecloud.databaseservice.updatedatabase.begin
Update End	com.oraclecloud.databaseservice.updatedatabase.end

Example 7-65 Database Example

This is a reference event for databases:

```
{
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType": "com.oraclecloud.databaseservice.restoredatabase.begin",
  "source": "databaseservice",
  "eventTypeVersion": "version",
  "eventTime": "2019-06-27T21:16:04Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Database-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-08-26T12:00:00.000Z",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "timeUpdated": "2019-08-26T12:30:00.000Z",
      "dbSystemId": "dbSystem-unique_ID",
      "displayName": "testDisplayName",
      "lifecycleDetails": "detail message",
      "vmClusterId": "VmCluster-unique_ID",
      "backupType": "FULL",
      "dbHomeId": "dbHome-unique_ID",
      "dbVersion": "19.0.0.0",
      "databaseEdition": "ENTERPRISE_EDITION_EXTREME",
      "autoBackupsEnabled": "true",
      "recoveryWindow": 30,
      "backupDestinationId": "backupDestination-unique_ID",
      "backupDestinationType": "OBJECT_STORAGE",
      "backupDestinationName": "my_backup_destination_name",
      "exadataInfrastructureId": "ExadataInfrastructure-unique_ID",
      "dbUniqueName": "akv_tgh_unqna"
    }
  }
}
```

Database and Grid Infrastructure Patching Event Types

Review the list of event types that Database and Grid Infrastructure Patching emit.

Table 7-10 Database and Grid Infrastructure Patching Events

Friendly Name	Event Type
VM Cluster - Patch Begin	com.oraclecloud.databaseservice.patchvmcluster.begin
VM Cluster - Patch End	com.oraclecloud.databaseservice.patchvmcluster.end
DB Home - Patch Begin	com.oraclecloud.databaseservice.patchdbhome.begin
DB Home - Patch End	com.oraclecloud.databaseservice.patchdbhome.end
Database - Move Begin	com.oraclecloud.databaseservice.movedatabase.begin
Database - Move End	com.oraclecloud.databaseservice.movedatabase.end

Example 7-66 Database Example

This is a reference event for VM Cluster - Patch Begin:

```
{
    "cloudEventsVersion": "0.1",
    "eventId": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.patchvmcluster.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "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,
        "dbVersion": "19.0.0.0",
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "giVersion": "19.0.0.0",
        "dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
        "timeZone": "US/Pacific"
      }
    }
  }
```


This is a reference event for DB Home - Patch Begin:


```
{
    "cloudEventsVersion": "0.1",
    "eventId": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.patchdbhome.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",
      "compartmentName": "example_name",
      "resourceName": "my_dbhome",
      "resourceId": "DbHome-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2019-08-29T12:00:00.000Z",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "dbSystemId": "DbSystem-unique_ID",
        "dbVersion": "19.0.0.0",
        "recordVersion": 4,
        "displayName": "testDisplayName"
      }
    }
  }
```

This is a reference event for Database - Move Begin:

```
{
    "cloudEventsVersion": "0.1",
    "eventId": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.movedatabase.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": "Database-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",
        "displayName": "testDisplayName",
        "lifecycleDetails": "detail message",
        "vmClusterId": "VmCluster-unique_ID",
        "dbSystemId": "dbSystem-unique_ID",
        "backupType": "FULL",
        "dbHomeId": "dbHome-unique_ID",
        "dbVersion": "19.0.0.0",
        "databaseEdition": "ENTERPRISE_EDITION_EXTREME",
        "autoBackupsEnabled": "true",
        "recoveryWindow": 30,
        "backupDestinationId": "backupDestination-unique_ID",
        "backupDestinationType": "OBJECT_STORAGE",
        "backupDestinationName": "my_backup_destination_name",
        "exadataInfrastructureId": "ExadataInfrastructure-unique-ID",
        "dbUniqueName": "akv_tgh_unqna"
      }
    }
  }
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Data Guard Event Types

Review the list of event types that Data Gurad groups 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 groups emit.

Table 7-11 Data Guard Group Events

Friendly Name	Event Type
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
Refresh Data Guard Health Status - Begin	com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.begin
Refresh Data Guard Health Status - End	com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.end

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-67 Data Guard Health Status

This is a reference event for Refresh Data Guard Health Status - Begin:


```
"exampleEvent": {
  "eventID": "022a63a4-ff77-11e9-a0af-f45c89b1cb17",
  "eventTime": "2024-12-04T10:06:21.000Z",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.begin",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "displayName": "Refresh Data Guard Health Status - Begin",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_id",
    "compartmentName": "example_name",
    "resourceName": "my_standby_database",
    "resourceId": "ocid1.database.oc1.phx.unique_id",
    "availabilityDomain": "AD1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.database.oc1.phx.unique_id",
      "timeCreated": "2025-10-04T10:06:21.000Z",
      "timeUpdated": "2025-10-04T10:06:21.000Z",
      "lifecycleState": "AVAILABLE",
      "lifecycleMessage": null,
      "dbSystemId": "ocid1.dbsystem.oc1.phx.unique_id",
      "databaseId": "ocid1.database.oc1.phx.unique_id",
      "dbHomeId": "ocid1.dbhome.oc1.phx.unique_id",
      "lastSyncedTime": "2025-10-04T10:06:21.000Z",
      "applyLag": "0 seconds",
      "syncState": "SYNCED",
      "switchoverReadiness": "HEALTHY",
      "switchoverReadinessMessage": null,
      "failoverReadiness": "HEALTHY",
      "failoverReadinessMessage": null,
      "dataLossExposure": "0 seconds",
      "memberTimeUpdated": "2025-10-04T10:06:21.000Z",
      "lastUpdatedIdentifier": "022a6912-ff77-11e9-9e77-f45c89b1cb17",
      "currentDataGuardRole" : "STANDBY",
      "noOfStandbyDatabases" : 3,
      "primaryDatabaseId" : "ocid1.database.oc1.iad.unique_id",
      "protectionMode" : "MAXIMUM_AVAILABILITY",
      "transportType" : "SYNC"
    }
  }
}
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

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"
    }
  }
}
```


Data Guard Event Types (Data Guard Associations resource)

Review the list of event types that Data Guard Associations emit.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-12 Data Guard Associations Events

Friendly Name	Event Type
Data Guard Association - Create Begin	com.oraclecloud.databaseservice.createdataguardassociation.begin
Data Guard Association - Create End	com.oraclecloud.databaseservice.createdataguardassociation.end
Switchover Begin	com.oraclecloud.databaseservice.switchoverdataguardassociation.begin
Switchover End	com.oraclecloud.databaseservice.switchoverdataguardassociation.end
Failover Begin	com.oraclecloud.databaseservice.failoverdataguardassociation.begin
Failover End	com.oraclecloud.databaseservice.failoverdataguardassociation.end
Reinstate Begin	com.oraclecloud.databaseservice.reinstatedataguardassociation.begin
Reinstate End	com.oraclecloud.databaseservice.reinstatedataguardassociation.end

Key Store Event Types

Review the list of event types that Key Store emits.

Table 7-13 Key Store Events

Friendly Name	Event Type
Key Store - Create	com.oraclecloud.databaseservice.createkeystore
Key Store - Update	com.oraclecloud.databaseservice.updatekeystore
Key Store - Terminate	com.oraclecloud.databaseservice.deletekeystore
Key Store - Change Compartment	com.oraclecloud.databaseservice.changekeystorecompartment
Autonomous Container Database - Rotate Key Begin	com.oraclecloud.databaseservice.rotatekeyautonomouscontainerdatabase.begin
Autonomous Container Database - Rotate Key End	com.oraclecloud.databaseservice.rotatekeyautonomouscontainerdatabase.end
Autonomous Database - Rotate Key Begin	com.oraclecloud.databaseservice.rotatekeyautonomousdatabase.begin
Autonomous Database - Rotate Key End	com.oraclecloud.databaseservice.rotatekeyautonomousdatabase.end

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-68 Key Store Example

This is a reference event for Create Key Store:


```
{
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.createkeystore",
    "source": "databaseservice",
    "eventTypeVersion": "version",
    "eventTime": "2020-10-27T21:16:04Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_keystore",
      "resourceId": "KeyStore-unique_ID",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "timeUpdated"="2020-10-27 21:16:34.864",
        "timeCreated"="2020-10-27 21:16:34.864",
        "keyStoreType": "all",
        "connectionIps": "ip1,ip2",
        "adminUsername": "username",
      }
    }
  }
```


Exadata Infrastructure Maintenance Event Types

Review the list of event types that Exadata Infrastructure Maintenance emits.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-14 Exadata Cloud@Customer Infrastructure Maintenance Events

Friendly Name	Event Type	Event Message
Exadata Infrastructure - Maintenance Reminder	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancereminder	This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-C@C Exadata Infrastructure <infra-name>, OCID <infra-ocid> part of Maintenance Run <mr-display-name>, OCID <mr-ocid> in approximately NO_OF_DAYS_LEFT days on SCHEDULE_TIME. The maintenance includes EA1, Full DB Server update for DB servers 1, 3, rolling, without custom action time. EA2, Full server update for DB server 2,4, non-rolling with 30 minutes custom action time between server selected per the scheduling plan. You will get a notification on the start of the quarterly maintenance update.
Exadata Infrastructure - Maintenance Begin	com.oraclecloud.databaseservice.exaccinfrastructuremaintenance.begin	This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-C@C 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 © 2008, 2026, Oracle and/or its affiliates.

Table 7-14 (Cont.) Exadata Cloud@Customer Infrastructure Maintenance Events

Friendly Name	Event Type	Event Message
Exadata Infrastructure - Maintenance End	com.oraclecloud.databaseservice.exaccinfrastructuremaintenance.end	**Success:** This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-C@C 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-C@C 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-C@C 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-C@C 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

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-14 (Cont.) Exadata Cloud@Customer Infrastructure Maintenance Events

Friendly Name	Event Type	Event Message
		was configured for a `WINDOW_DURATION` duration. Your maintenance is taking longer than the configured window duration. This window has 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.
`Exadata Infrastructure - Maintenance Custom action time Begin (ROLLING)`	`com.oraclecloud.databaseservice.exaccinfrastructuremaintenancecustomactiontime.begin`	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-C@C 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.
`Exadata Infrastructure - Maintenance Custom action time End (ROLLING)`	`com.oraclecloud.databaseservice.exaccinfrastructuremaintenancecustomactiontime.end`	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-C@C 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>`.
`Exadata Infrastructure - Maintenance Custom action time Begin (NONROLLING)`	`com.oraclecloud.databaseservice.exaccinfrastructuremaintenancecustomactiontime.begin`	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-C@C 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-14 (Cont.) Exadata Cloud@Customer Infrastructure Maintenance Events

Friendly Name	Event Type	Event Message
Exadata Infrastructure - Maintenance Custom action time End (NONROLLING)	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancecustomactiontime.end	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-C@C 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>.
Exadata Infrastructure - IB Switch maintenance Begin	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancenetworkswitches.begin	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Network Switches of your ExaDB-C@C Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Network switches. Your maintenance update started at <start-time>. You will get a notification on the completion of the quarterly maintenance update for the Network Switches.
Exadata Infrastructure - IB Switch maintenance End	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancenetworkswitches.end	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Network Switches of your ExaDB-C@C Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Network switches. Your maintenance update started at <start-time> and was completed successfully at <end-time>
Exadata Infrastructure - maintenance Storage servers Start	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancestorageservers.start	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Storage Servers of your ExaDB-C@C 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-14 (Cont.) Exadata Cloud@Customer Infrastructure Maintenance Events

Friendly Name	Event Type	Event Message
Exadata Infrastructure - maintenance Storage servers End	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancestorageservers.end	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Storage Servers of your ExaDB-C@C 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>.
Exadata Infrastructure - maintenance Database Servers Rolling Start	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancedbservers.start	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-C@C 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.
Exadata Infrastructure - maintenance Database Servers Rolling End	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancedbservers.end	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-C@C 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 © 2008, 2026, Oracle and/or its affiliates.

Table 7-14 (Cont.) Exadata Cloud@Customer Infrastructure Maintenance Events

Friendly Name	Event Type	Event Message
Exadata Infrastructure - maintenance Database Servers Non Rolling Start	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancedbservers.start	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-C@C 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.
Exadata Infrastructure - maintenance Database Servers Non Rolling End	com.oraclecloud.databaseservice.exaccinfrastructuremaintenancedbservers.end	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-C@C 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>.

Storage Expansion Event Types

Review the list of event types that storage expansion emits.

Table 7-15 Storage Expansion Events

Friendly Name	Event Type
Exadata Infrastructure - Add Storage Capacity Begin	com.oraclecloud.databaseservice.addstoragecapacityexadatainfrastructure.begin
Exadata Infrastructure - Add Storage Capacity End	com.oraclecloud.databaseservice.addstoragecapacityexadatainfrastructure.end

Example 7-69 Storage Expansion Events Examples

This is a reference event for Exadata Infrastructure - Add Storage Capacity Begin:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.addstoragecapacityexadatainfrastructure.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",
    "compartmentName": "example_name",
    "resourceName": "my_exadata_infra",
    "resourceId": "ExadataInfra-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "shape": "ExadataCC.X8",
      "timeZone": "US/Pacific",
      "displayName": "testDisplayName"
    }
  }
}
```


This is a reference event for Exadata Infrastructure - Add Storage Capacity End:


```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.addstoragecapacityexadatainfrastructure.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",
    "compartmentName": "example_name",
    "resourceName": "my_exadata_infra",
    "resourceId": "ExadataInfra-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "id": "ocid1.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "shape": "ExadataCC.X8",
      "timeZone": "US/Pacific",
      "displayName": "testDisplayName"
    }
  }
}
```

Database Software Images Event Types

Review the list of event types that Database Software Image emits.

Table 7-16 Database Software Image Events

Action	Event Name	Event Type
Create Database Software Image	CreateDatabaseSoftwareImage CreateDatabaseSoftwareImage	com.oraclecloud.DatabaseService.CreateDatabaseSoftwareImage.begin com.oraclecloud.DatabaseService.CreateDatabaseSoftwareImage.end
Delete Database Software Image	DeleteDatabaseSoftwareImage DeleteDatabaseSoftwareImage	com.oraclecloud.DatabaseService.DeleteDatabaseSoftwareImage.begin com.oraclecloud.DatabaseService.DeleteDatabaseSoftwareImage.end
List Database Software Images	ListDatabaseSoftwareImages	com.oraclecloud.DatabaseService.ListDatabaseSoftwareImages
Get Database Software Image	GetDatabaseSoftwareImage	com.oraclecloud.DatabaseService.GetDatabaseSoftwareImage
Change compartment of Database Software Image	MoveDatabaseSoftwareImage MoveDatabaseSoftwareImage	com.oraclecloud.DatabaseService.MoveDatabaseSoftwareImage.begin com.oraclecloud.DatabaseService.MoveDatabaseSoftwareImage.end
Update Database Software Image	UpdateDatabaseSoftwareImage	com.oraclecloud.DatabaseService.UpdateDatabaseSoftwareImage

Example 7-70 Database Software Image Events Examples

This is a reference event for Create Database Software Image - Begin:

```
{
  "eventType":
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
"com.oraclecloud.DatabaseService.CreateDatabaseSoftwareImage.begin",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "6dcad2c5-de0a-4e46-9a18-25c66f292dcf",
  "eventTime": "2021-06-18T04:04:35.451Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csid68e598ea4474b18860cdd476af4a/
b46a0a70da064d57a149c3c49b7cc588/FB22262C016611EFC401A9292187861D",
    "eventName": "CreateDatabaseSoftwareImage",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage_19c",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljsbpu3kxb54loym5sgkn2z4briz4xbokhou
xywkir764ht4txx4nya",
    "availabilityDomain": "AD2",
    "freeformTags": null,
    "definedTags": null,
    "identity": {
      "principalName": "xxxx",
      "principalId":
"ocid1.user.region1..aaaaaaaaedlsln4welqy3upwitxqrss3fnu6potqmds4xo3xld5sqyzwb
veq",
      "authType": "natv",
      "callerName": null,
      "callerId": null,
      "tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
      "ipAddress": "160.34.124.111",
      "credentials": ""
    },
    "request": {
      "id": "csid68e598ea4474b18860cdd476af4a/
b46a0a70da064d57a149c3c49b7cc588/FB22262C016611EFC401A9292187861D",
      "path": "/20160918/databaseSoftwareImages",
      "action": "POST",
      "parameters": {},
      "headers": {}
    },
    "response": {
      "status": "200",
      "responseTime": "2021-06-18T04:04:36.457Z",
      "headers": {},
      "payload": null,
      "message": "CreateDatabaseSoftwareImage succeeded"
    },
    "stateChange": {
      "previous": null,
      "current": {
        "displayName": "DBImage_19c",
        "lifecycleState": "PROVISIONING"
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      }
    },
    "additionalDetails": {
      "dbVersion": "19.0.0.0",
      "displayName": "DBImage_19c",
      "lifecycleState": "PROVISIONING",
      "timeCreated": "2021-06-18T04:04:35.739Z"
    }
  }
}
```


This is a reference event for Create Database Software Image – End


```
{
  "eventType":
"com.oraclecloud.DatabaseService.CreateDatabaseSoftwareImage.end",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "7bf14762-1ec7-4230-99d0-6323c5e8b3cb",
  "eventTime": "2021-06-18T04:26:24.119Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csid68e598ea4474b18860cdd476af4a/
b46a0a70da064d57a149c3c49b7cc588/FB22262C016611EFC401A9292187861D",
    "eventName": "CreateDatabaseSoftwareImage",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage_19c",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljsbpu3kxb54loym5sgkn2z4briz4xbokhou
xywkir764ht4txx4nya",
    "availabilityDomain": null,
    "freeformTags": null,
    "definedTags": null,
    "identity": {
      "principalName": null,
      "principalId": null,
      "authType": null,
      "callerName": null,
      "callerId": null,
      "tenantId": null,
      "ipAddress": null,
      "credentials": null,
      "userAgent": null,
      "consoleSessionId": null
    },
    "request": {
      "id": "d2030b31-16aa-4e71-9b06-dabc7545ad65",
      "path": null,
      "action": null,
      "parameters": null,
      "headers": null
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    },
    "response": {
      "status": null,
      "responseTime": "2021-06-18T04:26:24.119Z",
      "headers": null,
      "payload": null,
      "message": "CreateDatabaseSoftwareImage"
    },
    "stateChange": {
      "previous": null,
      "current": {
        "displayName": "DBImage_19c",
        "lifecycleState": "AVAILABLE"
      }
    },
    "additionalDetails": {
      "dbVersion": "19.0.0.0",
      "displayName": "DBImage_19c",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2021-06-18T04:04:35.739Z"
    }
  }
}
```


This is a reference event for Delete Database Software Image – Start


```
{
  "eventType":
"com.oraclecloud.DatabaseService.DeleteDatabaseSoftwareImage.begin",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "b789f327-ad93-45e1-a739-45fdb97c24d2",
  "eventTime": "2021-06-25T05:06:25.316Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csidd793a61a40689138814477fa5c6e/
ffe9b65c8dcf425c9b30ae1a30ae5687/B7CA65C1DDBEDB55052051EF1113DB73",
    "eventName": "DeleteDatabaseSoftwareImage",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage_202106152041",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljrzqt4tr326jtdmyudcziz6h5uhce36jbnd
xojgeg5kpwjcipxelga",
    "availabilityDomain": "AD2",
    "freeformTags": null,
    "definedTags": null,
    "identity": {
      "principalName": "xxx",
      "principalId":
"ocid1.user.region1..aaaaaaaaedlsln4welqy3upwitxqrss3fnu6potqmds4xo3xld5sqyzwb
veq",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "authType": "natv",
      "callerName": null,
      "callerId": null,
      "tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
      "ipAddress": "160.34.124.245",
      "credentials": ""
    },
    "request": {
      "id": "csidd793a61a40689138814477fa5c6e/
ffe9b65c8dcf425c9b30ae1a30ae5687/B7CA65C1DDBEDB55052051EF1113DB73",
      "path": "/20160918/databaseSoftwareImages/
ocid1.databasesoftwareimage.oc1.sea.abzwkljrzqt4tr326jtdmyudcziz6h5uhce36jbndx
ojgeg5kpwjcipxelga",
      "action": "DELETE",
      "parameters": {},
      "headers": {}
    },
    "response": {
      "status": "204",
      "responseTime": "2021-06-25T05:06:26.074Z",
      "headers": {},
      "payload": null,
      "message": "DeleteDatabaseSoftwareImage succeeded"
    },
    "stateChange": {
      "previous": null,
      "current": {
        "displayName": "DBImage_202106152041",
        "lifecycleState": "AVAILABLE"
      }
    },
    "additionalDetails": {
     "dbVersion": "19.0.0.0",
      "displayName": "DBImage_202106152041",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2021-06-15T15:11:53.672Z"
    }
  }
}
```


This is a reference event for Delete Database Software Image – End


```
{
 "eventType":
"com.oraclecloud.DatabaseService.DeleteDatabaseSoftwareImage.end",
 "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "6a1a808f-f5dd-4fb5-9d77-3bad38591998",
  "eventTime": "2021-06-25T05:07:21.454Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csidd793a61a40689138814477fa5c6e/
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
ffe9b65c8dcf425c9b30ae1a30ae5687/B7CA65C1DDBEDB55052051EF1113DB73",
    "eventName": "DeleteDatabaseSoftwareImage",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage_202106152041",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljrzqt4tr326jtdmyudcziz6h5uhce36jbnd
xojgeg5kpwjcipxelga",
    "availabilityDomain": null,
    "freeformTags": null,
    "definedTags": null,
    "identity": {
      "principalName": null,
      "principalId": null,
      "authType": null,
      "callerName": null,
      "callerId": null,
      "tenantId": null,
      "ipAddress": null,
      "credentials": null,
      "userAgent": null,
      "consoleSessionId": null
    },
    "request": {
      "id": "676465d5-e066-43e4-bb21-77dc468ce1f9",
      "path": null,
      "action": null,
      "parameters": null,
      "headers": null
    },
    "response": {
      "status": null,
      "responseTime": "2021-06-25T05:07:21.454Z",
      "headers": null,
      "payload": null,
      "message": "DeleteDatabaseSoftwareImage"
    },
    "stateChange": {
      "previous": null,
      "current": {
        "displayName": "DBImage_202106152041",
        "lifecycleState": "TERMINATED"
      }
    },
    "additionalDetails": {
      "dbVersion": "19.0.0.0",
      "displayName": "DBImage_202106152041",
      "lifecycleState": "TERMINATED",
      "timeCreated": "2021-06-15T15:11:53.672Z"
    }
  }
}
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

This is a reference event for List Database Software Images


```
{
  "eventType": "com.oraclecloud.DatabaseService.ListDatabaseSoftwareImages",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "762072e3-52a0-437c-9956-334aad767db8",
  "eventTime": "2021-06-24T05:39:41.344Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": null,
    "eventName": "ListDatabaseSoftwareImages",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "",
    "resourceId": null,
    "availabilityDomain": "AD2",
    "freeformTags": null,
    "definedTags": null,
    "identity": {
      "principalName": "xxx",
      "principalId":
"ocid1.user.region1..aaaaaaaaedlsln4welqy3upwitxqrss3fnu6potqmds4xo3xld5sqyzwb
veq",
      "authType": "natv",
      "callerName": null,
      "callerId": null,
      "tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
      "ipAddress": "160.34.125.213",
      "credentials": "",
      "userAgent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/91.0.4472.77 Safari/537.36",
      "consoleSessionId": "csidc25d57894f49986ae07157ab5e03"
    },
    "request": {
      "id": "csidc25d57894f49986ae07157ab5e03/
dc74ad5a75d7494bb35408a33a99c0ad/084F718D53FBC66D46B573CF8D5D1A5A",
      "path": "/20160918/databaseSoftwareImages",
      "action": "GET",
      "parameters": {},
      "headers": {}
    },
    "response": {
      "status": "200",
      "responseTime": "2021-06-24T05:39:41.998Z",
      "headers": {},
      "payload": null,
      "message": "ListDatabaseSoftwareImages succeeded"
    },
    "stateChange": {
      "previous": null,
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "current": null
    },
    "additionalDetails": null
  }
}
```


This is a reference event for Get Database Software Image


```
{
  "eventType": "com.oraclecloud.DatabaseService.GetDatabaseSoftwareImage",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "179471f3-c628-4e55-98d0-959b0913b327",
  "eventTime": "2021-06-24T04:11:10.738Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csid31b4236f4f1ebc4ee66169d17b0a/
1f2e5dbf57ca40c886b111616b02c33f/CA5B5AD4C8678E2C428DA9AF0C62A082",
    "eventName": "GetDatabaseSoftwareImage",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage_19c",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljsbpu3kxb54loym5sgkn2z4briz4xbokhou
xywkir764ht4txx4nya",
    "availabilityDomain": "AD2",
    "freeformTags": null,
    "definedTags": null,
    "identity": {
      "principalName": "xxx",
      "principalId":
"ocid1.user.region1..aaaaaaaalkib75hbawg6xh6rlys5ztumfx3cxf2ktwbon5ryv6xjev7oe
ctq",
      "authType": "natv",
      "callerName": null,
      "callerId": null,
      "tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
      "ipAddress": "192.188.170.109",
      "credentials": "",
      "userAgent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_15_7)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/91.0.4472.106 Safari/537.36",
      "consoleSessionId": "csid31b4236f4f1ebc4ee66169d17b0a"
    },
    "request": {
      "id": "csid31b4236f4f1ebc4ee66169d17b0a/
1f2e5dbf57ca40c886b111616b02c33f/CA5B5AD4C8678E2C428DA9AF0C62A082",
      "path": "/20160918/databaseSoftwareImages/
ocid1.databasesoftwareimage.oc1.sea.abzwkljsbpu3kxb54loym5sgkn2z4briz4xbokhoux
ywkir764ht4txx4nya",
      "action": "GET",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "parameters": {},
      "headers": {}
    },
    "response": {
      "status": "200",
      "responseTime": "2021-06-24T04:11:10.966Z",
      "headers": {},
      "payload": null,
      "message": "GetDatabaseSoftwareImage succeeded"
    },
    "stateChange": {
      "previous": null,
      "current": {
        "displayName": "DBImage_19c",
        "lifecycleState": "AVAILABLE"
      }
    },
    "additionalDetails": {
      "dbVersion": "19.0.0.0",
      "displayName": "DBImage_19c",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2020-09-24T18:06:00.395Z"
    }
  }
}
```


Database Upgrade Event Types

Review the list of event types that Database Upgrade emit.

Table 7-17 Database Upgrade Events

Friendly Name	Event Type
MoveDatabase	com.oraclecloud.DatabaseService.MoveDatabase.begin
	com.oraclecloud.DatabaseService.MoveDatabase.end
UpgradeDatabase	com.oraclecloud.DatabaseService.UpgradeDatabase.begin
	com.oraclecloud.DatabaseService.UpgradeDatabase.end

Example 7-71 Database Upgrade Examples

This is a reference event for Move Database Begin:


```
eventType": "com.oraclecloud.DatabaseService.MoveDatabase.begin"
"eventName": "MoveDatabase"
{
  "eventType": "com.oraclecloud.DatabaseService.MoveDatabase.begin",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "c78fb931-adda-4813-ab75-b9cc51b96847",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "eventTime": "2021-07-15T11:03:52.582Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csidff49902144a1a3c98bd4e5f49c1c/
37d3a7113f1e462fbd5a741dc5d52ef7/57472323A2FFABC31623FB7660F519AB",
    "eventName": "MoveDatabase",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DB211507",
    "resourceId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3p
yiqhpq",
    "availabilityDomain": "",
    "freeformTags": {},
    "definedTags": {
      "Operators": {
        "Created_By": "testuser@example.com",
        "Created_Time": "2021-07-15T11:01:24.085Z"
      }
    },
    "identity": {
      "principalName": null,
      "principalId": null,
      "authType": null,
      "callerName": null,
      "callerId": null,
      "tenantId": null,
      "ipAddress": null,
      "credentials": null,
      "userAgent": null,
      "consoleSessionId": null
    },
    "request": {
      "id": "8f7b5283-4896-4ce0-aefa-a4156971e07b",
      "path": null,
      "action": null,
      "parameters": null,
      "headers": null
    },
    "response": {
      "status": null,
      "responseTime": null,
      "headers": null,
      "payload": null,
      "message": "MoveDatabase"
    },
    "stateChange": {
      "previous": null,
      "current": {
        "definedTags": {
          "Operators": {
            "Created_By": "testuser@example.com",
            "Created_Time": "2021-07-15T11:01:24.085Z"
          }
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        },
        "displayName": "DB211507",
        "freeTags": {},
        "lifecycleState": "UPGRADING"
      }
    },
    "additionalDetails": {
      "databaseEdition": "ENTERPRISE_EDITION_EXTREME",
      "databaseId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3p
yiqhpq",
      "dbHomeId":
"ocid1.dbhome.oc1.sea.abzwkljskjhqcc4bkzuai6kybx3hp3jznvxcgwlghqjnklpchrktcynw
nloq",
      "dbUniqueName": "Audit_check",
      "dbVersion": "18.14.0.0",
      "exadataInfrastructureId":
"ocid1.exadatainfrastructure.oc1.sea.abzwkljsqtwpumvorlttjuxcx6qtkxhuasvmyu4qy
3bng74sqolpwjunsrqa",
      "lifecycleState": "UPGRADING",
      "timeCreated": "2021-07-15T11:01:24Z",
      "timeUpdated": "2021-07-15T11:02:34Z",
      "vmClusterId":
"ocid1.vmcluster.oc1.sea.abzwkljsghfg7e3q3xh3gzqqxijt72fepczowmk35t554wxbpew7t
xnooprq"
    }
  }
}
```


This is a reference event for Move Database End:


```
eventType": "com.oraclecloud.DatabaseService.MoveDatabase.end"
"eventName": "MoveDatabase"

{
  "eventType": "com.oraclecloud.DatabaseService.MoveDatabase.end",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "97429610-9b8f-4da7-8f99-2f7cd5455c24",
  "eventTime": "2021-07-15T11:03:56.359Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csidff49902144a1a3c98bd4e5f49c1c/
37d3a7113f1e462fbd5a741dc5d52ef7/57472323A2FFABC31623FB7660F519AB",
    "eventName": "MoveDatabase",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DB211507",
    "resourceId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3p
yiqhpq",
    "availabilityDomain": "",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    "freeformTags": {},
    "definedTags": {
      "Operators": {
        "Created_By": "testuser@example.com",
        "Created_Time": "2021-07-15T11:01:24.085Z"
      }
    },
    "identity": {
      "principalName": null,
      "principalId": null,
      "authType": null,
      "callerName": null,
      "callerId": null,
      "tenantId": null,
      "ipAddress": null,
      "credentials": null,
      "userAgent": null,
      "consoleSessionId": null
    },
    "request": {
      "id": "8f7b5283-4896-4ce0-aefa-a4156971e07b",
      "path": null,
      "action": null,
      "parameters": null,
      "headers": null
    },
    "response": {
      "status": null,
      "responseTime": null,
      "headers": null,
      "payload": null,
      "message": "MoveDatabase"
    },
    "stateChange": {
      "previous": null,
      "current": {
        "definedTags": {
          "Operators": {
            "Created_By": "testuser@example.com",
            "Created_Time": "2021-07-15T11:01:24.085Z"
          }
        },
        "displayName": "DB211507",
        "freeTags": {},
        "lifecycleState": "AVAILABLE"
      }
    },
    "additionalDetails": {
      "databaseEdition": "ENTERPRISE_EDITION_EXTREME",
      "databaseId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3p
yiqhpq",
      "dbHomeId":
"ocid1.dbhome.oc1.sea.abzwkljs7obwmwq2kgkjgeko4voudfcd5chu6iuonnew7uuxculsvdra
lzda",
      "dbUniqueName": "Audit_check",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "dbVersion": "19.10.0.0",
      "exadataInfrastructureId":
"ocid1.exadatainfrastructure.oc1.sea.abzwkljsqtwpumvorlttjuxcx6qtkxhuasvmyu4qy
3bng74sqolpwjunsrqa",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2021-07-15T11:01:24Z",
      "timeUpdated": "2021-07-15T11:03:56Z",
      "vmClusterId":
"ocid1.vmcluster.oc1.sea.abzwkljsghfg7e3q3xh3gzqqxijt72fepczowmk35t554wxbpew7t
xnooprq"
    }
  }
}
```


This is a reference event for Upgrade Database Begin:

```
eventType": " com.oraclecloud.DatabaseService.UpgradeDatabase.begin"
"eventName": " UpgradeDatabase"
```


```
{
  "eventType": "com.oraclecloud.DatabaseService.UpgradeDatabase.begin",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "96c66a03-fa53-4008-b86c-619b9f44ef31",
  "eventTime": "2021-07-15T11:02:33.948Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csidff49902144a1a3c98bd4e5f49c1c/
37d3a7113f1e462fbd5a741dc5d52ef7/57472323A2FFABC31623FB7660F519AB",
    "eventName": "UpgradeDatabase",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DB211507",
    "resourceId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3p
yiqhpq",
    "availabilityDomain": "",
    "freeformTags": {},
    "definedTags": {
      "Operators": {
        "Created_By": "testuser@example.com",
        "Created_Time": "2021-07-15T11:01:24.085Z"
      }
    },
    "identity": {
      "principalName": "testuser@example.com",
      "principalId":
"ocid1.user.region1..aaaaaaaac5dskxmhmemrbyhaeabru4kgpituqmvgxjcm5y22mlxre26om
4zq",
      "authType": "natv",
      "callerName": null,
      "callerId": null,
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
      "ipAddress": "160.34.92.141",
      "credentials":
"ST$eyJraWQiOiJhc3dfcmVnaW9uMV8yMDIxLTA1LTE0IiwiYWxnIjoiUlMyNTYifQ.eyJzdWIiOiJ
vY2lkMS51c2VyLnJlZ2lvbjEuLmFhYWFhYWFhYzVkc2t4bWhtZW1yYnloYWVhYnJ1NGtncGl0dXFtd
md4amNtNXkyMm1seHJlMjZvbTR6cSIsIm1mYV92ZXJpZmllZCI6ImZhbHNlIiwiaXNzIjoiYXV0aFN
lcnZpY2Uub3JhY2xlLmNvbSIsInB0eXBlIjoidXNlciIsInNlc3NfZXhwIjoiRnJpLCAxNiBKdWwgM
jAyMSAxMDo1OToxNyBVVEMiLCJhdWQiOiJvY2kiLCJwc3R5cGUiOiJuYXR2IiwidHR5cGUiOiJsb2d
pbiIsImV4cCI6MTYyNjM1MDM1OSwiaWF0IjoxNjI2MzQ2NzU5LCJqdGkiOiJhNWE2YmRhNC02OGJkL
TQ0ZjktODFlNS0yZGRiZGMyZWU4MTQiLCJ0ZW5hbnQiOiJvY2lkMS50ZW5hbmN5LnJlZ2lvbjEuLmF
hYWFhYWFhZTM0Z3diZXozZGJvN3VydGNmNDYyd2ozbW5tYW5iYmE1eHFpbmN6Z20yejVkNjRjbmw1c
SIsImp3ayI6IntcImtpZFwiOlwicHVia2V5LWU5Y2RmMWY4LTJjNTEtNGEyNi1hZjkyLTNlZjlkNTl
hNWEzN1wiLFwiblwiOlwicjE1azFxR0s1MXJaVUZzLVdFNDFrU2ZLRWozanlndlZ2bG9pZVM5Y3RLO
DJ4QmxZT1F3SkZwX19HUXVGTFktWk45MVZabXdSQkpFeHJYbE1hNHV6aE5VS1BNQ3M5OG9XTENpQmo
2UlVMekFLZHVtbHZlbHEwTUcwVmtMMjVXb1hONzU2N19wUUx3NDc0T2w2VzNJUl9Jd2tBVy1QZnF5Q
zdpY1NOekI2UXh2aGlVMXRpTUFUZ09RUjQ1Ty1md3NnVnl5WW14eVNCbm1sN1BDUEtoaGpBLTlFOFo
zeV94cnlPWFBuYjl1QlFrczZOblgzX08yRkZ1UGxualN5QXdlSEZLSzd6MUJNVVdONUMtLV9Uemw5T
0VaOENMeGdKNkpzeUUtRUVMeEZEUkZ0MlVHZGNubUVBY2hNQ1gxNEtnWEFELVk5a2dkVS1HczNaUU1
VTW5pYmFRXCIsXCJlXCI6XCJBUUFCXCIsXCJrdHlcIjpcIlJTQVwiLFwiYWxnXCI6XCJSUzI1NlwiL
FwidXNlXCI6bnVsbH0ifQ.GGPB9TMU_CkqOfLNCgO2J6-
QwvHEEaY6jky5T92PBoPOBUkXb_ruNOe6C8RwQhs2pmIsIBybTP7Ueene4cJ9nNKzsoTzNAUrSp971
4Be2noQgqJL-ZtFIApThz1vXdfFWZtQQInxx-jALpsvkhvD_EVdbfx6FvlwDfm-
ReNDSnrUoFkfx8NpIrkqP8KUFhnX-
TRy0nUvqlaiB9C1L8FZGOJbBn0uiyloZe4DNP3Z39jViy37UntHeuMPnIqOYEsmPrSMTiJUBAypy4_
I5irWuR0PpxfrttP6nlEi80V5YrmO3l2-ZJBRDKz8TdMkxclw3IT0qulTzFPqB3cXg45PGQ",
      "userAgent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10.15; rv:78.0)
Gecko/20100101 Firefox/78.0",
      "consoleSessionId": "csidff49902144a1a3c98bd4e5f49c1c"
    },
    "request": {
      "id": "csidff49902144a1a3c98bd4e5f49c1c/
37d3a7113f1e462fbd5a741dc5d52ef7/57472323A2FFABC31623FB7660F519AB",
      "path": "/20160918/databases/
ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3py
iqhpq/actions/upgrade",
      "action": "POST",
      "parameters": {},
      "headers": {}
    },
    "response": {
      "status": "200",
      "responseTime": "2021-07-15T11:02:36.072Z",
      "headers": {},
      "payload": null,
      "message": "UpgradeDatabase succeeded"
    },
    "stateChange": {
      "previous": null,
      "current": {
        "definedTags": {
          "Operators": {
            "Created_By": "testuser@example.com",
            "Created_Time": "2021-07-15T11:01:24.085Z"
          }
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        },
        "displayName": "DB211507",
        "freeTags": {},
        "lifecycleState": "UPGRADING"
      }
    },
    "additionalDetails": {
      "action": "UPGRADE",
      "currentDbVersion": "18.14.0.0",
      "lifecycleState": "IN_PROGRESS",
      "source": "DATABASE_HOME",
      "targetDbVersion": "19.10.0.0",
      "timeStarted": "2021-07-15T11:02:34Z",
      "upgradeHistoryEntryId":
"ocid1.dbupgradehistory.oc1.sea.abzwkljstw7dymj6mfnaotkxpuie2bdxmlizqr5ie46dew
cuzdh35md7fuhq"
    }
  }
}
```


This is a reference event for Upgrade Database End:


```
eventType": " com.oraclecloud.DatabaseService.UpgradeDatabase.end"
"eventName": " UpgradeDatabase"

{
  "eventType": "com.oraclecloud.DatabaseService.UpgradeDatabase.end",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "DatabaseService",
  "eventId": "e6c8a9e4-53b0-4bb8-be65-a39d157f04b2",
  "eventTime": "2021-07-15T11:03:56.028Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csidff49902144a1a3c98bd4e5f49c1c/
37d3a7113f1e462fbd5a741dc5d52ef7/57472323A2FFABC31623FB7660F519AB",
    "eventName": "UpgradeDatabase",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d6
4cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DB211507",
    "resourceId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3p
yiqhpq",
    "availabilityDomain": "",
    "freeformTags": {},
    "definedTags": {
      "Operators": {
        "Created_By": "testuser@example.com",
        "Created_Time": "2021-07-15T11:01:24.085Z"
      }
    },
    "identity": {
      "principalName": null,
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "principalId": null,
        "authType": null,
        "callerName": null,
        "callerId": null,
        "tenantId": null,
        "ipAddress": null,
        "credentials": null,
        "userAgent": null,
        "consoleSessionId": null
      },
      "request": {
        "id": "8f7b5283-4896-4ce0-aefa-a4156971e07b",
        "path": null,
        "action": null,
        "parameters": null,
        "headers": null
      },
      "response": {
        "status": null,
        "responseTime": null,
        "headers": null,
        "payload": null,
        "message": "UpgradeDatabase"
      },
      "stateChange": {
        "previous": null,
        "current": {
          "definedTags": {
            "Operators": {
              "Created_By": "testuser@example.com",
              "Created_Time": "2021-07-15T11:01:24.085Z"
            }
          },
          "displayName": "DB211507",
          "freeTags": {},
          "lifecycleState": "AVAILABLE"
        }
      },
      "additionalDetails": {
        "databaseEdition": "ENTERPRISE_EDITION_EXTREME",
        "databaseId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3p
yiqhpq",
        "dbHomeId":
"ocid1.dbhome.oc1.sea.abzwkljs7obwmwq2kgkjgeko4voudfcd5chu6iuonnew7uuxculsvdra
lzda",
        "dbUniqueName": "Audit_check",
        "dbVersion": "19.10.0.0",
        "exadataInfrastructureId":
"ocid1.exadatainfrastructure.oc1.sea.abzwkljsqtwpumvorlttjuxcx6qtkxhuasvmyu4qy
3bng74sqolpwjunsrqa",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2021-07-15T11:01:24Z",
        "timeUpdated": "2021-07-15T11:03:56Z",
        "vmClusterId":
"ocid1.vmcluster.oc1.sea.abzwkljsghfg7e3q3xh3gzqqxijt72fepczowmk35t554wxbpew7t
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
xnooprq"
    }
  }
}
```

Pluggable Database Event Types

Review the list of event types that Pluggable Databases emit.

Table 7-18 Pluggable Database Events

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 - Start	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 - Local Clone Begin	com.oraclecloud.databaseservice.pluggabledatabase.localclone.begin
Pluggable Database - Local Clone End	com.oraclecloud.databaseservice.pluggabledatabase.localclone.end

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-18 (Cont.) Pluggable Database Events

Friendly Name	Event Type
Pluggable Database - Refresh Begin	com.oraclecloud.databaseservice.pluggabledatabase.refresh.begin
Pluggable Database - Refresh End	com.oraclecloud.databaseservice.pluggabledatabase.refresh.end
Pluggable Database - Relocate Begin	com.oraclecloud.databaseservice.pluggabledatabase.relocate.begin
Pluggable Database - Relocate End	com.oraclecloud.databaseservice.pluggabledatabase.relocate.end
Pluggable Database - Remote Clone Begin	com.oraclecloud.databaseservice.pluggabledatabase.remoteclone.begin
Pluggable Database - Remote Clone End	com.oraclecloud.databaseservice.pluggabledatabase.remoteclone.end

Example 7-72 Pluggable Database Examples

This is a reference event for Pluggable Database - Create Begin:

```
{
  “eventID”: “af1cdf4e-4001-11eb-9fb6-f45c89b1cb17",
  “eventTime”: “2020-12-17T00:49:14.123Z”,
  “extensions”: {
    “compartmentId”: “ocid1.compartment.oc1..unique_id”
  },
  “eventType”:
“com.oraclecloud.databaseservice.createexternalpluggabledatabase”,
  “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.externalpluggabledatabases.oc1.phx.unique_id”,
    “availabilityDomain”: “XXIT:PHX-AD-1”,
    “freeFormTags”: {},
    “definedTags”: {},
    “additionalDetails”: {
      “id”: “ocid1.externalpluggabledatabases.oc1.phx.unique_id”,
      “timeCreated”: “2020-11-13T21:15:59.000Z”,
      “timeUpdated”: “2020-11-13T21:15:59.000Z”,
      “externalCDBId”: “ocid1.externalcontainerdatabase.oc1.....unique_id”,
      “lifecycleState”: “AVAILABLE”,
      “lifecycleDetails”: “External Pluggable Database is available”,
      “dbUniqueName”: “CDB122_phx16q”,
      “dbId”: “3455094890",
      “dbVersion”: “12.2.0.1.0",
      “dbEdition”: “ENTERPRISE_EDITION_EXTREME_PERFORMANCE”,
      “timeZone”: “US/Pacific”,
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      “databaseManagementServiceStatus”: “ENABLED”,
      “databaseManagementServiceConnectorId”:
“ocid1.externaldatabaseconnector.oc1.....unique_id”,
      “displayName”: “External Pluggable Database - Create”
    }
  }
}
```

This is a reference event for Pluggable Database - Convert to Regular Begin:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "isRefreshableClone": true,
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "UPDATING",
        "displayName": "Pluggable Database - Convert to Regular Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Convert to Regular End:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
"com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "isRefreshableClone": false,
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "displayName": "Pluggable Database - Convert to Regular End"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```


This is a reference event for Pluggable Database - Inplace Restore Begin:


```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "RESTORE_IN_PROGRESS",
        "isRefreshableClone": false,
        "displayName": "Pluggable Database - Inplace Restore Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Inplace Restore End:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "isRefreshableClone": false,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Inplace Restore End"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Local Clone Begin:

```
"exampleEvent": {
    "eventID": "unique_id",
```


Copyright © 2008, 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.localclone.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 - Local Clone Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```


This is a reference event for Pluggable Database - Local Clone End:


```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.localclone.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",
```


Copyright © 2008, 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",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "isRefreshableClone": false,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Local Clone End"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```


This is a reference event for Pluggable Database - Refresh Begin:


```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.refresh.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "isRefreshableClone": true,
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Refresh Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Refresh End:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.refresh.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "isRefreshableClone": true,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Refresh End"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Relocate Begin:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.relocate.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
```


Copyright © 2008, 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",
        "isRefreshableClone": false,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Relocate Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```


This is a reference event for Pluggable Database - Relocate End:


```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.relocate.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.......unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Relocate End"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Remote Clone Begin:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.remoteclone.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 - Remote Clone Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Remote Clone End:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_id"
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.remoteclone.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 - Remote Clone End"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```


VM Node Subsetting Event Types

Review the list of event types that VM Node Subsetting emits.

Table 7-19 VM Node Subsetting Events

Friendly Name	Event Type
VM Cluster - Add Virtual Machine Begin	com.oraclecloud.databaseservice.vmclustereraddvirtualmachine.begin
VM Cluster - Add Virtual Machine End	com.oraclecloud.databaseservice.vmclustereraddvirtualmachine.end
VM Cluster - Terminate Virtual Machine Begin	com.oraclecloud.databaseservice.vmclustererterminatevirtualmachine.begin
VM Cluster - Terminate Virtual Machine End	com.oraclecloud.databaseservice.vmclustererterminatevirtualmachine.end

Copyright © 2008, 2026, Oracle and/or its affiliates.

Example 7-73 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",
      "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",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "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 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",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    "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 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",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "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"
    }
  }
}
```

Database Service Events

The Database Service emits events, which are structured messages that indicate changes in resources.

- Overview of Database Service Events
- 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.
- Remediation
  These topics cover some common issues you might run into and how to address them.

Overview of Database Service Events

Database Service Events feature implementation enables you to get notified about health issues with your Oracle Databases or 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. Customers are not notified of this situation.

Database Service Events feature implementation generates events for Data Plane 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 ExaDB-C@C 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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Manage Oracle Trace File Analyzer

- To check the run status of Oracle Trace File Analyzer, run the `tfactl status` command as `root` or a non-root user:

```
# tfactl status
.-----------------------------------------------------------------------------
----------------------.
| Host    | Status of TFA | PID    | Port | Version    | Build ID
| Inventory Status|
+---------------+--------------+--------+------+------------
+---------------------+------------+
| node1      | RUNNING    | 41312  | 5000 | 22.1.0.0.0 |
22100020220310214615 | COMPLETE        |
| node2      | RUNNING    | 272300 | 5000 | 22.1.0.0.0 |
22100020220310214615 | COMPLETE        |
'---------------+--------------+--------+------+------------
+---------------------+------------'
```

- To start the Oracle Trace File Analyzer daemon on the local node, run the `tfactl start` command as `root`:

```
# tfactl start
Starting TFA..
Waiting up to 100 seconds for TFA to be started..
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
Successfully started TFA Process..
. . . . .
TFA Started and listening for commands
```

- To stop the Oracle Trace File Analyzer daemon on the local node, run the `tfactl stop` command as `root`:

```
# tfactl stop
Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..
```

Manage sysLens

- If sysLens is running, then once every 15 minutes data is collected in the local Guest VM to discover the events to be reported.

Copyright © 2008, 2026, Oracle and/or its affiliates.

To check if sysLens is running, run the `systemctl status syslens` command as `root` in the Guest VM:

```
# 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 node1 su[336679]: pam_unix(su-l:session): session closed
for user oracle
Mar 19 20:53:13 node1 su[336728]: (to oracle) root on none
Mar 19 20:53:13 node1 su[336728]: pam_unix(su-l:session): session opened
for user oracle by (uid=0)
Mar 19 20:53:13 node1 su[336728]: pam_unix(su-l:session): session closed
for user oracle
Mar 19 20:53:14 node1 su[336808]: (to oracle) root on none
Mar 19 20:53:14 node1 su[336808]: pam_unix(su-l:session): session opened
for user oracle by (uid=0)
Mar 19 20:53:14 node1 su[336808]: pam_unix(su-l:session): session closed
for user oracle
Mar 19 20:53:15 node1 su[336869]: (to oracle) root on none
Mar 19 20:53:15 node1 su[336869]: pam_unix(su-l:session): session opened
for user oracle by (uid=0)
Mar 19 20:53:15 node1 su[336869]: pam_unix(su-l:session): session closed
for user oracle
```

- If sysLens is enabled, when there is a reboot of the Guest VM, then sysLens starts automatically.
  To validate if sysLens is enabled to collect telemetry, run the `systemctl is-enabled syslens` command as `root` in the Guest VM:

  ```
  # 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 `/var/opt/oracle/syslens/bin/applin_ctl config_file get config.telemetry.enable_telemetry --json /etc/oracle/syslens/config/host.config` command as `root` in the Guest VM:

  ```
  # /var/opt/oracle/syslens/bin/applin_ctl config_file get
  config.telemetry.enable_telemetry --json /etc/oracle/syslens/config/
  ```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
host.config
on
```

Manage Database Service Agent

View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

- To check the status of the Database Service Agent, run the `systemctl status` command:

```
# systemctl status dbcsagent.service
dbcsagent.service
Loaded: loaded (/usr/lib/systemd/system/dbcsagent.service; enabled; vendor
preset: disabled)
Active: active (running) since Fri 2022-04-01 13:40:19 UTC; 6min ago
Process: 9603 ExecStopPost=/bin/bash -c kill `ps -fu opc |grep "java.*dbcs-
agent.*jar" |awk '{print $2}' ` (code=exited, status=0/SUCCESS)
Main PID: 10055 (sudo)
CGroup: /system.slice/dbcsagent.service
 10055 sudo -u opc /bin/bash -c umask 077; /bin/java -
Doracle.security.jps.config=/opt/oracle/...
```

- To start the agent if it is not running, run the `systemctl start` command as the `root` user:

```
systemctl start dbcsagent.service
```

Related Topics

- Using the Console to Create an ASM VM Cluster
  To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection
  You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- Overview of Events
- Notifications Overview

Receive Notifications about Database Service Events

Subscribe to the Database Service Events and get notified.

To receive notifications, subscribe to Database Service Events and get notified using the Oracle Notification service, see *Notifications Overview*. For more information about Oracle Cloud Infrastructure Events, see *Overview of Events*.

Events Service - Event Types:

- Database - Critical
- DB Node - Critical
- DB Node - Error
- DB Node - Warning
- DB Node - Info

Copyright © 2008, 2026, Oracle and/or its affiliates.

- DB System - Critical

Related Topics

- Overview of Events
- Notifications Overview

Database Service Event Types

Review the list of event types that the Database Service emits.

> **Note**
>
> - Critical events are triggered due to several types of critical conditions and errors that cause disruption to the database and other critical components. For example, database hang errors, and availability errors for databases, database nodes, and database systems to let you know if a resource becomes unavailable.
> - Information events are triggered when the database and other critical components work as expected. For example, a clean shutdown of CRS, CDB, client, or scan listener, or a startup of these components will create an event with the severity of INFO.
> - Threshold limits reduce the number of notifications customers will receive for similar incident events whilst at the same time ensuring they receive the incident events and are reminded in a timely fashion.

Table 7-20 Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
Resource Utilization - Disk Usage	`HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE` This event is reported when VM guest file system free space falls below 10% free, as determined by the operating system `df(1)` command, for the following file systems: • `/` • `/u01` • `/u02` • `/var` (X8M and later only) • `/tmp` (X8M and later only)	HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE	`com.oraclecloud.databaseservice.dbnode.critical`	Critical threshold: 90%

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-20 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
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`	HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION	`com.oraclecloud.databaseservice.dbnode.critical`	N/A
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.	CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE	`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.	CONFIGURATION.DB_GUEST.SSHD.INVALID	`com.oraclecloud.databaseservice.dbnode.critical`	N/A

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-20 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
Oracle EXAchk Reported Issues	`HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT` Oracle EXAchk is Exadata database platform's holistic health check that includes software, infrastructure and database configuration checks. CRITICAL check alerts should be addressed in 24 hours to maintain the maximum stability and availability of your system. This database service event alerts every 24 hours whenever there are any CRITICAL checks that are flagged in the most recent Oracle EXAchk report. The event will point to the latest Oracle EXAchk zip report.	HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT	`com.oraclecloud.databaseservice.dbnode.critical`	N/A
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.	AVAILABILITY.DB_GUEST.CRS_INSTANCE.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.	N/A	`com.oraclecloud.databaseservice.dbnode.information` (if .DOWN_CLEARED)	

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-20 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	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].	AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN	`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.	N/A	`com.oraclecloud.databaseservice.dbnode.information` (if .DOWN_CLEARED)	

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-20 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	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.	AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN	`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.	N/A	`com.oraclecloud.databaseservice.database.information` (if .DOWN_CLEARED)	

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-20 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	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.	AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN	`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.	N/A	`com.oraclecloud.databaseservice.database.information` (if .DOWN_CLEARED)	

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-20 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	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.	HEALTH.DB_CLUSTER.CDB.CORRUPTION	`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.	HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG	`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.	HEALTH.DB_CLUSTER.CDB.DATABASE_HANG		

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-20 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
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.	HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE	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%.	HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE	com.oraclecloud.databaseservice.dbsystem.critical com.oraclecloud.databaseservice.dbsystem.information (if < 90%)	Notifications are sent when the usage hits 70%, 80%, 90%, and 100% with a corresponding severity of 4, 3, 2, and 1.
DB I/O Latency	HEALTH.DB_GUEST.PDB.HIGH_IO_LATENCY	HEALTH.DB_GUEST.PDB.HIGH_IO_LATENCY	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >16ms
DB IO Latency Distribution	HEALTH.DB_GUEST.CDB.HIGH_IO_LATENCY_HISTOGRAM	HEALTH.DB_GUEST.CDB.HIGH_IO_LATENCY_HISTOGRAM	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >32ms
Unusual CPU Waits	HEALTH.DB_GUEST.PDB.HIGH_CPU_WAITS	HEALTH.DB_GUEST.PDB.HIGH_CPU_WAITS	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >5 minutes
DB Waits	HEALTH.DB_GUEST.PDB.EXCESSIVE_WAITS	HEALTH.DB_GUEST.PDB.EXCESSIVE_WAITS	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >20% Threshold: > 35% "Exadata User I/O Wait-event"
CPU Throttling	HEALTH.DB_GUEST.PDB.EXCESSIVE_CPU_THROTTLING	HEALTH.DB_GUEST.PDB.EXCESSIVE_CPU_THROTTLING	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >0.25 * CPU_COUNT

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-20 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
FAST RECOVERY AREA (FRA) usage	`HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_1` When the VM cluster's RECO disk group or the database's FAST RECOVERY AREA (FRA) usage exceeds 90%, Oracle preserves archive logs locally to simplify the fully cloud-automated conversion of a Snapshot Standby to a Physical Standby. If the RECO disk group or FRA usage exceeds 95%, an additional alert is triggered, and archived logs that have already been backed up are purged. At this stage, converting to a Physical Standby will require manual intervention. Reverting to a Physical Standby before reaching this threshold is recommended to avoid manual intervention.	HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_1	`com.oraclecloud.databaseservice.database.warning`	Threshold: VM cluster RECO disk group > 90%. Threshold: Database FAST RECOVERY AREA (FRA) > 90%.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-20 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
FAST RECOVERY AREA (FRA) space usage	HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_2 VM Cluster's RECO disk group or database's RECOVERY AREA space exceeds 95% used, archives that have been backed up will be purged to reduce space utilization. To convert the Snapshot Standby to Physical Standby will require following these manual steps described in *How to Roll Forward a Standby Database Using Recover Database From Service (Doc ID 2850185.1)*.	HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_2	com.oraclecloud.databaseservice.database.critical	Threshold: VM cluster RECO disk group > 95%. Threshold: Database FAST RECOVERY AREA (FRA) > 95%.
DB Blocker Detector	HEALTH.DB_GUEST.PDB.DB_BLOCKER_STALL_CHECK	HEALTH.DB_GUEST.PDB.DB_BLOCKER_STALL_CHECK	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >= 10 number of Session waiting Threshold: >= 600 sec wait time for the chain
Exadata Cache Efficiency	HEALTH.DB_GUEST.PDB.TOTAL_EXADATA_CACHE_HIT_RATIO	HEALTH.DB_GUEST.PDB.TOTAL_EXADATA_CACHE_HIT_RATIO	com.oraclecloud.databaseservice.dbnode.warning	Thresold: < 90%
Redo Log Wait	HEALTH.DB_GUEST.CDB.HIGH_REDO_LOG_WAITS	HEALTH.DB_GUEST.CDB.HIGH_REDO_LOG_WAITS	com.oraclecloud.databaseservice.dbnode.warning	Thresold: > 50% of Overall Background Processes Time

Example 7-74 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" : "2022-03-07T23:17:47Z",
 "contentType" : "application/json",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
 "data" : {
   "compartmentId" : "compartment_ID",
   "compartmentName" : "compartment_Name",
   "resourceName" : "resource_Name",
   "resourceId" : "resource_ID",
   "additionalDetails" : {
     "serviceType" : "EXACC",
     "hostName" : "host_Name",
     "description" : "EXACHK is reporting 6 checks in critical status.
Results in: /u02/oracle.ahf/data/host_Name/exachk/user_root/output/
exachk_host_Name_v1c2_030722_150746.zip",
     "eventName" : "HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT",
     "status" : "online"
   }
 },
 "eventID" : "cce55ca2-9e6c-11ec-90e6-00163e9b4de8",
 "extensions" : {
   "compartmentId" : "compartment_ID"
 }
}


{
 "eventType" : "com.oraclecloud.databaseservice.dbnode.critical",
 "cloudEventsVersion" : "0.1",
 "eventTypeVersion" : "2.0",
 "source" : "SYSLENS/host_Name/DomU",
 "eventTime" : "2022-03-06T18:14:57Z",
 "contentType" : "application/json",
 "data" : {
   "compartmentId" : "compartment_ID",
   "compartmentName" : "compartment_Name",
   "resourceName" : "resource_Name",
   "resourceId" : "resource_ID",
   "additionalDetails" : {
     "serviceType" : "EXACC",
     "hostName" : "host_Name",
     "description" : "Parameter ClientAliveCountMax has incorrect value.",
     "eventName" : "CONFIGURATION.DB_GUEST.SSHD.INVALID",
     "status" : "online"
   }
 },
 "eventID" : "5453554e-9d79-11ec-8096-00163eb980bb",
 "extensions" : {
   "compartmentId" : "compartment_ID"
 }
}


{
 "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",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
 "data" : {
   "compartmentId" : "compartment_ID",
   "compartmentName" : "compartment_Name",
   "resourceName" : "resource_Name",
   "resourceId" : "resource_ID",
   "additionalDetails" : {
     "serviceType" : "EXACC",
     "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"
 }
}


{
 "eventType" : "com.oraclecloud.databaseservice.dbnode.critical",
 "cloudEventsVersion" : "0.1",
 "eventTypeVersion" : "2.0",
 "source" : "SYSLENS/host_Name/DomU",
 "eventTime" : "2022-03-04T18:49:25Z",
 "contentType" : "application/json",
 "data" : {
   "compartmentId" : "compartment_ID",
   "compartmentName" : "compartment_Name",
   "resourceName" : "resource_Name",
   "resourceId" : "resource_ID",
   "additionalDetails" : {
     "serviceType" : "EXACC",
     "hostName" : "host_Name",
     "description" : "Huge Pages is configured more than 90% of total memory
amount.",
     "eventName" : "CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE",
     "status" : "online"
   }
 },
 "eventID" : "d0724fac-9beb-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.

Copyright © 2008, 2026, Oracle and/or its affiliates.

You can also restrict automatic diagnostic collection at a granular level, for example, only for **ORA-00600** or even only **ORA-00600** with specific arguments.

Syntax

```
tfactl blackout add|remove|print
-targettype host|crs|asm|asmdg|database|dbbackup|db_dataguard|db_tablespace|
pdb_tablespace|pdb|listener|service|os
-target all|name
[-container name]
[-pdb pdb_name]
-event all|"event_str1,event_str2"|availability
[-timeout nm|nh|nd|none]
[-c|-local|-nodes "node1,node2"]
[-reason "reason for blackout"]
[-docollection]
```

Parameters

Table 7-21 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`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-21 (Cont.) tfactl blackout Command Parameters

Parameter	Description
`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.
`docollection`	Use this option to do an automatic diagnostic collection even if a blackout is set for this target.

Example 7-75 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
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

- To blackout **ALL events** on **target type:** host, **target:** all

```
tfactl blackout add -targettype host -event all -target all -timeout 1h -
reason "Disabling all events during patching"
```

- To print blackout details

```
tfactl blackout print
```

```
.------------------------------------------------------------------------
------------------------------------------------------------------------
-------------------------.
|

myhostname
                    |
+--------------+--------------------+----------
+-----------------------------+-----------------------------+--------
+--------------+--------------------------------------+
| Target Type   | Target              | Events    | Start
Time                   | End Time                    | Status | Do
Collection | Reason                               |
+--------------+--------------------+----------
+-----------------------------+-----------------------------+--------
+--------------+--------------------------------------+
| HOST          | ALL                 | ALL       | Thu Mar 24 16:48:39
UTC 2022 | Thu Mar 24 17:48:39 UTC 2022 | ACTIVE | false         |
Disabling all events during patching |
| DATABASE      | MYDB                | ORA-00600 | Thu Mar 24 16:39:03
UTC 2022 | Fri Mar 25 16:39:03 UTC 2022 | ACTIVE | false         |
NA                                   |
| DATABASE      | ALL                 | ORA-04031 | Thu Mar 24 16:39:54
UTC 2022 | Thu Mar 24 17:39:54 UTC 2022 | ACTIVE | false         |
NA                                   |
| DB_DATAGUARD  | MYDB                | ALL       | Thu Mar 24 16:41:38
UTC 2022 | Thu Mar 24 17:11:38 UTC 2022 | ACTIVE | false         |
NA                                   |
| DBBACKUP      | MYDB                | ALL       | Thu Mar 24 16:40:47
UTC 2022 | Fri Mar 25 16:40:47 UTC 2022 | ACTIVE | false         |
NA                                   |
| DB_TABLESPACE | SYSTEM_CDBNAME_MYDB | ALL       | Thu Mar 24 16:45:56
UTC 2022 | Thu Mar 24 17:15:56 UTC 2022 | ACTIVE | false         |
NA                                   |
'--------------+--------------------+----------
+-----------------------------+-----------------------------+--------
+--------------+--------------------------------------'
```

- To remove blackout for **event:** ORA-00600 on **target type:** database, **target:** mydb

```
tfactl blackout remove -targettype database -event "ORA-00600" -target mydb
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

- To remove blackout for **db backup events** on **target type:** dbbackup, **target:** mydb

  ```
  tfactl blackout remove -targettype dbbackup -target mydb
  ```

- To remove blackout for **db tablespace events** on **target type:** db_tablespace, **target:** system, **container:** mydb

  ```
  tfactl blackout remove -targettype db_tablespace -target system -container
  mydb
  ```

- To remove blackout for **host events** on **target type:** host, **target:** all

  ```
  tfactl blackout remove -targettype host -event all -target all
  ```

Remediation

These topics cover some common issues you might run into and how to address them.

- HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE
- HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION
- CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE
- CONFIGURATION.DB_GUEST.SSHD.INVALID
- HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT
- AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN
- AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN
- AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN
- AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN
- HEALTH.DB_CLUSTER.CDB.CORRUPTION
- HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG
- HEALTH.DB_CLUSTER.CDB.DATABASE_HANG
- HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE
- HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE
- Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Cloud@Customer
- HEALTH.DB_GUEST.PDB.HIGH_IO_LATENCY
- HEALTH.DB_GUEST.CDB.HIGH_IO_LATENCY_HISTOGRAM
- HEALTH.DB_GUEST.PDB.HIGH_CPU_WAITS
- HEALTH.DB_GUEST.PDB.EXCESSIVE_WAITS
- HEALTH.DB_GUEST.PDB.EXCESSIVE_CPU_THROTTLING
- HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_1
- HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_2
- HEALTH.DB_GUEST.PDB.DB_BLOCKER_STALL_CHECK
- HEALTH.DB_GUEST.PDB.TOTAL_EXADATA_CACHE_HIT_RATIO

Copyright © 2008, 2026, Oracle and/or its affiliates.

- HEALTH.DB_GUEST.CDB.HIGH_REDO_LOG_WAITS

HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE

Problem Statement: One or more VM guest file systems has free space below 10% free.

Risk: Insufficient VM guest file system free space can cause disk space allocation failure, which can result in wide-ranging errors and failures in Oracle software (Database, Clusterware, Cloud, Exadata).

Action:

Oracle Cloud and Exadata utilities run automatically to purge old log files and trace files created by Oracle software to reclaim file system space.

If the automatic file system space reclamation utilities cannot sufficiently purge old files to clear this event, then perform the following actions:

1. Remove unneeded files and/or directories created manually or by customer-installed applications or utilities. Files created by customer-installed software are outside the scope of Oracle's automatic file system space reclamation utilities. The following operating system command, run as the `opc` user, is useful for identifying directories consuming excessive disk space:

   ```
   $ sudo du -hx file-system-mount-point | sort -hr
   ```

   Only remove files or directories you are certain can be safely removed.
2. Reclaim `/u02` file system disk space by removing Database Homes that have no databases. For more information about managing Database Homes, see *Manage Oracle Database Homes on Exadata Database Service on Cloud@Customer Systems*.
3. (Exadata Cloud@Customer only) Increase `/u02` file system size. For more information about scaling up local storage, see *Introduction to Scale Up or Scale Down Operations*.
4. Open service request to receive additional guidance about reducing file system space use.

Related Topics

- Manage Oracle Database Homes on Oracle Exadata Database Service on Cloud@Customer Systems
  Learn to manage Oracle Database homes on Oracle Exadata Database Service on Cloud@Customer.
- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.

HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION

Problem Statement: A file system that is expected to be read-write can no longer be written to.

Risk: Oracle software (Linux, Database, Clusterware, Cloud, Exadata) requires write access to file systems to operate correctly.

Action:

`/u01` and `/u02` file systems:

Copyright © 2008, 2026, Oracle and/or its affiliates.

1. Stop running services, if any, that are using the file system, such as Oracle Clusterware, Trace File Analyzer (TFA), and Enterprise Manager (EM) agent.
2. Unmount the file system.
3. Run file system check and repair.
   - **ext4:** Refer to *Checking and Repairing a File System*.
   - **xfs:** Refer to *Checking and Repairing an XFS File System*.
   - If the file system cannot be repaired then open a service request with Oracle Support for assistance.
4. Mount the file system.
5. Start the services.

/ (root) file system:

Open a service request with Oracle Support for assistance.

- If there is VM access, then collect full `dmesg(1)` command output and provide it to Oracle Support.
- Note that / (root) file system repair is possible only with console access.

Related Topics

- Checking and Repairing a File System
- Checking and Repairing an XFS File System

CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE

Problem Statement: Too much VM memory is allocated for HugePages use.

Risk: Excessive memory allocated to HugePages may result in poor database performance, or the system running out of memory, experiencing excessive swapping, or having crucial system services fail, causing system crash or node eviction.

Action:

1. Reduce HugePages memory use. To determine the proper setting for operating system parameter vm.nr_hugepages, see My Oracle Support document 361323.1.
2. Scale up VM memory. For more information about scaling VM memory, see *Introduction to Scale Up or Scale Down Operations*.

Related Topics

- https://support.oracle.com/rs?type=doc&id=361323.1
- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.

CONFIGURATION.DB_GUEST.SSHD.INVALID

Problem Statement: SSHD configuration is unexpected.

SSHD Configuration Setting	Expected Value
`PubkeyAuthentication`	yes

Copyright © 2008, 2026, Oracle and/or its affiliates.

SSHD Configuration Setting	Expected Value
AuthorizedKeysFile	.ssh/authorized_keys This file must exist in root user home directory.
HostbasedAuthentication	no
IgnoreUserKnownHosts	yes
IgnoreRhosts	yes
PermitEmptyPasswords	no
PasswordAuthentication	no
ChallengeResponseAuthentication	no
GSSAPIAuthentication	no
UsePAM	yes
PrintMotd	no
UsePrivilegeSeparation	yes
PermitUserEnvironment	no
Compression	delayed
MaxStartups	100
AcceptEnv	Must contain one of the following: • LANG • LC_CTYPE • LC_NUMERIC • LC_TIME • LC_COLLATE • LC_MONETARY • LC_MESSAGES • LC_PAPER • LC_NAME • LC_ADDRESS • LC_TELEPHONE • LC_MEASUREMENT • LC_IDENTIFICATION • LC_ALL • LANGUAGE • XMODIFIERS
Subsystem	**sftp** /usr/libexec/openssh/sftp-server
Protocol	2
AddressFamily	inet

Risk: SSHD configuration is unexpected which may cause Oracle Cloud automation failure or prevent customer SSH access to the VM.

Action: Change SSHD to match expected configuration.

Copyright © 2008, 2026, Oracle and/or its affiliates.

1. Verify SSHD service is active.

```
$ sudo systemctl is-active sshd.service
active
```

If SSHD service is inactive, then start it.

```
$ sudo systemctl start sshd.service
```

2. Verify SSHD service is enabled.

```
$ sudo /opt/oracle.cellos/host_access_control ssh-service --status
[INFO] [IMG-SEC-1201] Service sshd is enabled {1}
```

If SSHD service is disabled, then enable it.

```
$ sudo /opt/oracle.cellos/host_access_control ssh-service --enable
```

3. Change SSHD configuration to match the expected values according to the table shown in the Problem Statement section above.

SSHD Configuration Setting	How to Change Current setting
`Ciphers`	`/opt/oracle.cellos/host_access_control sshciphers --help`
`MACs`	`/opt/oracle.cellos/host_access_control ssh-macs --help`
`PermitRootLogin`	`/opt/oracle.cellos/host_access_control rootssh --help`
`ClientAliveInterval`	`/opt/oracle.cellos/host_access_control idle-timeout --help`
`ClientAliveCountMax`	`/opt/oracle.cellos/host_access_control idle-timeout --help`
`ListenAddress`	`/opt/oracle.cellos/host_access_control ssh-listen --help`
ALL OTHER PARAMETERS	• Edit `/etc/ssh/sshd_config`. • Restart `sshd.service`.`$ sudo systemctl restart sshd.service`

HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT

Problem Statement: A CRITICAL Exachk check failed and should be reviewed and addressed as soon as possible.

Risk: A CRITICAL check is expected to impact a large number of customers AND should be addressed immediately (for example, within 24 hours) AND meets one or more of the following criteria:

1. On-disk corruption or data loss
2. Intermittent wrong results with Exadata feature usage (e.g. smart scan)
3. System wide availability impact

Copyright © 2008, 2026, Oracle and/or its affiliates.

4. Severe system wide performance impact seriously affecting application service Service Level Agreements (SLAs)
5. Compromised redundancy and inability to restore redundancy
6. Inability to update software in a rolling manner
7. Configuration error that could lead to an unexpected or unknown impact

Action:

Recommend that you bring up the EXAchk HTML report from the latest EXAchk zip file and click "**view**" on each CRITICAL check and follow the recommendation guidance that contains: Benefit/Impact, Risk, and Action/Repair guidance. Once the CRITICAL check is addressed, the next EXAchk run will pass that check. For more information about Oracle EXAchk, see *Oracle Exadata Database Machine Exachk (Doc ID 1070954.1)*.

As the `root` user, you can re-run EXAchk command by issuing:

```
/usr/bin/exachk -profile exatier1 -noupgrade -dball
```

If the check results are returning false data, then log a Service Request.

If there is a CRITICAL check that needs to be temporarily excluded, then follow the "**Skipping Specific Best Practice Checks in Exadata Cloud**" section of *Oracle Exadata Database Machine Exachk (Doc ID 1070954.1)*.

Related Topics

- Oracle Exadata Database Machine Exachk (Doc ID 1070954.1)

AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN

Problem Statement: The Cluster Ready Stack is in an offline state or has failed.

Risk: If the Cluster Ready Service is offline on a node, then the node cannot provide database services for the application.

Action:

1. Check if CRS was stopped by your administrator, as part of a planned maintenance event, or a scale up or down of local storage.
   a. The following patching events will stop CRS:
      i. GRID Patching
      ii. Exadata VM patching of Guest
      iii. Exadata VM Patching of Host
2. If CRS has stopped unexpectedly, then the current status can be checked by issuing the `crsctl check crs` command.
   a. If the node is not responding, then the VM node may be rebooting. Wait for the node reboot to finish, CRS will normally be started through the `init` process.
3. If CRS is still down, then investigate the cause of the failure by referring to the `alert.log` found in `/u01/app/grid/diag/crs/<node_name>/crs/trace`.
   Review the log entries corresponding to the date/time of the down event. Act on any potential remediation.
4. Restart the CRS, by issuing the `crsctl start crs` command.

Copyright © 2008, 2026, Oracle and/or its affiliates.

5. A successful restart of CRS will generate the clearing event: `AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN_CLEARED`.

AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN

Problem Statement: A SCAN listener is down and unable to accept application connections.

Risk: If all SCAN listeners are down, then application connections to the database through the SCAN listener will fail.

Action:

Start the SCAN listener to receive the `DOWN_CLEARED` event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Grid Infrastructure software update, then no action is required. The affected SCAN listener will automatically failover to an available instance.
2. If the event was caused by user action, then start the SCAN listener at the next opportunity.

DOWN event of type CRITICAL

Check SCAN status and restart the SCAN listener.

1. Login to the VM as `opc` user and `sudo` to the `grid` user:

   ```
   [opc@vm ~] sudo su - grid
   ```

2. Check the SCAN listener status on any node:

   ```
   [grid@vm ~] srvctl status scan_listener
   ```

3. Start the SCAN listener:

   ```
   [grid@vm ~] srvctl start scan_listener
   ```

4. Recheck the SCAN listeners status on any node:
   If the `scan_listener` is still down, then investigate the cause of the scan listener failure:

   a. Collect both the CRS and operating system logs 30 minutes prior and 10 minutes for the <*hostName*>indicated in the log. Note the time in the event payload is always provided in UTC. For `tfactl` collection, adjust the time to the timezone of the VM Cluster.

      ```
      [grid@vm ~] tfactl diagcollect -crs -os -node <hostName> –from
      "<eventTime adjusted for local vm timezone> - 30 minute " -to
      "<eventTime adjusted for local vm timezone> + 10 minutes"
      ```

   b. Review the SCAN listener log located under `/u01/app/grid/diag/tnslsnr/<hostName>/<listenerName>/trace`

AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN

Problem Statement: A client listener is down and unable to accept application connections.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Risk:

- If the node's client listener is down, then the database instances on the node cannot provide services for the application.
- If the client listener is down on all nodes, then any application that connects to any database using the SCAN or VIP will fail.

Action:

Start the client listener to receive the `DOWN_CLEARED` event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Grid Infrastructure software update, then no action is required. The affected client listener will automatically restart when maintenance affecting the grid instance is complete.
2. If the event was caused by user action, then start the client listener at the next opportunity.

DOWN event of type CRITICAL

Check the client listener status and then restart the client listener.

1. Login to the VM as `opc` user and `sudo` to the `grid` user:

   ```
   [opc@vm ~] sudo su - grid
   ```

2. Check the client listener status on any node:

   ```
   [grid@vm ~] srvctl status listener
   ```

3. Start the client listener:

   ```
   [grid@vm ~] srvctl start listener
   ```

4. Recheck the client listener status on any node:
   If the client listener is still down, then investigate the cause of the client listener failure:

   a. Use tfactl to collect both the CRS and operating system logs 30 minutes prior and 10 minutes for the <*hostName*> indicated in the log. Note the time in the event payload is always provided in UTC. For tfactl collection, adjust the time to the timezone of the VM Cluster.

      ```
      [grid@vm ~] tfactl diagcollect -crs -os -node <hostName> –from
      "<eventTime adjusted for local vm timezone> - 30 minute " -to
      "<eventTime adjusted for local vm timezone> + 10 minutes"
      ```

   b. Review the listener log located under `/u01/app/grid/diag/tnslsnr/<hostName>/<listenerName>/trace`

AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN

Problem Statement: A database instance has gone down.

Risk: A database instance has gone down, which may result in reduced performance if database instances are available on other nodes in the cluster, or complete downtime if database instances on all nodes are down.

Action:

Copyright © 2008, 2026, Oracle and/or its affiliates.

Start the database instance to receive the `DOWN_CLEARED` event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Database Home software update, then no action is required. The affected database instance will automatically restart when maintenance affecting the instance is complete.
2. If the event was caused by user action, then start the affected database instance at the next opportunity.

DOWN event of type CRITICAL

1. Check database status and restart the down database instance.
   a. Login to the VM as `oracle` user:
   b. Set the environment:

      ```
      [oracle@vm ~] . <dbName>.env
      ```

   c. Check the database status:

      ```
      [oracle@vm ~] srvctl status database -db <dbName>
      ```

   d. Start the database instance:

      ```
      [oracle@vm ~] srvctl start instance -db <dbName> -instance
      <instanceName>
      ```

2. Investigate the cause of the database instance failure.
   a. Review Trace File Analyzer (TFA) events for the database:

      ```
      [oracle@vm ~] tfactl events -database <dbName> -instance <instanceName>
      ```

   b. Review the database alert log located at `$ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/alert_<instanceName>.log`

HEALTH.DB_CLUSTER.CDB.CORRUPTION

Problem Statement: Corruptions can lead to application or database errors and in worse case result in significant data loss if not addressed promptly.

A corrupt block is a block that was changed so that it differs from what Oracle Database expects to find. Block corruptions can be categorized as physical or logical:

- In a physical block corruption, which is also called a media corruption, the database does not recognize the block at all; the checksum is invalid or the block contains all zeros. An example of a more sophisticated block corruption is when the block header and footer do not match.
- In a logical block corruption, the contents of the block are physically sound and pass the physical block checks; however, the block can be logically inconsistent. Examples of logical block corruption include incorrect block type, incorrect data or redo block sequence number, corruption of a row piece or index entry, or data dictionary corruptions.

For more information, see *Physical and Logical Block Corruptions. All you wanted to know about it. (Doc ID 840978.1)*.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Block corruptions can also be divided into interblock corruption and intrablock corruption:

- In an intrablock corruption, the corruption occurs in the block itself and can be either a physical or a logical block corruption.
- In an interblock corruption, the corruption occurs between blocks and can only be a logical block corruption.

Oracle checks for the following errors in the `alert.log`:

- ORA-01578
- ORA-00752
- ORA-00753
- ORA-00600 [3020]
- ORA-00600 [kdsgrp1]
- ORA-00600 [kclchkblk_3]
- ORA-00600 [13013]
- ORA-00600 [5463]

Risk: A data corruption outage occurs when a hardware, software, or network component causes corrupt data to be read or written. The service-level impact of a data corruption outage may vary, from a small portion of the application or database (down to a single database block) to a large portion of the application or database (making it essentially unusable). If remediation action is not taken promptly, then potential downtime and data loss can increase.

Action:

The current event notification currently triggers on physical block corruptions (ORA-01578), lost writes (ORA-00752, ORA-00753 and ORA-00600 with first argument 3020), and logical corruptions (typical detected from ORA-00600 with first argument of kdsgrp1, kdsgrp1, kclchkblk_3, 13013 OR 5463).

Oracle recommends the following steps:

1. Confirm that these corruptions were reported in the alert.log trace file. Log a Service Request (SR) with latest EXAchk report, excerpt of the alert.log and trace file containing the corruption errors, any history of recent application, database or software changes and any system, clusterware and database logs for the same time period. For all these cases, a TFA collection should be available and should be attached to the SR.
2. For repair recommendations, refer to *Handling Oracle Database Corruption Issues (Doc ID 1088018.1)*.

For physical corruptions or ORA-1578 errors, the following notes will be helpful:

- Doc ID 1578.1 : OERR: ORA-1578 "ORACLE data block corrupted (file # %s, block # %s)" Primary Note
- Doc ID 472231.1 : How to identify all the Corrupted Objects in the Database reported by RMAN
- Doc ID 819533.1 : How to identify the corrupt Object reported by ORA-1578 / RMAN / DBVERIFY
- Depending on the object that has the corruption, follow the guidance in Doc ID 1088018.1. Note RMAN can be used to recover one or many data block that are physically corrupted. Also using Active Data Guard with real time apply, auto block repair of physical data corruptions would have occurred automatically.

Copyright © 2008, 2026, Oracle and/or its affiliates.

For logical corruptions caused by lost writes (ORA-00752, ORA-00753 and ORA-00600 with first argument 3020) on the primary or standby databases, they will be detected on the primary or with standby's redo apply process. The following notes will be helpful:

- Follow the guidance, follow Doc ID 1088018.1.
- If you have a standby and lost write corruption on the primary or standby, refer to Resolving ORA-00752 or ORA-00600 [3020] During Standby Recovery (Doc ID 1265884.1)

For logical corruptions (typical detected from ORA-00600 with arguments of kdsgrp1, kclchkblk_3, 13013 OR 5463):

- Follow the guidance, follow Doc ID 1088018.1 for specific guidance on the error that was detected.
- If you have a standby and logical corruption on the primary, refer to Resolving Logical Block Corruption Errors in a Physical Standby Database (Doc ID 2821699.1)

Related Topics

- Physical and Logical Block Corruptions. All you wanted to know about it. (Doc ID 840978.1)
- OERR: ORA-1578 "ORACLE data block corrupted (file # %s, block # %s)" Primary Note (Doc ID 1578.1)
- How to identify all the Corrupted Objects in the Database with RMAN (Doc ID 472231.1)
- How to identify the corrupt Object reported by ORA-1578 / RMAN / DBVERIFY (Doc ID 819533.1)
- Resolving ORA-00752 or ORA-600 [3020] During Standby Recovery (Doc ID 1265884.1)
- Resolving Logical Block Corruption Errors in a Physical Standby Database (Doc ID 2821699.1)

HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG

Problem Statement: CDB RAC Instance may temporarily or permanently stall due to the log writer's (LGWR) inability to write the log buffers to an online redo log. This occurs because all online logs need archiving. Once the archiver (ARC) can archive at least one online redo log, LGWR will be able to resume writing the log buffers to online redo logs and the application impact will be alleviated.

Risk: If the archiver hang is temporary, then this can result in a small application brown out or stall for application processes attempting to commit their database changes. If the archiver is not unblocked, applications can experience extended delay in processing.

Action:

- See, *Script To Find Redo log Switch History And Find Archivelog Size For Each instance In RAC (Doc ID 2373477.1)* to determine the hourly frequency for each thread/instance.
- If any hourly bucket is greater than 12, then consider resizing the online redo logs. See item 2 below for resizing steps.
- If the database hangs are temporary, then the archiver may be unable to keep up with the redo log generated. Check the `alert.log`, `$ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/alert_<instanceName>.log`, for "All online logs need archiving", multiple events in a short period can indicate 2 possible solutions.
  - If the number of redo logs groups per thread is less than 4, then consider adding additional logs groups to reach 4, see item 1 below for add redo log steps.

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - The other possible solution is to resize the redo logs, see item 2 below for resizing steps.

- For Data Guard and Non Data Guard review the *Configure Online Redo Logs Appropriately* of section Oracle Database High Availability Overview and Best Practices for sizing guidelines.

1. Add a redo log group for each thread. The additional redo log should equal the current log size.
   a. Use the following query:

      ```
      select max(group#) Ending_group_number, thread#, count(*)
      number_of_groups_per_thread, bytes redo_size_in_bytes from v$log group
      by thread#,bytes
      ```

   b. Add one new group per thread using the same size as the current redo logs.

      ```
      alter database add logfile thread <thread_number> Group <max group + 1>
      ('<DATA_DISKGROUP>') size <redo_size_in_bytes>
      ```

2. Resize the online redo logs by adding larger redo logs and dropping the current smaller redo logs.
   a. Use the following query:

      ```
      select max(group#) Ending_group_number, thread#, count(*)
      number_of_groups_per_thread, bytes redo_size_in_bytes from v$log group
      by thread#,bytes
      ```

   b. Add the same number of redo logs for each thread <*number_of_groups_per_thread*> that currently exist. The <*new_redo_size_in_bytes*> should be based on *Configure Online Redo Logs Appropriately* of section Oracle Database High Availability Overview and Best Practices.

      i. `alter database add logfile thread <thread_number> Group <max group + 1> ('<DATA_DISKGROUP>') size <new_redo_size_in_bytes>`

      ii. The original smaller redo logs should be deleted. A redo log can only be deleted if its status is inactive.
      To determine the status of a redo logs issue:

      ```
      select group#, thread#, status, bytes from v$log order by bytes,
      group#, thread#;
      ```

      To delete the original smaller redo logs:

      ```
      alter database drop logfile <group#>
      ```

- If the database is hung, the primary log archive destination and alternate may be full. Review the *HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE* for details on freeing space in RECO and DATA disk groups.

Related Topics

- Script To Find Redolog Switch History And Find Archivelog Size For Each Instances In RAC (Doc ID 2373477.1)

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Configure Online Redo Logs Appropriately
- HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

HEALTH.DB_CLUSTER.CDB.DATABASE_HANG

Problem Statement: Hang management detected a process hang and generated a ORA-32701 error message. Additionally, this event may be raised if Diagnostic Process (DIA0) process detects a hang in a critical database process.

Risk: A hang can indicate resource, operating system, or application coding related issues.

Action:

Investigate the cause of the session hang.

1. Review TFA events for the database for the following message patterns corresponding to the date/time of the event: ORA-32701, "DIA0 Critical Database Process Blocked" or "DIA0 Critical Database Process As Root".

   ```
   [oracle@vm ~] tfactl events -database <dbName> -instance <instanceName>
   ```

2. Review the `alert.log` file.

   ```
   $ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/
   alert_<instanceName>.log
   ```

3. **For ora-32701:** An overloaded system can cause slow progress, which can be interpreted as a hang.
   The hang manager may attempt to resolve the hang by terminating the final blocker process.
4. **For DIA0 Critical Database Process messages:** Review the related diagnostic lines indicating the process and the reason for the hang.

HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE

Problem Statement: A daily incremental BACKUP of the CDB failed.

Risk: A failure of the backup can compromise the ability to use the backups for restore/recoverability of the database. Recoverability Point Object (RPO) and the Recoverability Time Object (RTO) can be impacted.

Action:

Review the RMAN logs corresponding to the date/time of the event. Note the event time stamp <*eventTime*> is in UTC, adjust as necessary for the VM's timezone.

- For Exadata Database Service on Cloud@Customer Oracle Managed Backups or User Configured Backups under `dbaascli`:
  - RMAN output can be found at `/var/opt/oracle/log/<DB_NAME>/obkup`. Daily incremental logs have the following format `obkup_yyyy-mm-dd_24hh:mm:ss.zzzzzzzzzzzz.log` within the `obkup` directory. The logs are located on the lowest active node/instance of the database when the backup was initiated.
  - Review the log for any failures:

Copyright © 2008, 2026, Oracle and/or its affiliates.

    - If the failure is due to an external event outside of RMAN, for example the backup location was full or a networking issue, resolve the external issue.
    - For other RMAN script errors, collect the diagnostic logs and open a Service Request. See *DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check*.
  - If the issue is transient or is resolved, take a new incremental backup: See *dbaascli database backup*.
- For Customer owned and managed backup taken through RMAN:
  - Review the RMAN logs for the backup.

Related Topics

- DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check
- dbaascli database backup
  To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

Problem Statement: ASM disk group space usage is at or exceeds 90%.

Risk: Insufficient ASM disk group space can cause database creation failure, tablespace and data file creation failure, automatic data file extension failure, or ASM rebalance failure.

Action:

ASM disk group used space is determined by the running the following query while connected to the ASM instance.

```
[opc@node ~] sudo su - grid
[grid@node ~] sqlplus / as sysasm

SQL> select 'ora.'||name||'.dg', total_mb, free_mb, round ((1-(free_mb/
total_mb))*100,2) pct_used from v$asm_diskgroup;

NAME                           TOTAL_MB    FREE_MB   PCT_USED
------------------------------ ---------- ---------- ----------
ora.DATAC1.dg                   75497472    7408292      90.19
ora.RECOC1.dg                   18874368   17720208       6.11
```

ASM disk group capacity can be increased in the following ways:

1. Scale Exadata VM Cluster storage to add more ASM disk group capacity. See *Introduction to Scale Up or Scale Down Operations*.
2. Scale Exadata Infrastructure storage to add more ASM disk group capacity. See *Overview of Elastic Storage Expansion*.

DATA disk group space use can be reduced in the following ways:

1. Drop unused data files and temp files from databases. See *Dropping Data Files*.
2. Terminate unused databases (e.g. test databases). See *Using the Console to Terminate a Database*.

RECO disk group space use can be reduced in the following ways:

Copyright © 2008, 2026, Oracle and/or its affiliates.

1. Drop unnecessary Guaranteed Restore Points. See *Using Normal and Guaranteed Restore Points*.
2. Delete archived redo logs or database backups already backed up outside the Flash Recovery Area (FRA). See *Maintaining the Fast Recovery Area*.

SPARSE disk group space use can be reduced in the following ways:

1. Move full copy test master databases to another disk group (e.g. DATA).
2. Drop unused snapshot databases or test master databases. See *Managing Exadata Snapshots*.

For more information about managing the log and diagnostic files, see *Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Cloud@Customer*.

Related Topics

- Introduction to Scale Up or Scale Down Operations
  With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.
- Overview of Elastic Storage Expansion
  With elastic storage expansion, you can dynamically increase your storage capacity to meet your growing workload requirements.
- Dropping Data Files
- Using the Console to Terminate a Database
  You can terminate a database and thereby remove the terminated database from the Cloud Control Plane.
- Using Normal and Guaranteed Restore Points
- Maintaining the Fast Recovery Area
- Managing Exadata Snapshots
- Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Cloud@Customer

Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Cloud@Customer

The software components in Oracle Exadata Database Service on Cloud@Customer generate a variety of log and diagnostic files, and not all these files are automatically archived and purged. Thus, managing the identification and removal of these files to avoid running out of file storage space is an important administrative task.

Database deployments on ExaDB-C@C include the `cleandblogs` script to simplify this administrative task. The script runs daily as a `cron` job on each compute node to archive key files and remove old log and diagnostic files.

The `cleandblogs` script operates by using the `adrci` (Automatic Diagnostic Repository [ADR] Command Interpreter) tool to identify and purge target diagnostic folders and files for each Oracle Home listed in `/etc/oratab`. It also targets Oracle Net Listener logs, audit files, and core dumps.

On ExaDB-C@C, the script is run separately as the `oracle` user to clean log and diagnostic files that are associated with Oracle Database, and as the `grid` user to clean log and diagnostic files that are associated with Oracle Grid Infrastructure.

Copyright © 2008, 2026, Oracle and/or its affiliates.

The `cleandblogs` script uses a configuration file to determine how long to retain each type of log or diagnostic file. You can edit the file to change the default retention periods. The file is located at `/var/opt/oracle/cleandb/cleandblogs.cfg` on each compute node.

> **Note**
>
> Configure an optimal retention period for each type of log or diagnostic file. An insufficient retention period will hinder root cause analysis and problem investigation.

Parameter	Description and Default Value
`AlertRetention`	Alert log (`alert_instance.log`) retention value in days. Default value: 14
`ListenerRetention`	Listener log (`listener.log`) retention value in days. Default value: 14
`AuditRetentionDB`	Database audit (`*.aud`) retention value in days. Default value: 1
`CoreRetention`	Core dump/files (`*.cmdp*`) retention value in days. Default value: 7
`TraceRetention`	Trace file (`*.tr*` and `*.prf`) retention value in days. Default value: 7
`longpRetention`	Data designated in the Automatic Diagnostic Repository (ADR) as having a long life (the `LONGP_POLICY` attribute). For information about ADR, see *Automatic Diagnostic Repository (ADR)* in the *Oracle Database Administrator's Guide*. Default value: 14
`shortpRetention`	Data designated in the Automatic Diagnostic Repository (ADR) as having a short life (the `SHORTP_POLICY` attribute). For information about ADR, see *Automatic Diagnostic Repository (ADR)* in the *Oracle Database Administrator's Guide*. Default value: 7
`LogRetention`	Log file retention in days for files under `/var/opt/oracle/log` and log files in ACFS under `/var/opt/oracle/dbaas_acfs/log`. Default value: 14
`LogDirRetention`	`cleandblogs` logfile retention in days. Default value: 14
`ScratchRetention`	Temporary file retention in days for files under `/scratch`. Default value: 7

Copyright © 2008, 2026, Oracle and/or its affiliates.

Archiving Alert Logs and Listener Logs

When cleaning up alert and listener logs, `cleandblogs` first archives and compresses the logs, operating as follows:

1. The current log file is copied to an archive file that ends with a date stamp.
2. The current log file is emptied.
3. The archive file is compressed using `gzip`.
4. Any existing compressed archive files older than the retention period are deleted.

Running the cleandblogs Script Manually

The `cleandblogs` script automatically runs daily on each compute node, but you can also run the script manually if the need arises.

1. Connect to the compute node as the `oracle` user to clean log and diagnostic files that are associated with Oracle Database, or connect as the `grid` user to clean log and diagnostic files that are associated with Oracle Grid Infrastructure.
   For detailed instructions, see *Connecting to a Virtual Machine with SSH*.

   Change to the directory containing the cleandblogs script:

   ```
   $ cd /var/opt/oracle/cleandb
   ```

2. Run the `cleandblogs` script:

   ```
   $ ./cleandblogs.pl
   ```

   When running the script manually, you can specify an alternate configuration file to use instead of `cleandblogs.cfg` by using the `--pfile` option:

   ```
   $ ./cleandblogs.pl --pfile config-file-name
   ```

3. Close your connection to the compute node:

   ```
   $ exit
   ```

Related Topics

- Automatic Diagnostic Repository (ADR)
- Connecting to a Virtual Machine with SSH
  You can connect to the virtual machines in an Oracle Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

HEALTH.DB_GUEST.PDB.HIGH_IO_LATENCY

Problem Statement: Higher I/O latency can lead to significant IO issues.

Risk: Increased I/O latency directly translates to slower database operations and SQL query response times. High latency can cause significant database performance degradation, especially for OLTP workloads.

Action: To investigate high cell single block physical read waits on Exadata, follow My Oracle Support (MOS) notes: 2119510.1 and 2530864.1.

Copyright © 2008, 2026, Oracle and/or its affiliates.

HEALTH.DB_GUEST.CDB.HIGH_IO_LATENCY_HISTOGRAM

Problem Statement: High wait count values pertaining to IO wait class exceeding 32 ms, can indicate a potential problem with underlying Exadata Infrastructure.

Risk: Increased I/O latency directly translates to slower database operations and SQL query response times. High latency can cause significant.

Action: To investigate high cell single block physical read waits on Exadata, follow My Oracle Support (MOS) notes: 2119510.1 and 2530864.1.

HEALTH.DB_GUEST.PDB.HIGH_CPU_WAITS

Problem Statement: CPU intensive CDB/PDB sessions waiting on latch/cursor related wait-events for more than 300 seconds and CPU intensive CDB/PDB sessions waiting on Library cache related wait-events for more than 900 seconds.

Risk: These CPU intensive Wait events generally impact the overall system performance. Since mutexes are CPU intensive resource, in the event of mutex contention, CPU usage can rise and can impact user sessions.

Action: To investigate high CPU intensive waits in the Database, follow My Oracle Support (MOS) notes: 1349387.1, 1357946.1, 1377998.1 and 444560.1, which provide specific troubleshooting guidance for various scenarios.

HEALTH.DB_GUEST.PDB.EXCESSIVE_WAITS

Problem Statement: The Foreground Wait Event is consuming over 20% of total DB time, excluding the 'IDLE' wait-event class and Exadata User I/O wait-events. The Exadata User I/O wait-event alone is accounting for more than 35% of total DB time.

Risk: High wait event times indicate that the database is spending a significant amount of time waiting for resources, processes, or I/O operations, leading to performance bottlenecks and potentially impacting application responsiveness.

Action: To investigate high waits in the Database, follow My Oracle Support (MOS) notes: 1377446.1.

HEALTH.DB_GUEST.PDB.EXCESSIVE_CPU_THROTTLING

Problem Statement: Excessive CPU throttling observed in the Database.

Risk: Excessive CPU throttling in Database can occur when the resource manager limits do not match the consumer group utilization, potentially leading to performance issues. If the throttling occurs frequently, it can lead to performance problems, such as slow queries or application response times.

Action: To troubleshoot excessive CPU throttling in the Database, follow My Oracle Support (MOS) notes: 1339769.1.

HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_1

Problem Statement: VM cluster's RECO disk group or the database's FAST RECOVERY AREA (FRA) usage exceeds 90%.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Action: Convert snapshot back to physical standby using `dbaascli`:

```
dbaascli dataguard convertStandby --standbyType physical --dbname <dbname>
```

HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_2

Problem Statement: VM cluster's RECO disk group or the database's FAST RECOVERY AREA (FRA) usage exceeds 95%.

Action: The standby database must be rolled forward using the SQL*PLUS `RECOVER DATABASE FROM SERVICE` statement.

Convert back to physical database

```
dbaascli dataguard convertStandby --dbname <db name> --standbyType physical
```

Stop managed recovery and redo transport

Stop transport and apply of redologs on the standby database.

```
DGMGRL> edit database <primary> set state=TRANSPORT-OFF;
DGMGRL> edit database <standby> set state=APPLY-OFF;
```

Save the existing logfile names

The roll forward process will create new online redo logs and standby redo logs, leaving the existing logs as orphans that consume space. Use the following query to save the names of these logfiles so they can be removed later in the process.

```
SQL> set heading off linesize 999 pagesize 0 feedback off trimspool on
SQL> spool /tmp/delete_logfiles.log
SQL> select member from v$logfile;
SQL> spool off;
```

Save the current RMAN configuration settings

Save the standby database RMAN configuration so it can be reapplied after restoring the control file.

```
$ rman target / nocatalog log=/tmp/RMAN_settings.log <<EOF
show all;
$ grep ^CONFIGURE /tmp/RMAN_settings.log | grep -v 'RETENTION POLICY' >/tmp/
RMAN_settings.rman
$ rm /tmp/RMAN_settings.log
```

Refresh the standby control file from the primary database

1. Before deleting backups, determine whether AL backups taken by the snapshot database need to be retained.
   - For OSS, NAS, and disk backups, these backups will no longer be tracked by the database and therefore will not be included in retention-based deletion.
   - Such backups must be manually deleted directly from their backup destination if no longer needed.

Copyright © 2008, 2026, Oracle and/or its affiliates.

2. If backups are no longer needed, they can be deleted using RMAN. This avoids the need for manual deletion in the future.

```
RMAN> delete force noprompt backup device type 'sbt_tape';
RMAN> delete force noprompt archivelog all device type 'sbt_tape';
RMAN> delete force noprompt datafilecopy all device type 'sbt_tape';
```

3. Restore the control file from the primary database.

```
$ srvctl stop database -d <db> -o immediate
$ rman target / nocatalog
RMAN> startup nomount
RMAN> restore standby controlfile from service <tns alias for primary
database>;
RMAN> alter database mount;
```

4. Replace the RMAN configuration settings and remove the orphaned online and standby redo log files.

```
RMAN> @/tmp/RMAN_settings.rman
<output from CONFIGURE commands in the script>
RMAN> exit
$ rm /tmp/RMAN_settings.rman
```

5. Restore the datafiles:

   a. Switch to the primary database's current incarnation. From the primary database, run:

```
RMAN> list incarnation;
list incarnation;
using target database control file instead of recovery catalog

List of Database Incarnations
DB Key  Inc Key DB Name  DB ID            STATUS  Reset SCN  Reset Time
------- ------- -------- ---------------- --- ---------- ----------
1       1       OGG_DG   1153531161       PARENT  1          21-JAN-25
2       2       OGG_DG   1153531161       CURRENT 1514702    01-JUL-25
```

   b. From the standby database, run:

```
RMAN> reset database to incarnation <>;
```

6. Roll the Standby Database Forward: The standby database is now ready to perform the incremental roll forward using the recover-from-service operation.
Restart all instances to mount.

```
$ srvctl stop database -db <dbname> -o immediate
$ srvctl start database -db <dbname> -o mount
```

7. Re-Enable Redo Transport: Archived logs generated during the roll forward are required to bring the database to a consistent state. It is more efficient to allow the primary database to ship redo to the standby while the roll forward is in progress, rather than waiting for the logs to be transported after the roll forward completes.

Copyright © 2008, 2026, Oracle and/or its affiliates.

a. Check the log archive destination for this standby database and note the destination number.

```
SQL> show parameter LOG_ARCHIVE_DEST_;
NAME                      TYPE    VALUE
------------------------------------ -----------
------------------------------
log_archive_dest_2  string  service="snby002", ASYNC NOAFF
                            IRM delay=0 optional compressi
                            on=disable max_failure=0 reope
                            n=300 db_unique_name="snby002"
                            net_timeout=30, valid_for=(on
                            line_logfile,all_roles)
```

b. Enable redo transport.

```
DGMGRL> ALTER SYSTEM SET LOG_ARCHIVE_DEST_STATE_2=ENABLE;
```

c. Using Data Guard Broker, perform this operation from either the primary or the standby database.

```
DGMGRL> edit database <primary> set state=TRANSPORT-ON;
```

8. Run Recover From Service: The recovery command can use all available standby instances during the recovery process.

```
$ rman target sys/<password>  <- It is necessary to connect with the
password
RMAN > run {

allocate channel c1 type disk connect '/@<standby instance 1 SID_NAME>';
allocate channel c2 type disk connect '/@<standby instance 1 SID_NAME>';
allocate channel c3 type disk connect '/@<standby instance 1 SID_NAME>';
allocate channel c4 type disk connect '/@<standby instance 1 SID_NAME>';
allocate channel c5 type disk connect '/@<standby instance 2 SID_NAME>';
allocate channel c6 type disk connect '/@<standby instance 2 SID_NAME>';
allocate channel c7 type disk connect '/@<standby instance 2 SID_NAME>';
allocate channel c8 type disk connect '/@<standby instance 2 SID_NAME>';
restore database from service '<primary unique name>' section size 64G;
recover database from service '<primary unique name>' section size 64G;
}
```

9. Recover Until Consistent: At the completion of the `RECOVER DATABASE FROM SERVICE` command, additional recovery is required to make the standby database consistent before you can re-enable Flashback Database and open it in read-only mode.
A database is considered consistent when the control file and all datafiles are at the same SCN. The amount of additional recovery needed depends on how long the initial recovery took and the level of activity on the primary during that period.

   As a first step, switch a log on the primary database to ensure that the last of the redo generated during the recovery process is archived.

   From the primary database, run:

```
SQL> alter system archive log current;
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

10. Re-Enable Flashback Database: Restoring the standby control file automatically disables Flashback Database. Once the standby is brought to a consistent state, you can re-enable Flashback Database.

    ```
    SQL> alter database flashback on;
    ```

11. Restart the standby database and re-enable managed recovery.

    ```
    $ srvctl stop database -db <dbname> -o immediate
    $ srvctl start database -db <dbname> -o 'read only'
    ```

12. Enable redo apply.

    ```
    DGMGRL> edit database <standby> set state=APPLY-ON;
    ```

HEALTH.DB_GUEST.PDB.DB_BLOCKER_STALL_CHECK

Problem Statement: Identify Stalled DB Sessions Blocker

Risk: A high number of blocked or hanging sessions can lead to a database hang if critical operations wait indefinitely, and can also negatively impact application performance.

Action: To troubleshoot DB Blocker, follow My Oracle Support (MOS) notes: 1378583.1.

HEALTH.DB_GUEST.PDB.TOTAL_EXADATA_CACHE_HIT_RATIO

Problem Statement: Identify Exadata Cache Efficiency

Risk: If the Consolidated Exadata Cache Hit Ratio (%) drops below 90%, it suggests suboptimal cache utilization, potentially contributing to increased physical I/O and overall database performance degradation.

Action: To troubleshoot the low Exadata Cache Hit Ratio, refer to Monitoring Smart Flash Cache – Common Issues.

HEALTH.DB_GUEST.CDB.HIGH_REDO_LOG_WAITS

Problem Statement: Elevated redo log write waits observed

Risk: Prolonged redo log write waits can significantly delay transaction commits, degrade throughput, and lead to overall database performance degradation.

Action: To troubleshoot high redo log waits, refer to the following My Oracle Support (MOS) notes: 223117.1, 2634755.1, and 1376916.1.

Compute Expansion Event Types

Review the list of event types that compute expansion emits.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-22 Compute Expansion Events

Friendly Name	Event Type
VM Cluster Network - Add DB Server Network Begin	com.oraclecloud.databaseservice.adddbservervmclusternetwork.begin
VM Cluster Network - Add DB Server Network Begin	com.oraclecloud.databaseservice.adddbservervmclusternetwork.end
VM Cluster Network - Remove DB Server Network Begin	com.oraclecloud.databaseservice.removedbservervmclusternetwork.begin
VM Cluster Network - Remove DB Server Network End	com.oraclecloud.databaseservice.removedbservervmclusternetwork.end

Example 7-76 Compute Expansion Event Types Examples

This is a reference event for VM Cluster Network - Add DB Server Network Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrdsa7ybmw6mjndfutmibmw6a4ztvu7sgm562ojbudo3l
e6qwsspdq",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.adddbservervmclusternetwork.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",
      "compartmentName": "example_name",
      "resourceName": "my_vmcluster_network",
      "resourceId": "VmClusterNetwork-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "UPDATING",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "exadataInfrastructureId": "ExadataInfra-unique_ID",
        "displayName": "testDisplayName"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "VM Cluster Network - Add DB Server Network Begin",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "eventType":
"com.oraclecloud.databaseservice.adddbservervmclusternetwork.begin",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "timeUpdated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "lifecycleDetails", "type": ["null", "string"] },
    { "name": "exadataInfrastructureId", "type": ["null", "string"] }
  ],
  "timeCreated": "2019-06-26T13:31:31.979Z"
}
```

This is a reference event for VM Cluster Network - Add DB Server Network Begin:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrupbmw5ksic23bmwnqc4rssbmwc23eslydoumcnr6rcu
m6mtmgg3a",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.adddbservervmclusternetwork.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",
      "compartmentName": "example_name",
      "resourceName": "my_vmcluster_network",
      "resourceId": "VmClusterNetwork-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "NEEDS_ATTENTION",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "displayName": "testDisplayName",
        "exadataInfrastructureId": "ExadataInfra-unique_ID"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "VM Cluster Network - Add DB Server Network End",
  "eventType":
"com.oraclecloud.databaseservice.adddbservervmclusternetwork.end",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "timeUpdated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "lifecycleDetails", "type": ["null", "string"] },
    { "name": "exadataInfrastructureId", "type": ["null", "string"] }
  ],
  "timeCreated": "2019-06-26T13:31:31.979Z"
}
```

This is a reference event for VM Cluster Network - Remove DB Server Network Begin:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrdsa7ybmw6mjndfutmifv46a4ztvu7sgm562ojbudo3l
e6qwsspdq",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.removedbservervmclusternetwork.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",
      "compartmentName": "example_name",
      "resourceName": "my_vmcluster_network",
      "resourceId": "VmClusterNetwork-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "UPDATING",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "exadataInfrastructureId": "ExadataInfra-unique_ID",
        "displayName": "testDisplayName"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "VM Cluster Network - Remove DB Server Network Begin",
  "eventType":
"com.oraclecloud.databaseservice.removedbservervmclusternetwork.begin",
  "additionalDetails": [
    { "name": "id", "type": "string" },
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    { "name": "timeCreated", "type": "string" },
    { "name": "timeUpdated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "lifecycleDetails", "type": ["null", "string"] },
    { "name": "exadataInfrastructureId", "type": ["null", "string"] }
  ],
  "timeCreated": "2019-06-26T13:31:31.979Z"
}
```


This is a reference event for VM Cluster Network - Remove DB Server Network End:


```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrupbmw5ksic23bmwnqc4rsskxwc23eslydoumcnr6rcu
m6mtmgg3a",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.removedbservervmclusternetwork.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",
      "compartmentName": "example_name",
      "resourceName": "my_vmcluster_network",
      "resourceId": "VmClusterNetwork-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "ALLOCATED",
        "timeCreated": "2019-08-29T12:00:00.000Z",
        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "displayName": "testDisplayName",
        "exadataInfrastructureId": "ExadataInfra-unique_ID"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "VM Cluster Network - Remove DB Server Network End",
  "eventType":
"com.oraclecloud.databaseservice.removedbservervmclusternetwork.end",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "timeUpdated", "type": "string" },
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    { "name": "displayName", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "lifecycleDetails", "type": ["null", "string"] },
    { "name": "exadataInfrastructureId", "type": ["null", "string"] }
  ],
  "timeCreated": "2019-06-26T13:31:31.979Z"
}
```

Serial Console Connection Event Types

Review the list of event types that serial console connection emits.

Table 7-23 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 7-77 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",
```


Copyright © 2008, 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 - 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"
  },
```


Copyright © 2008, 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": "DELETING",
      "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 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",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "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"
  },
  "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",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
      "vmClusterId": "VmCluster-unique_ID",
      "dbHostId": "dbHost-unique_ID",
      "nodeNumber": 2,
      "powerAction": "HardReset",
      "hostName": "testHostName"
    }
  }
}
```

- Viewing Audit Log Events
  Oracle Cloud Infrastructure Audit service provides records of API operations performed against supported services as a list of log events.

Viewing Audit Log Events

Oracle Cloud Infrastructure Audit service provides records of API operations performed against supported services as a list of log events.

An audit event is generated when you connect to the serial console using a Secure Shell (SSH) connection. Navigate to Audit in the Console and search for `VmConsoleConnected`. When you navigate to Audit in the Console, a list of results is generated for the current compartment. Audit logs are organized by compartment, so if you are looking for a particular event, you must know which compartment the event occurred in. You can filter the list in the following ways:

- Date and time
- Request Action Types (operations)
- Keywords

For more information, see *Viewing Audit Log Events*.

Example 7-78 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,
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    "identity": {
      "principalName": "dsaes",
      "principalId": "ocid1.user.oc1..<TRUNCATED>aaaaaaaaaxxxxxxxxxx",
      "authType": "Native",
      "callerName": null,
      "callerId": null,
      "tenantId": "ocid1.tenancy.oc1..<TRUNCATED>aaaaaaxxxxx",
      "ipAddress": null,
      "credentials": null,
      "userAgent": null,
      "consoleSessionId": null
    },

    "request": {
      "id": "",
      "path": "",
      "action": "",
      "parameters": null,
      "headers": null
   },

    "response": {
      "status": "",
      "responseTime": "0001-01-01T00:00:00.000Z",
      "headers": null,
      "payload": null,
      "message": ""
    },

    "stateChange": null,

    "additionalDetails": {
      "DBNodeId": "ocid1.dbnode.oc1.iad.<TRUNCATED>aaaaaxxxxxxx"
    }
  }
}
```


Related Topics

- Overview of Audit
- Viewing Audit Log Events
- Setting Audit Log Retention Period

Serial Console History Event Types

Review the list of new event types that serial console history emits.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-24 Serial Console History Events

User Action	Event Type	Friendly Name	Event Type
Create Console History	Async	DB Node Console History - Create Begin	com.oraclecloud.databaseservice.createdbnodeconsolehistory.begin
		DB Node Console History - Create End	com.oraclecloud.databaseservice.createdbnodeconsolehistory.end
Terminate Console History	Async	DB Node Console History - Delete Begin	com.oraclecloud.databaseservice.deletedbnodeconsolehistory.begin
		DB Node Console History - Delete End	com.oraclecloud.databaseservice.deletedbnodeconsolehistory.end
Update Console History	Sync	DB Node Console History - Update	com.oraclecloud.databaseservice.updatedbnodeconsolehistory
Get Console History Content	Sync	DB Node Console History - Get Content	com.oraclecloud.databaseservice.getdbnodeconsolehistorycontent

Example 7-79 Serial Console History Event Types Examples

This is a reference event for DB Node Console History - Create Begin:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsolehistory.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.dbnodeconsolehistory.oc1..unique_ID",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbnodeconsolehistory.oc1..unique_ID",
        "displayName": "console-history-20230808-0012",
        "lifecycleState": "REQUESTED",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "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 History - Create End:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsolehistory.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.dbnodeconsolehistory.oc1..unique_ID",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbnodeconsolehistory.oc1..unique_ID",
        "displayName": "console-history-20230808-0012",
        "lifecycleState": "SUCCEEDED",
        "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 History - Delete Begin:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsolehistory.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2019-08-29T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "resourceId": "ocid1.dbnodeconsolehistory.oc1..unique_ID",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbnodeconsolehistory.oc1..unique_ID",
        "displayName": "console-history-20230808-0012",
        "lifecycleState": "DELETING",
        "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 History - Delete End:


```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsolehistory.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.dbnodeconsolehistory.oc1..unique_ID",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbnodeconsolehistory.oc1..unique_ID",
        "displayName": "console-history-20230808-0012",
        "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 History - Update:


```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    "eventType": "com.oraclecloud.databaseservice.updatedbnodeconsolehistory",
    "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.dbnodeconsolehistory.oc1..unique_ID",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbnodeconsolehistory.oc1..unique_ID",
        "displayName": "console-history-20230808-0012",
        "lifecycleState": "SUCCEEDED",
        "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 History - Get Content:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.getdbnodeconsolehistorycontent",
    "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.dbnodeconsolehistory.oc1..unique_ID",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dbnodeconsolehistory.oc1..unique_ID",
        "displayName": "console-history-20230808-0012",
        "lifecycleState": "SUCCEEDED",
        "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 © 2008, 2026, Oracle and/or its affiliates.

```
    }
  }
```

Interim Software Updates Event Types

Table 7-25 Interim Software Updates Events

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

Example 7-80 One-Off Patches Event Types Examples

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",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Create Begin",
  "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.begin",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```


This is a reference event for Oneoff Patch - Create End:


```
{
  "id":
"ocid1.eventschema.oc1.phx.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"
      }
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Create End",
  "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.end",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```


This is a reference event for Oneoff Patch - Delete Begin:


```
{
  "id":
"ocid1.eventschema.oc1.phx.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",
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.begin",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

This is a reference event for Oneoff Patch - Delete End:

```
{
  "id":
"ocid1.eventschema.oc1.phx.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" },
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

This is a reference event for Oneoff Patch - Download Begin:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljr3vkb7klt5hkbsnqzjaxmszsqomanlbqmr2tsrcq7xaf
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" },
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

This is a reference event for Oneoff Patch - Download End:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrn2lruez55ah56kqksi5qfg6m7igvven7o2qkahlk5tk
wrj5ll3oa",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Download End",
  "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.end",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

Policy Details for Oracle Exadata Database Service on Cloud@Customer

Learn to write policies to control access to Oracle Exadata Database Service on Cloud@Customer 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 Cloud@Customer instances".

- About Resource-Types
  Learn about resource-types you can use in your policies.
- Resource-Types for Oracle Exadata Database Service on Cloud@Customer
  Review the list of resource-types specific to Oracle Exadata Database Service on Cloud@Customer.
- 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.
- Permissions Required for Each API Operation
  Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.

Related Topics

- How Policies Work
- Let database admins manage Exadata Database Service on Cloud@Customer instances

About Resource-Types

Learn about resource-types you can use in your policies.

An aggregate resource-type covers the list of individual resource-types that directly follow.

For example, writing one policy to allow a group to have access to the `database-family` is equivalent to writing eight separate policies for the group that would grant access to the `exadata-infrastructures`, `vmcluster-networks`, `vmclusters`, `backup-destinations`, `db-nodes`, `dbnode-console-connection`, and the rest of the individual resource-types.

For example, writing one policy to allow a group to have access to the `autonomous-database-family` is equivalent to writing four separate policies for the group that would grant access to the `autonomous-databases`, `autonomous-backups`, `autonomous-container-databases`, and `cloud-autonomous-vmclusters` resource-types.

For more information, see Resource-Types.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Resource-Types for Oracle Exadata Database Service on Cloud@Customer

Review the list of resource-types specific to Oracle Exadata Database Service on Cloud@Customer.

Aggregate Resource-Type

```
database-family
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Individual Resource-Types

```
exadata-infrastructures

vmclusters

backup-destinations

db-nodes

db-homes

databases

backups

database-software-images

exascale-db-storage-vaults

keystores

dbnode-console-connection

dbnode-console-history

scheduling-policies

scheduling-windows

scheduling-plan

scheduling-action

execution-windows

execution-action
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Supported Variables

Use variables when adding conditions to a policy.

Exadata Database Service on Cloud@Customer 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 Exadata Infrastructures
- Permissions and API operation details for exascale-db-storage-vaults
- Permissions and API operation details for VM Cluster Networks
- Permissions and API operation details for VM Clusters
- Permissions and API operation details for Backup Destinations
- Permissions and API operation details for DB Nodes
- Permissions and API operation details for DB Homes
- Permissions and API operation details for Databases
- Permissions and API operation details for Backups
- Permissions and API operation details for Database Software Image
- Permissions and API operation details for Key Stores
- Permissions and API operation details for Pluggable Databases (PDBs)
- Permissions and API operation details for DB Servers
- Permissions and API operation details for DB Node Console Connection
- Permissions and API operation details for DB Node Console History
- Permissions and API operation details for Interim Software Updates
- Permissions and API operation details for Scheduling Policies
- Permissions and API operation details for Scheduling Windows
- Permissions and API operation details for Scheduling Plan
- Permissions and API operation details for Scheduled Action
- Permissions and API operation details for Execution Windows
- Permissions and API operation details for Execution Action
- Permissions and API operation details for Data Guard Group

Copyright © 2008, 2026, Oracle and/or its affiliates.

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
    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)

Copyright © 2008, 2026, Oracle and/or its affiliates.

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 Exadata Infrastructures

Granting permissions on `exadata-infrastructure` resources grants permissions on associated `vmcluster-network` resources.

The table below lists permissions and API operations for `exadata-infrastructures`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	EXADATA_INFRASTRUCTURE_INSPECT	ListExadataInfrastructures GetExadataInfrastructure GenerateRecommendedNetworkDetails ListVmClusterNetworks GetVmClusterNetwork ValidateVmClusterNetwork DownloadExadataInfrastructureConfigFile DownloadVmClusterNetworkConfigFile	ChangeExadataInfrastructureCompartment
READ	*INSPECT* + EXADATA_INFRASTRUCTURE_CONTENT_READ	*none*	*none*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

USE	*READ* + `EXADATA_INFRASTRUCTURE_UPDATE`	`ActivateExadataInfrastructure`
		`UpdateExadataInfrastructure`
		`ChangeExadataInfrastructureCompartment`
		`AddStorageCapacityExadataInfrastructure`
		`CreateVmClusterNetwork`
		`UpdateVmClusterNetwork`
		`DeleteVmClusterNetwork`
		`ConfigureExascaleExadataInfrastructure`
			`CreateVmCluster` (also needs `manage vmclusters`)
			`UpdateVmCluster` (also needs `use vmclusters`)
			`ChangeExadataInfrastructureCompartment`

MANAGE	*USE* + `EXADATA_INFRASTRUCTURE_CREATE`
	`EXADATA_INFRASTRUCTURE_DELETE`
		`CreateExadataInfrastructure`
		`DeleteExadataInfrastructure`
		`downloadExadataInfrastructureConfigFile`
			none

Permissions and API operation details for exascale-db-storage-vaults

Granting permissions on `exadata-infrastructure` resources grants permissions on associated `vmcluster-network` resources.

The table below lists permissions and API operations for `exadata-infrastructures`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`EXASCALE_DB_STORAGE_VAULT_INSPECT`	`ListExascaleDbStorageVaults` `GetExascaleDbStorageVault`	*none*
READ	*no extra*	*no extra*	*none*
USE	*READ* + `EXASCALE_DB_STORAGE_VAULT_UPDATE`	`ChangeExascaleDbStorageVaultCompartment` `UpdateExascaleDbStorageVault`	*none*
MANAGE	*USE* + `EXASCALE_DB_STORAGE_VAULT_CREATE` `EXASCALE_DB_STORAGE_VAULT_DELETE`	`CreateExascaleDbStorageVault` `DeleteExascaleDbStorageVault`	*none*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Permissions and API operation details for VM Cluster Networks

`vmcluster-network` resources inherit permissions from the `exadata-infrastructure` resources with which they are associated. You cannot grant permissions to `vmcluster-network` resources explicitly.

The table below lists permissions and API operations for `vmcluster-networks`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

INSPECT	`EXADATA_INFRASTRUCTURE_INSPECT`	`ListVmClusterNetworks`
		`GetVmClusterNetwork`
		`ValidateVmClusterNetwork`
			none

READ	*INSPECT* + `EXADATA_INFRASTRUCTURE_CONTENT_READ`	`DownloadVmClusterNetworkConfigFile`	*none*

USE	*READ* +
	`EXADATA_INFRASTRUCTURE_UPDATE`
		`CreateVmClusterNetwork`
		`UpdateVmClusterNetwork`
		`DeleteVmClusterNetwork`
			none

MANAGE	*USE* +
	`EXADATA_INFRASTRUCTURE_CREATE`
	`EXADATA_INFRASTRUCTURE_DELETE`
		none	*none*

Permissions and API operation details for VM Clusters

The table below lists permissions and API operations for `vmclusters`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	VM_CLUSTER_INSPECT	ListVmClusters GetVmCluster ListVmClusterPatches ListVmClusterPatchHistoryEntries GetVmClusterPatch GetVmClusterPatchHistoryEntry ListVmClusterUpdates ListVmClusterUpdateHistoryEntries GetVmClusterUpdate GetVmClusterUpdateHistoryEntry	*none*
READ	*No extra*	*No extra*	*No extra*
USE	*READ* + VM_CLUSTER_UPDATE VM_CLUSTER_UPDATE_TAGS VM_CLUSTER_UPDATE_COMPARTMENT VM_CLUSTER_UPDATE_SSH_KEY VM_CLUSTER_UPDATE_LICENSE VM_CLUSTER_UPDATE_CPU VM_CLUSTER_UPDATE_MEMORY VM_CLUSTER_UPDATE_LOCAL_STORAGE VM_CLUSTER_UPDATE_EXADATA_STORAGE VM_CLUSTER_UPDATE_GI_SOFTWARE VM_CLUSTER_UPDATE_FILE_SYSTEM VM_CLUSTER_UPDATE_CLOUD_AUTOMATION_PREFERENCES	ChangeVmClusterCompartment	UpdateVmCluster (also needs use exadata-infrastructures) CreateDbHome, (also needs manage db-homes and manage databases). If automatic backups are enabled on the default database, also needs manage backups DeleteDbHome, (also needs manage db-homes and manage databases. If automatic backups are enabled on the default database, also needs manage backups
MANAGE	*USE* + VM_CLUSTER_CREATE VM_CLUSTER_DELETE	*No extra*	CreateVmCluster (also needs use exadata-infrastructures) DeleteVmCluster (also needs use exadata-infrastructures)

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> The `VM_CLUSTER_UPDATE_SSH_KEY` permission is a highly privileged permission that allows the user to be a root user on the guest VM and gives them the ability to run other cluster update operations on the guest VM using `dbaascli`.

Using fine-grained permissions, you can write policies as follows:

- To allow any update operations:

```
allow group abc to use vmclusters in compartment comp1
```

- To allow only scale CPU:

```
allow group abc to use vmclusters in compartment comp1 where
request.permission = 'VM_CLUSTER_UPDATE_CPU'
```

- To allow GI update and any scale operations:

```
allow group abc to use vmclusters in compartment comp1
where any
     { request.permission = 'VM_CLUSTER_UPDATE_CPU', request.permission =
'VM_CLUSTER_UPDATE_EXADATA_STORAGE',
         request.permission = 'VM_CLUSTER_UPDATE_MEMORY',
request.permission = 'VM_CLUSTER_UPDATE_LOCAL_STORAGE', request.permission
= 'VM_CLUSTER_UPDATE_GI_SOFTWARE'}
```

- To allow any operations except add SSH key:

```
allow group abc to use vmclusters in compartment comp1 where all
{ request.permission != 'VM_CLUSTER_UPDATE_SSH_KEY' , request.permission !
= 'VM_CLUSTER_UPDATE' }
```

Permissions and API operation details for Backup Destinations

The table below lists permissions and API operations for `backup-destinations`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`BACKUP_DESTINATION_INSPECT`	`ListBackupDestinations` `GetBackupDestination`	*none*
READ	*no extra*	*none*	*none*
USE	*READ* + `BACKUP_DESTINATION_UPDATE`	`UpdateBackupDestination` `ChangeBackupDestinationCompartment`	*none*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
MANAGE	*USE* + `BACKUP_DESTINATION_CREATE` `BACKUP_DESTINATION_DELETE`	`CreateBackupDestination` `DeleteBackupDestination`	*none*

Permissions and API operation details for DB Nodes

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	*READ* + `DB_NODE_UPDATE`	`UpdateDbNode`	*none*
MANAGE	*USE* + `DB_NODE_POWER_ACTIONS`	`DbNodeAction`	*none*

Permissions and API operation details for DB Homes

The table below lists permissions and API operations for `db-homes`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`DB_HOME_INSPECT`	`ListDBHome` `GetDBHome` `ListDbHomePatches` `ListDbHomePatchHistoryEntries` `GetDbHomePatch` `GetDbHomePatchHistoryEntry`	*none*
READ	*No extra*	*No extra*	*none*
USE	`DB_HOME_UPDATE`	`UpdateDBHome`	*none*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
MANAGE	*USE* + DB_HOME_CREATE DB_HOME_DELETE	*No extra*	CreateDbHome, (also needs manage db-homes, manage backups, manage db-nodes, read vmclusters, read work-requests, and inspect exadata-infrastructures). If automatic backups are enabled on the default database, also needs manage backups. DeleteDbHome, (also needs manage db-homes, manage backups, manage db-nodes, read vmclusters, read work-requests, and inspect exadata-infrastructures). If automatic backups are enabled on the default database, also needs manage backups.

Permissions and API operation details for Databases

The table below lists permissions and API operations for databases.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	DATABASE_INSPECT	ListDatabases GetDatabase getDatabaseUpgradeHistoryEntry ListDataGuardAssociations GetDataGuardAssociation	ListDatabaseUpgradeHistoryEntries UpgradeDatabase
READ	*No extra*	*No extra*	*none*
USE	*READ* + DATABASE_UPDATE DB_HOME_UPDATE	UpdateDatabase SwitchoverDataGuardAssociation FailoverDataGuardAssociation ReinstateDataGuardAssociation	If enabling automatic backups, also needs manage backups.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

MANAGE	*USE* +
	`DATABASE_CREATE`
	`DATABASE_DELETE`
		No extra	`CreateDbHome`, (also needs `use vmclusters` and `manage db-homes`). If automatic backups are enabled on the default database, also needs `manage backups`
			`DeleteDbHome`, (also needs `use vmclusters` and `manage db-homes`). If automatic backups are enabled on the default database, also needs `manage backups`

Permissions and API operation details for Backups

The table below lists permissions and API operations for `backups`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`DB_BACKUP_INSPECT`	`GetBackup` `ListBackups`	*none*
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`
		no extra	*none*

Permissions and API operation details for Database Software Image

The table below lists permissions and API operations for `database-software-image`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`DB_SOFTWARE_IMG_INSPECT`	`ListDatabaseSoftwareImages` `GetDatabaseSoftwareImage`	*none*
READ	*No extra*	*No extra*	*none*
USE	*READ* + `DB_SOFTWARE_IMG_UPDATE`	`UpdateDatabaseSoftwareImage` `ChangeDatabaseSoftwareImageCompartment`	*none*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
MANAGE	*USE* + DB_SOFTWARE_IMG_CREATE + DB_SOFTWARE_IMG_DELETE	CreateDatabaseSoftwareImage DeleteDatabaseSoftwareImage	*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*

Permissions and API operation details for Pluggable Databases (PDBs)

The table below lists permissions and API operations for `pluggable-databases`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	PLUGGABLE_DATABASE_INSPECT	ListPluggableDatabases GetPluggableDatabase	UpdatePluggableDatabase StartPluggableDatabase StopPluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase RefreshPluggableDatabase ConvertRefreshablePluggableDatabase

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
	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
	PLUGGABLE_DATABASE_UPDATE	*no extra*	UpdatePluggableDatabase StartPluggableDatabase StopPluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase RefreshPluggableDatabase ConvertRefreshablePluggableDatabase
	DATABASE_UPDATE	*no extra*	CreatePluggableDatabase DeletePluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

manage	*USE* +
	`PLUGGABLE_DATABASE_CREATE`
		no extra	`CreatePluggableDatabase`
			`LocalClonePluggableDatabase`
			`RemoteClonePluggableDatabase`

	`PLUGGABLE_DATABASE_DELETE`	*no extra*	`DeletePluggableDatabase`

Permissions and API operation details for DB Servers

The table below lists permissions and API operations for `dbServers`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	`EXADATA_INFRASTRUCTURE_INSPECT`	*none*	`GetDbServer` `ListDbServers`
READ	*no extra*	*no extra*	*none*

USE	*READ* +
	`VM_CLUSTER_UPDATE`
	`EXADATA_INFRASTRUCTURE_UPDATE`
		none	`AddVirtualMachineToVmCluster`, `RemoveVirtualMachineFromVmCluster`

MANAGE	*No extra*	*No extra*	*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

Copyright © 2008, 2026, Oracle and/or its affiliates.

Permissions and API operation details for DB Node Console History

The table below lists permissions and API operations for `dbnode-console-history`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	`DBNODE_CONSOLE_HISTORY_INSPECT`	`GetDbNodeConsoleHistory` `ListDbNodeConsoleHistories`	*none*
read	*INSPECT* + `DBNODE_CONSOLE_HISTORY_CONTENT_READ`	`getDbNodeConsoleHistoryContent`	*none*

use	*READ* +
	`DBNODE_CONSOLE_HISTORY_UPDATE`
	`PLUGGABLE_DATABASE_UPDATE`
		`UpdateDbNodeConsoleHistory`	*none*

manage	*USE* +
	`DBNODE_CONSOLE_HISTORY_CREATE`
	`DBNODE_CONSOLE_HISTORY_DELETE`
		`CreateDbNodeConsoleHistory`
		`DeleteDbNodeConsoleHistory`
			none

Permissions and API operation details for Interim Software Updates

The table below lists permissions and API operations for `oneoffPatch`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

inspect	`ONEOFF_PATCH_INSPECT`	`DownloadOneoffPatch`
		`GetOneoffPatch`
		`ListOneoffPatches`
			`CreateOneoffPatch`
			`DeleteOneoffPatch`
			`UpdateOneoffPatch`
			`ChangeOneoffPatchCompartment`

read	*INSPECT* + no extra	`DownloadOneoffPatch`	*none*
use	*READ* + `ONEOFF_PATCH_UPDATE`	*none*	`UpdateOneoffPatch` `ChangeOneoffPatchCompartment`

manage	*USE* +
	`ONEOFF_PATCH_CREATE`
	`ONEOFF_PATCH_DELETE`
		none	`CreateOneoffPatch`
			`DeleteOneoffPatch`

Related Topics

- OneoffPatch Reference

Copyright © 2008, 2026, Oracle and/or its affiliates.

Permissions and API operation details for Scheduling Policies

The table below lists permissions and API operations for `scheduling-policies`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

inspect	`SCHEDULING_POLICY_INSPECT`	`ListSchedulingPolicies`
		`GetSchedulingPolicy`
		`ListRecommendedScheduledActions`
			`CreateSchedulingPolicy`
			`UpdateSchedulingPolicy`
			`DeleteSchedulingPolicy`
			`ChangeSchedulingPolicyCompartment`

read	*INSPECT + No extra*	*No extra*	*none*
use	*READ +* `SCHEDULING_POLICY_UPDATE`	*No extra*	`UpdateSchedulingPolicy` `ChangeSchedulingPolicyCompartment`

manage	*USE +*
	`SCHEDULING_POLICY_CREATE`
	`SCHEDULING_POLICY_DELETE`
		No extra	`CreateSchedulingPolicy`
			`DeleteSchedulingPolicy`

Permissions and API operation details for Scheduling Windows

The table below lists permissions and API operations for `scheduling-windows`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered

inspect	`SCHEDULING_WINDOW_INSPECT`	`ListSchedulingWindows`
		`GetSchedulingWindow`
			`CreateSchedulingWindow`
			`UpdateSchedulingWindow`
			`DeleteSchedulingWindow`

read	*INSPECT + No extra*	*No extra*	*none*
use	*READ +* `SCHEDULING_WINDOW_UPDATE`	*No extra*	`UpdateSchedulingWindow`

manage	*USE +*
	`SCHEDULING_WINDOW_CREATE`
	`SCHEDULING_WINDOW_DELETE`
		No extra	`CreateSchedulingWindow`
			`DeleteSchedulingWindow`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Permissions and API operation details for Scheduling Plan

`scheduling-plan` resources inherit permissions from the `exadata-infrastructure` resources with which they are associated. You cannot grant permissions to `scheduling-plan` resources explicitly.

The table below lists permissions and API operations for `scheduling-plan`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	`EXADATA_INFRASTRUCTURE_INSPECT`	• `ListExadataInfrastructures` • `GetExadataInfrastructure` • `GenerateRecommendedNetworkDetails` • `ListVmClusterNetworks` • `GetVmClusterNetwork` • `ValidateVmClusterNetwork` • `DownloadExadataInfrastructureConfigFile` • `DownloadVmClusterNetworkConfigFile` • `ListSchedulingPlans` • `GetSchedulingPlan`	`ChangeExadataInfrastructureCompartment`
read	*INSPECT + No extra*	*No extra*	*none*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
use	*READ* + EXADATA_INFRASTRUCTURE_UPDATE	• ActivateExadataInfrastructure • UpdateExadataInfrastructure • ChangeExadataInfrastructureCompartment • AddStorageCapacityExadataInfrastructure • CreateVmClusterNetwork • UpdateVmClusterNetwork • DeleteVmClusterNetwork • ChangeSchedulingPlanCompartment • CascadingDeleteSchedulingPlan	• CreateVmCluster (also needs manage vmclusters) • UpdateVmCluster (also needs use vmclusters) • ChangeExadataInfrastructureCompartment
manage	*USE* + EXADATA_INFRASTRUCTURE_CREATE EXADATA_INFRASTRUCTURE_DELETE	• CreateExadataInfrastructure • DeleteExadataInfrastructure • downloadExadataInfrastructureConfigFile • CreateSchedulingPlan • DeleteSchedulingPlan	*none*

Permissions and API operation details for Scheduled Action

`scheduled-action` resources inherit permissions from the `exadata-infrastructure` resources with which they are associated. You cannot grant permissions to `scheduled-action` resources explicitly.

The table below lists permissions and API operations for `scheduled-action`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	EXADATA_INFRASTRUCTURE_INSPECT	• ListExadataInfrastructures • GetExadataInfrastructure • GenerateRecommendedNetworkDetails • ListVmClusterNetworks • GetVmClusterNetwork • ValidateVmClusterNetwork • DownloadExadataInfrastructureConfigFile • DownloadVmClusterNetworkConfigFile • ListScheduledActions • GetScheduledAction • ListParamsForActionType	ChangeExadataInfrastructureCompartment
read	*INSPECT + No extra*	*No extra*	*none*
use	*READ* + EXADATA_INFRASTRUCTURE_UPDATE	• ActivateExadataInfrastructure • UpdateExadataInfrastructure • ChangeExadataInfrastructureCompartment • AddStorageCapacityExadataInfrastructure • CreateVmClusterNetwork • UpdateVmClusterNetwork • DeleteVmClusterNetwork • UpdateScheduledAction • ReorderScheduledActions	• CreateVmCluster (also needs manage vmclusters) • UpdateVmCluster (also needs use vmclusters) • ChangeExadataInfrastructureCompartment

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	*USE* + `EXADATA_INFRASTRUCTURE_CREATE` `EXADATA_INFRASTRUCTURE_DELETE`	• `CreateExadataInfrastructure` • `DeleteExadataInfrastructure` • `downloadExadataInfrastructureConfigFile` • `CreateScheduledAction` • `DeleteScheduledAction`	*none*

Permissions and API operation details for Execution Windows

`execution-windows` resources inherit permissions from the `exadata-infrastructure` resources with which they are associated. You cannot grant permissions to `execution-windows` resources explicitly.

The table below lists permissions and API operations for `execution-windows`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	`EXADATA_INFRASTRUCTURE_INSPECT`	• `ListExadataInfrastructures` • `GetExadataInfrastructure` • `GenerateRecommendedNetworkDetails` • `ListVmClusterNetworks` • `GetVmClusterNetwork` • `ValidateVmClusterNetwork` • `DownloadExadataInfrastructureConfigFile` • `DownloadVmClusterNetworkConfigFile` • `ListExecutionWindows` • `GetExecutionWindow`	`ChangeExadataInfrastructureCompartment`
read	*INSPECT + No extra*	*No extra*	*none*

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
use	*READ* + EXADATA_INFRASTRUCTURE_UPDATE	• ActivateExadataInfrastructure • UpdateExadataInfrastructure • ChangeExadataInfrastructureCompartment • AddStorageCapacityExadataInfrastructure • CreateVmClusterNetwork • UpdateVmClusterNetwork • DeleteVmClusterNetwork • UpdateExecutionWindow • CancelExecutionWindow	• CreateVmCluster (also needs manage vmclusters) • UpdateVmCluster (also needs use vmclusters) • ChangeExadataInfrastructureCompartment
manage	*USE* + EXADATA_INFRASTRUCTURE_CREATE EXADATA_INFRASTRUCTURE_DELETE	• CreateExadataInfrastructure • DeleteExadataInfrastructure • downloadExadataInfrastructureConfigFile • CreateExecutionWindow • DeleteExecutionWindow	*none*

Permissions and API operation details for Execution Action

`execution-action` resources inherit permissions from the `exadata-infrastructure` resources with which they are associated. You cannot grant permissions to `execution-action` resources explicitly.

The table below lists permissions and API operations for `execution-action`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	EXADATA_INFRASTRUCTURE_INSPECT	• ListExadataInfrastructures • GetExadataInfrastructure • GenerateRecommendedNetworkDetails • ListVmClusterNetworks • GetVmClusterNetwork • ValidateVmClusterNetwork • DownloadExadataInfrastructureConfigFile • DownloadVmClusterNetworkConfigFile • ListExecutionActions • GetExecutionAction	ChangeExadataInfrastructureCompartment
read	*INSPECT + No extra*	*No extra*	*none*
use	*READ* + EXADATA_INFRASTRUCTURE_UPDATE	• ActivateExadataInfrastructure • UpdateExadataInfrastructure • ChangeExadataInfrastructureCompartment • AddStorageCapacityExadataInfrastructure • CreateVmClusterNetwork • UpdateVmClusterNetwork • DeleteVmClusterNetwork • UpdateExecutionAction • MoveExecutionActionMember	• CreateVmCluster (also needs manage vmclusters) • UpdateVmCluster (also needs use vmclusters) • ChangeExadataInfrastructureCompartment

Copyright © 2008, 2026, Oracle and/or its affiliates.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	*USE* + `EXADATA_INFRASTRUCTURE_CREATE` `EXADATA_INFRASTRUCTURE_DELETE`	• `CreateExadataInfrastructure` • `DeleteExadataInfrastructure` • `downloadExadataInfrastructureConfigFile` • `CreateExecutionAction` • `DeleteExecutionAction`	*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`	• `ListDatabases` • `GetDatabase`	`CreateDatabase`
read	*No extra*	*No extra*	*none*
use	READ+ Standby: `VM_CLUSTER_INSPECT` + `DB_HOME_INSPECT` + `DATABASE_CREATE` Primary: `DATABASE_INSPECT` + `DATABASE_CONTENT_READ` + `DATABASE_UPDATE`	`CreateDatabase` (standby) `DeleteDatabase`	`DataguardAction` (Switchover, Failover, Reinstate, `UpdateDatabase`) Needs only `DATABASE_INSPECT` + `DATABASE_UPDATE`
manage	USE + `DATABASE_DELETE`	`DeleteDatabase`	*none*

Permissions Required for Each API Operation

Review the list of API operations for Oracle Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.

For information about permissions, see *Permissions*.

Table 7-26 Database API Operations

API Operation	Permissions Required to Use the Operation
`ListExadataInfrastructures`	`EXADATA_INFRASTRUCTURE_INSPECT`
`GetExadataInfrastructure`	`EXADATA_INFRASTRUCTURE_INSPECT`
`CreateExadataInfrastructure`	`EXADATA_INFRASTRUCTURE_CREATE`
`UpdateExadataInfrastructure`	`EXADATA_INFRASTRUCTURE_UPDATE`

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-26 (Cont.) Database API Operations

API Operation	Permissions Required to Use the Operation
ChangeExadataInfrastructureCompartment	EXADATA_INFRASTRUCTURE_INSPECT and EXADATA_INFRASTRUCTURE_UPDATE
DeleteExadataInfrastructure	EXADATA_INFRASTRUCTURE_DELETE
DownloadExadataInfrastructureConfigFile	EXADATA_INFRASTRUCTURE_CONTENT_READ
AddStorageCapacityExadataInfrastructure	EXADATA_INFRASTRUCTURE_UPDATE
ActivateExadataInfrastructure	EXADATA_INFRASTRUCTURE_UPDATE
GenerateRecommendedNetworkDetails	EXADATA_INFRASTRUCTURE_INSPECT
ListVmClusterNetworks	EXADATA_INFRASTRUCTURE_INSPECT
GetVmClusterNetwork	EXADATA_INFRASTRUCTURE_INSPECT
CreateVmClusterNetwork	EXADATA_INFRASTRUCTURE_INSPECT and EXADATA_INFRASTRUCTURE_UPDATE
UpdateVmClusterNetwork	EXADATA_INFRASTRUCTURE_INSPECT and EXADATA_INFRASTRUCTURE_UPDATE
DeleteVmClusterNetwork	EXADATA_INFRASTRUCTURE_UPDATE
DownloadVmClusterNetworkConfigFile	EXADATA_INFRASTRUCTURE_INSPECT and EXADATA_INFRASTRUCTURE_CONTENT_READ
ValidateVmClusterNetwork	EXADATA_INFRASTRUCTURE_INSPECT
ListVmClusters	VM_CLUSTER_INSPECT
GetVmCluster	VM_CLUSTER_INSPECT
CreateVmCluster	EXADATA_INFRASTRUCTURE_INSPECT and EXADATA_INFRASTRUCTURE_UPDATE and VM_CLUSTER_CREATE
UpdateVmCluster	EXADATA_INFRASTRUCTURE_INSPECT and EXADATA_INFRASTRUCTURE_UPDATE and VM_CLUSTER_UPDATE
ChangeVmClusterCompartment	VM_CLUSTER_INSPECT and VM_CLUSTER_UPDATE
DeleteVmCluster	VM_CLUSTER_DELETE
ListVmClusterPatches	VM_CLUSTER_INSPECT
ListVmClusterPatchHistoryEntries	VM_CLUSTER_INSPECT
GetVmClusterPatch	VM_CLUSTER_INSPECT
GetVmClusterPatchHistoryEntry	VM_CLUSTER_INSPECT
ListVmClusterUpdates	VM_CLUSTER_INSPECT
ListVmClusterUpdateHistoryEntries	VM_CLUSTER_INSPECT
GetVmClusterUpdate	VM_CLUSTER_INSPECT
GetVmClusterUpdateHistoryEntry	VM_CLUSTER_INSPECT
ListBackupDestination	BACKUP_DESTINATION_INSPECT
GetBackupDestination	BACKUP_DESTINATION_INSPECT
CreateBackupDestination	BACKUP_DESTINATION_CREATE
UpdateBackupDestination	BACKUP_DESTINATION_UPDATE

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-26 (Cont.) Database API Operations

API Operation	Permissions Required to Use the Operation
DeleteBackupDestination	BACKUP_DESTINATION_DELETE
ChangeBackupDestinationCompartment	BACKUP_DESTINATION_INSPECT and BACKUP_DESTINATION_UPDATE
GetDbNode	DB_NODE_INSPECT
DbNodeAction	DB_NODE_POWER_ACTIONS
ListDbHomes	DB_HOME_INSPECT
GetDbHome	DB_HOME_INSPECT
CreateDbHome	VM_CLUSTER_INSPECT and VM_CLUSTER_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	VM_CLUSTER_UPDATE and DB_HOME_UPDATE and DATABASE_DELETE
ListDbHomePatches	DB_HOME_INSPECT
ListDbHomePatchHistoryEntries	DB_HOME_INSPECT
GetDbHomePatch	DB_HOME_INSPECT
GetDbHomePatchHistoryEntry	DB_HOME_INSPECT
CreateDatabase	VM_CLUSTER_INSPECT, VM_CLUSTER_UPDATE, DB_HOME_INSPECT, DB_HOME_UPDATE, DATABASE_CREATE DB_BACKUP_CREATE and DATABASE_CONTENT_READ DB_BACKUP_INSPECT, DB_BACKUP_CONTENT_READ
ListDatabases	DATABASE_INSPECT
GetDatabase	DATABASE_INSPECT
UpdateDatabase	DATABASE_UPDATE To enable automatic backups, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
DeleteDatabase	VM_CLUSTER_UPDATE, DB_HOME_UPDATE, DATABASE_DELETE DB_BACKUP_INSPECT, DB_BACKUP_DELETE DB_BACKUP_CREATE and DATABASE_CONTENT_READ
UpgradeDatabase	DATABASE_INSPECT DATABASE_UPDATE DB_HOME_INSPECT DB_HOME_UPDATE
getDatabaseUpgradeHistoryEntry	DATABASE_INSPECT

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-26 (Cont.) Database API Operations

API Operation	Permissions Required to Use the Operation
getDatabaseUpgradeHistoryEntry	DATABASE_INSPECT
ListDbVersions	(no permissions required; available to anyone)
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
CreateDataGuardAssociation	VM_CLUSTER_INSPECT and DATABASE_INSPECT and DATABASE_UPDATE
GetDataGuardAssociation	DATABASE_INSPECT
ListDataGuardAssociations	DATABASE_INSPECT
SwitchoverDataGuardAssociation	DATABASE_UPDATE
FailoverDataGuardAssociation	DATABASE_UPDATE
ReinstateDataGuardAssociation	DATABASE_UPDATE
DeleteDatabase	VM_CLUSTER_UPDATE and DB_HOME_UPDATE and DATABASE_DELETE
CreateKeyStore	KEY_STORE_CREATE
GetKeyStore	KEY_STORE_INSPECT
UpdateKeyStore	KEY_STORE_UPDATE
DeleteKeyStore	KEY_STORE_DELETE
ChangeKeyStoreCompartment	KEY_STORE_INPSECT and KEY_STORE_UPDATE
ListDatabaseSoftwareImages	DB_SOFTWARE_IMG_INSPECT
GetDatabaseSoftwareImage	DB_SOFTWARE_IMG_INSPECT
UpdateDatabaseSoftwareImage	DB_SOFTWARE_IMG_INSPECT and DB_SOFTWARE_IMG_UPDATE
ChangeDatabaseSoftwareImageCompartment	DB_SOFTWARE_IMG_INSPECT and DB_SOFTWARE_IMG_UPDATE
CreateDatabaseSoftwareImage	DB_SOFTWARE_IMG_INSPECT and DB_SOFTWARE_IMG_CREATE
DeleteDatabaseSoftwareImage	DB_SOFTWARE_IMG_INSPECT and DB_SOFTWARE_IMG_DELETE
ListPluggableDatabase	PLUGGABLE_DATABASE_INSPECT
GetPluggableDatabase	PLUGGABLE_DATABASE_INSPECT
CreatePluggableDatabase	PLUGGABLE_DATABASE_CREATE, DATABASE_INSPECT and DATABASE_UPDATE.
UpdatePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-26 (Cont.) Database API Operations

API Operation	Permissions Required to Use the Operation
StartPluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
StopPluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
DeletePluggableDatabase	PLUGGABLE_DATABASE_DELETE, DATABASE_INSPECT, and DATABASE_UPDATE
LocalClonePluggableDatabase	PLUGGABLE_DATABASE_INSPECT, PLUGGABLE_DATABASE_UPDATE, PLUGGABLE_DATABASE_CONTENT_READ, PLUGGABLE_DATABASE_CONTENT_WRITE, PLUGGABLE_DATABASE_CREATE, DATABASE_INSPECT, and DATABASE_UPDATE
RemoteClonePluggableDatabase	PLUGGABLE_DATABASE_INSPECT, PLUGGABLE_DATABASE_UPDATE, PLUGGABLE_DATABASE_CONTENT_READ, PLUGGABLE_DATABASE_CONTENT_WRITE, PLUGGABLE_DATABASE_CREATE, DATABASE_INSPECT, and DATABASE_UPDATE
RefreshPluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
ConvertRefreshablePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
GetDbServer	DB_SERVER_INSPECT
ListDbServers	DB_SERVER_INSPECT
AddVirtualMachineToVmCluster	VM_CLUSTER_UPDATE EXADATA_INFRASTRUCTURE_UPDATE
RemoveVirtualMachineFromVmCluster	VM_CLUSTER_UPDATE EXADATA_INFRASTRUCTURE_UPDATE
CreateOneoffPatch	ONEOFF_PATCH_INSPECT ONEOFF_PATCH_CREATE
DeleteOneoffPatch	ONEOFF_PATCH_INSPECT ONEOFF_PATCH_DELETE
DownloadOneoffPatch	ONEOFF_PATCH_INSPECT
UpdateOneoffPatch	ONEOFF_PATCH_INSPECT ONEOFF_PATCH_UPDATE
ListOneoffPatches	ONEOFF_PATCH_INSPECT
GetOneoffPatch	ONEOFF_PATCH_INSPECT
ChangeOneoffPatchCompartment	ONEOFF_PATCH_INSPECT ONEOFF_PATCH_UPDATE
CreateDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_CREATE DBNODE_CONSOLE_CONNECTION_INSPECT
GetDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_INSPECT

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-26 (Cont.) Database API Operations

API Operation	Permissions Required to Use the Operation
ListDbNodeConsoleConnections	DBNODE_CONSOLE_CONNECTION_INSPECT
DeleteDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_DELETE
UpdateDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_UPDATE
UpdateDbNode	DB_NODE_UPDATE
createDbNodeConsoleHistory	DBNODE_CONSOLE_HISTORY_CREATE & DBNODE_CONSOLE_HISTORY_INSPECT
getDbNodeConsoleHistory	DBNODE_CONSOLE_HISTORY_INSPECT
getDbNodeConsoleHistoryContent	DBNODE_CONSOLE_HISTORY_CONTENT_READ
listDbNodeConsoleHistories	DBNODE_CONSOLE_HISTORY_INSPECT
updateDbNodeConsoleHistory	DBNODE_CONSOLE_HISTORY_UPDATE
deleteDbNodeConsoleHistory	DBNODE_CONSOLE_HISTORY_DELETE
ListSchedulingPolicies	SCHEDULING_POLICY_INSPECT
GetSchedulingPolicy	SCHEDULING_POLICY_INSPECT
UpdateSchedulingPolicy	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_UPDATE
ChangeSchedulingPolicyCompartment	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_UPDATE
ListRecommendedScheduledActions	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_UPDATE
CreateSchedulingPolicy	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_CREATE
DeleteSchedulingPolicy	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_DELETE
ListSchedulingWindows	SCHEDULING_WINDOW_INSPECT
GetSchedulingWindow	SCHEDULING_WINDOW_INSPECT
UpdateSchedulingWindow	SCHEDULING_WINDOW_INSPECT & SCHEDULING_WINDOW_UPDATE
CreateSchedulingWindow	SCHEDULING_WINDOW_INSPECT & SCHEDULING_WINDOW_CREATE
DeleteSchedulingWindow	SCHEDULING_WINDOW_INSPECT & SCHEDULING_WINDOW_DELETE
CreateSchedulingWindow	SCHEDULING_WINDOW_INSPECT & SCHEDULING_WINDOW_CREATE
DeleteSchedulingWindow	SCHEDULING_WINDOW_INSPECT & SCHEDULING_WINDOW_DELETE
ListSchedulingPlans	EXADATA_INFRASTRUCTURE_INSPECT
GetSchedulingPlan	EXADATA_INFRASTRUCTURE_INSPECT
CreateSchedulingPlan	EXADATA_INFRASTRUCTURE_CREATE
ChangeSchedulingPlanCompartment	EXADATA_INFRASTRUCTURE_UPDATE
ReorderScheduledActions	EXADATA_INFRASTRUCTURE_UPDATE

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-26 (Cont.) Database API Operations

API Operation	Permissions Required to Use the Operation
CascadingDeleteSchedulingPlan	EXADATA_INFRASTRUCTURE_UPDATE
DeleteSchedulingPlan	EXADATA_INFRASTRUCTURE_DELETE
ListScheduledActions	EXADATA_INFRASTRUCTURE_INSPECT
GetScheduledAction	EXADATA_INFRASTRUCTURE_INSPECT
ListParamsForActionType	EXADATA_INFRASTRUCTURE_INSPECT
UpdateScheduledAction	EXADATA_INFRASTRUCTURE_UPDATE
ReorderScheduledActions	EXADATA_INFRASTRUCTURE_UPDATE
CreateScheduledAction	EXADATA_INFRASTRUCTURE_CREATE
DeleteScheduledAction	EXADATA_INFRASTRUCTURE_DELETE
ListExecutionWindows	EXADATA_INFRASTRUCTURE_INSPECT
GetExecutionWindow	EXADATA_INFRASTRUCTURE_INSPECT
UpdateExecutionWindow	EXADATA_INFRASTRUCTURE_UPDATE
ReorderExecutionActions	EXADATA_INFRASTRUCTURE_UPDATE
CancelExecutionWindow	EXADATA_INFRASTRUCTURE_UPDATE
CreateExecutionWindow	EXADATA_INFRASTRUCTURE_CREATE
DeleteExecutionWindow	EXADATA_INFRASTRUCTURE_DELETE
ListExecutionActions	EXADATA_INFRASTRUCTURE_INSPECT
GetExecutionAction	EXADATA_INFRASTRUCTURE_INSPECT
UpdateExecutionAction	EXADATA_INFRASTRUCTURE_UPDATE
MoveExecutionActionMember	EXADATA_INFRASTRUCTURE_UPDATE
CreateExecutionAction	EXADATA_INFRASTRUCTURE_CREATE
DeleteExecutionAction	EXADATA_INFRASTRUCTURE_DELETE

Related Topics

- Permissions

Managing Exadata Resources with Oracle Enterprise Manager Cloud Control

To manage and monitor Exadata Cloud and Exadata 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 and Exadata Cloud@Customer services.

Copyright © 2008, 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 Exadata Cloud and Exadata Cloud@Customer resources.

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 and Exadata Cloud@Customer services.

Enterprise Manager Cloud Control discovers Exadata Cloud and Exadata Cloud@Customer 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, Exadata Cloud and Exadata Cloud@Customer systems, along with any other targets, from a single interface
- Visualize storage and compute data
- View performance metrics of your Exadata components

Features of Enterprise Manager Cloud Control

Familiarize yourself with the features of Enterprise Manager Cloud Control to manage and monitor Exadata Cloud and Exadata Cloud@Customer resources.

Enterprise Manager Target for Exadata Cloud

The target for Oracle Cloud Infrastructure Exadata resources, which covers both Exadata Cloud and Exadata Cloud@Customer does the following:

- Automatically identifies and organizes related targets.
- Provides a high-level integration point for Enterprise Manager framework features such as incident rules, groups, notifications, and monitoring templates.

Improved Performance Monitoring

Enterprise Manager Cloud Control enhances performance monitoring in the following ways:

- Adds Exadata Storage Server and Exadata Storage Grid targets.
- Offers visualization of storage and compute performance for your Exadata Cloud and Exadata Cloud@Customer resources.
- Enables use of the same Maximum Availability Architecture (MAA) key performance indicators (KPI) developed for Oracle Exadata Database Machine.

Scripted CLI-based Discovery

Enterprise Manager Cloud Control uses scripts to discover Oracle Cloud Infrastructure Exadata resources. Scripts comb the existing hosts, clusters, ASM, databases and related targets, as well as adding the storage server targets.

Copyright © 2008, 2026, Oracle and/or its affiliates.

"Single Pane of Glass" View of On-Premises and Oracle Cloud Infrastructure Exadata Resources

Enterprise Manager Cloud Control 's use of a single Exadata target type provides a consistent Enterprise Manager experience across on-premises, Exadata Cloud, and Exadata Cloud@Customer resources. The common Exadata target menu allows you to easily navigate to, monitor and manage all of your Exadata systems.

Visualization

Enterprise Manager Cloud Control allows you to visualize the database and related targets associated with each Exadata Cloud and Exadata Cloud@Customer system.

Observability and Management for Oracle Exadata Database Service on Cloud@Customer

Database Management and Logging Analytics can now be used with Oracle Exadata Database Service on Cloud@Customer by installing and configuring the Management Agent RPM.

For documentation on installing the Management Agent, see the following documentation resources:

- Oracle Cloud Database-related prerequisite tasks (MOS Document ID 3015115.1)
- https://docs.oracle.com/en-us/iaas/management-agents/doc/install-management-agent-chapter.html

For documentation on the Database Management service, see Oracle Cloud Database-related Prerequisite Tasks.

For documentation on the Logging Analytics service, see Enable Access to Logging Analytics and Its Resources.

Security Guide for Oracle Exadata Database Service on Cloud@Customer Systems

This guide describes security for an Oracle Exadata Database Service on Cloud@Customer system. It includes information about the best practices for securing the Oracle Exadata Database Service on Cloud@Customer system.

- Security Configurations and Default Enabled Features
- Additional Procedures for Updating Security Posture

Security Configurations and Default Enabled Features

- Responsibilities
- Guiding Principles Followed for Security Configuration Defaults
- Security Features
- Guest VM Default Fixed Users

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Guest VM Default Fixed Groups
- Guest VM Default Security Settings
- Guest VM Default Processes
- Guest VM Network Security
- OS/VM backups

Responsibilities

Exadata Cloud@Customer is jointly managed by the customer and Oracle. The Exadata Cloud@Customer deployment is divided into two major areas of responsibilities:

- Customer accessible services:
  Components that the customer can access as part of their subscription to Exadata Cloud@Customer:
  - Customer accessible virtual machines (VM)
  - Customer accessible database services
- Oracle-managed infrastructure:
  - Power Distribution Units (PDUs)
  - Out-of-band (OOB) management switches » InfiniBand switches
  - Exadata Storage Servers
  - Physical Exadata database servers

Customers control and monitor access to customer services, including network access to their VMs (through layer 2 VLANs and firewalls implemented in the customer VM), authentication to access the VM, and authentication to access databases running in the VMs. Oracle controls and monitors access to Oracle-managed infrastructure components. Oracle staff are not authorized to access customer services, including customer VMs and databases.

Customers access Oracle Databases running on Exadata Cloud@Customer through a layer 2 (tagged VLAN) connection from customer equipment to the databases running in the customer VM using standard Oracle Database connection methods, such as Oracle Net on port 1521. Customers access the VM running the Oracle Databases through standard Oracle Linux methods, such as token based SSH on port 22.

Guiding Principles Followed for Security Configuration Defaults

- **Defense in Depth**
  Exadata Cloud@Customer offers a number of controls to ensure confidentiality, integrity, and accountability throughout the service.

  First, Exadata Cloud@Customer is built from the hardened operating system image provided by Exadata Database Machine. For more information, see *Overview of Oracle Exadata Database Machine Security*. This secures the core operating environment by restricting the installation image to only the required software packages, disabling unnecessary services, and implementing secure configuration parameters throughout the system.

  In addition to inheriting all the strength of Exadata Database Machine's mature platform, because Exadata Cloud@Customer provisions systems for customers, additional secure default configuration choices are implemented in the service instances. For example, all database tablespaces require transparent data encryption (TDE), strong password

Copyright © 2008, 2026, Oracle and/or its affiliates.

enforcement for initial database users and superusers, and enhanced audit and event rules.

Exadata Cloud@Customer also constitutes a complete deployment and service, so it is subjected to industry-standard external audits such as PCI, HIPPA and ISO27001. These external audit requirements impose additional value-added service features such as antivirus scanning, automated alerting for unexpected changes to the system, and daily vulnerability scans for all Oracle-managed infrastructure systems in the fleet.

- **Least Privilege**
  To ensure that processes only have access to the privileges they require, Oracle Secure Coding Standards require the paradigm of least privilege.

  Each process and daemon, must run as a normal, unprivileged user unless it can prove a requirement for a higher level of privilege. This helps contain any unforeseen issues or vulnerabilities to unprivileged user space and not compromise an entire system.

  This principle also applies to Oracle operations team members who use individual named accounts to access the Oracle Exadata Cloud@Customer infrastructure for maintenance or troubleshooting. Only when necessary will they use the audited access to higher levels of privilege to solve or resolve an issue. Most issues are resolved through automation, so we also employ least privilege by not permitting human operators to access a system unless the automation is unable to resolve the issue.
- **Auditing and Accountability**
  When required, access to the system is allowed, but all access and actions are logged and tracked for accountability.

  This ensures that both Oracle and customers know what was done on the system and when that occurred. These facts not only ensure we remain compliant with reporting needs for external audits, but also can help match up some change with a change in perceived behavior in the application.

  Auditing capabilities are provided at all infrastructure components to ensure all actions are captured. Customers also have ability to configure auditing for their database and Guest VM configuration and may choose to integrate those with other enterprise auditing systems.
- **Automating Cloud Operations**
  By eliminating manual operations required to provision, patch, maintain, troubleshoot, and configure systems, the opportunity for error is reduced and a secure configuration is ensured.

  The service is designed to be secure and by automating all provisioning, configuration, and most other operational tasks, it is possible to avoid missed configurations and ensure all necessary paths into the system are closed tightly.

Related Topics

- Overview of Oracle Exadata Database Machine Security

Security Features

- **Hardened Operating System Image**
  - Minimal package installation:
    Only the necessary packages required to run an efficient system are installed. By installing a smaller set of packages, the attack surface of the operating system is reduced and the system remains more secure.
  - Secure configuration:

Copyright © 2008, 2026, Oracle and/or its affiliates.

Many non-default configuration parameters are set during installation to enhance the security posture of the system and its content. For example, SSH is configured to only listen on certain network interfaces, sendmail is configured to only accept localhost connections, and many other similar restrictions are implemented during installation.

  - Run only necessary services:
    Any services that may be installed on the system, but not required for normal operation are disabled by default. For example, while NFS is a service often configured by customers for various application purposes, it is disabled by default as it is not required for normal database operations. Customers may choose to optionally configure services per their requirements.

- **Minimized Attack Surface**
  As part of the hardened image, attack surface is reduced by disabling services.
- **Additional Security Features Enabled (grub passwords, secure boot)**
  - Leveraging Exadata image capabilities, Exadata Cloud@Customer enjoys the features integrated into the base image such as grub passwords and secure boot in addition to many others.
  - Through customization, customers may wish to further enhance their security posture with additional configurations detailed later in this chapter.
- **Secure Access Methods**
  - Accessing database servers through SSH using strong cryptographic ciphers. Weak ciphers are disabled by default.
  - Accessing databases via encrypted Oracle Net connections. By default, our services are available using encrypted channels and a default configured Oracle Net client will use encrypted sessions.
  - Accessing diagnostics through Exadata MS web interface (https).
- **Auditing and Logging**
  By default, auditing is enabled for administrative operations and those audit records are communicated to OCI internal systems for automated review and alerting when required.
- **Deployment-Time Considerations**
  - Wallet file download contents and sensitivity: The wallet file download that is obtained by a customer to enable the deployment to occur contains sensitive information and should be handled appropriately to ensure the contents are protected. The content of the wallet file download is not needed after deployment is completed, so it should be destroyed to ensure no information leakage occurs.
  - Control Plane Server (CPS):
    * Deployment requirements for the CPS include permitting outbound HTTPS access so the CPS can connect to Oracle and enable remote administration, monitoring, and maintenance. Find more details in the *Deployment Guide*.
    * Customer responsibilities include providing physical security to the Exadata Cloud@Customer equipment in their data center. While Exadata Cloud@Customer employs many software security features, physical server compromise must be addressed by customer physical security to ensure the safety of the servers' contents.

Guest VM Default Fixed Users

Copyright © 2008, 2026, Oracle and/or its affiliates.

Several user accounts regularly manage the components of Oracle Exadata Cloud@Customer. In all Exadata Cloud@Customer machines, Oracle uses and recommends SSH based login only. No Oracle user or processes use password based authentication system.

Below described are the different kind of users created by default.

- **Default Users: No Logon Privileges**
  This user list consists of default operating system users along with some specialized users like `exawatch` and `dbmsvc`. These users should not be altered. These users cannot login to the system as all are set to `/bin/false`.
  - `exawatch`:
    The `exawatch` user is responsible for collecting and archiving system statistics on both the database servers and the storage servers.
  - `dbmsvc`:
    User is used for Management Server on the database node service in Oracle Exadata System.

  NOLOGIN Users

```
bin:x:1:1:bin:/bin:/bin/false
daemon:x:2:2:daemon:/sbin:/bin/false
adm:x:3:4:adm:/dev/null:/bin/false
mail:x:8:12:mail:/var/spool/mail:/bin/false
nobody:x:99:99:Nobody:/:/bin/false
systemd-network:x:192:192:systemd Network Management:/:/bin/false
dbus:x:81:81:System message bus:/:/bin/false
rpm:x:37:37::/var/lib/rpm:/bin/false
sshd:x:74:74:Privilege-separated SSH:/var/empty/sshd:/bin/false
rpc:x:32:32:Rpcbind Daemon:/var/lib/rpcbind:/bin/false
nscd:x:28:28:NSCD Daemon:/:/bin/false
saslauth:x:999:76:Saslauthd user:/run/saslauthd:/bin/false
mailnull:x:47:47::/var/spool/mqueue:/bin/false
smmsp:x:51:51::/var/spool/mqueue:/bin/false
chrony:x:998:997::/var/lib/chrony:/bin/false
rpcuser:x:29:29:RPC Service User:/var/lib/nfs:/bin/false
uucp:x:10:14:Uucp user:/var/spool/uucp:/bin/false
nslcd:x:65:55:LDAP Client User:/:/bin/false
tcpdump:x:72:72::/:/bin/false
exawatch:x:1010:510::/:/bin/false
sssd:x:997:508:User for sssd:/:/bin/false
tss:x:59:59:Account used by the trousers package to sandbox the tcsd
daemon:/dev/null:/bin/false
dbmsvc:x:12137:11137::/home/dbmsvc:/bin/false
```

- **Default Users With Restricted Shell Access**
  These users are used for accomplishing a defined task with a restricted shell login. These users should not be altered as the defined task will fail in case these users are altered or deleted.

  `dbmmonitor`: The `dbmmonitor` user is used for DBMCLI Utility. For more information, see *Using the DBMCLI Utility*.

  Restricted Shell Users

```
dbmmonitor:x:2003:2003::/home/dbmmonitor:/bin/rbash
```

- **Default Users with Login permissions**

Copyright © 2008, 2026, Oracle and/or its affiliates.

These privileged users are used for accomplishing most of the tasks in the system. These users should never be altered or deleted as it would have significant impact on the running system.

SSH keys are used for login by customer staff and cloud automation.

Customer-added SSH keys may be added by the `UpdateVmCluster` method, or by customers directly accessing the customer VM and managing SSH keys inside of the customer VM. Customers are responsible for adding comments to keys to make them identifiable. When a customer adds the SSH key by the `UpdateVmCluster` method, the key is only added to the `authorized_keys` file of the `opc` user.

Cloud automation keys are temporary, specific to a given cloud automation task, for example, VM Cluster Memory resize, and unique. Cloud automation access keys can be identified by the following comments: `OEDA_PUB`, `EXACLOUD_KEY`, `ControlPlane`. Cloud automation keys are removed after the cloud automation task completes so the `authorized_keys` files of the `root`, `opc`, `oracle`, and `grid` accounts should only contain cloud automation keys while the cloud automation actions are running.

Privileged Users

```
root:x:0:0:root:/root:/bin/bash
oracle:x:1001:1001::/home/oracle:/bin/bash
grid:x:1000:1001::/home/grid:/bin/bash
opc:x:2000:2000::/home/opc:/bin/bash
dbmadmin:x:2002:2002::/home/dbmadmin:/bin/bash
```

- `root`: Linux requirement, used sparingly to run local privileged commands. `root` is also used for some processes like Oracle Trace File Analyzer Agent and `ExaWatcher`.
- `grid`: Owns Oracle Grid Infrastructure software installation and runs Grid Infastructure processes.
- `oracle`: Owns Oracle database software installation and runs Oracle Database processes.
- `opc`:
  - Used by Oracle Cloud automation for automation tasks.
  - Has the ability to run certain privileged commands without further authentication (to support automation functions).
  - Runs the local agent, also known as “DCS Agent” that performs lifecycle operations for Oracle Database and Oracle Grid Infastructure software (patching, create database, and so on).
- `dbmadmin`:
  - The `dbmadmin` user is used for Oracle Exadata Database Machine Command-Line Interface (DBMCLI) utility.
  - The `dbmadmin` user should be used to run all services on the database server. For more information, see *Using the DBMCLI Utility*.

Related Topics

- Using the DBMCLI Utility

Guest VM Default Fixed Groups

Copyright © 2008, 2026, Oracle and/or its affiliates.

The following groups are created by default on the guest virtual machine.

Oracle Linux 7

```
root:x:0:
bin:x:1:
daemon:x:2:
sys:x:3:
adm:x:4:
tty:x:5:
disk:x:6:
lp:x:7:
mem:x:8:
kmem:x:9:
wheel:x:10:
cdrom:x:11:
mail:x:12:
man:x:15:
dialout:x:18:
tape:x:33:
video:x:39:
lock:x:54:
audio:x:63:
nobody:x:99:
users:x:100:
utmp:x:22:
utempter:x:35:
input:x:999:
systemd-journal:x:190:
systemd-network:x:192:
dbus:x:81:
ssh_keys:x:998:
sshd:x:74:
rpm:x:37:
rpc:x:32:
unbound:x:997:
nscd:x:28:
tss:x:59:
saslauth:x:76:
mailnull:x:47:
smmsp:x:51:
chrony:x:996:
rpcuser:x:29:
ldap:x:55:
slocate:x:21:
uucp:x:14:
tcpdump:x:72:
exawatch:x:510:
cgred:x:509:
fuse:x:508:
screen:x:84:
sssd:x:507:
printadmin:x:506:
docker:x:505:
dbmsvc:x:2001:
dbmadmin:x:2002:
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
dbmmonitor:x:2003:
dbmusers:x:2004:
floppy:x:19:
oinstall:x:1001:
asmdba:x:1004:
asmoper:x:1005:
asmadmin:x:1006:
dba:x:1002:
racoper:x:1003:
opc:x:2000:
```

Oracle Linux 8

```
root:x:0:
bin:x:1:
daemon:x:2:
sys:x:3:
adm:x:4:
tty:x:5:
disk:x:6:
lp:x:7:
mem:x:8:
kmem:x:9:
wheel:x:10:
cdrom:x:11:
mail:x:12:
man:x:15:
dialout:x:18:
tape:x:33:
video:x:39:
lock:x:54:
audio:x:63:
users:x:100:
nobody:x:2001:
utmp:x:22:
utempter:x:35:
dbus:x:81:
input:x:999:
kvm:x:36:
render:x:998:
systemd-journal:x:190:
systemd-coredump:x:997:
systemd-resolve:x:193:
tss:x:59:
ssh_keys:x:996:
sshd:x:74:
unbound:x:995:
nscd:x:28:
rpc:x:32:
saslauth:x:76:
mailnull:x:47:
smmsp:x:51:
rpcuser:x:29:
ldap:x:55:
slocate:x:21:
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

```
chrony:x:994:
tcpdump:x:72:
cgred:x:993:
fapolicyd:x:992:
printadmin:x:991:
polkitd:x:990:
sssd:x:989:
rpm:x:37:
screen:x:84:
dnsmasq:x:988:
exawatch:x:510:
dbmsvc:x:2002:
dbmadmin:x:2003:
dbmmonitor:x:2004:
dbmusers:x:2005:
uucp:x:14:
floppy:x:19:
mausers:x:2006:
secscan:x:1009:
oinstall:x:1001:
asmdba:x:1004:
asmoper:x:1005:
asmadmin:x:1006:
dba:x:1002:
racoper:x:1003:
opc:x:2000:
```

Guest VM Default Security Settings

In addition to all of the Exadata features explained in *Security Features of Oracle Exadata Database Machine*, the following security settings are also applicable.

- Custom database deployment with non-default parameters.
  The command `host_access_control` is to configure Exadata security settings:
  - Implementing password aging and complexity policies.
  - Defining account lockout and session timeout policies.
  - Restricting remote root access.
  - Restricting network access to certain accounts.
  - Implementing login warning banner.
- `account-disable`: Disables a user account when certain configured conditions are met.
- `pam-auth`: Various PAM settings for password changes and password authentication.
- `rootssh`: Adjusts the `PermitRootLogin` value in `/etc/ssh/sshd_config`, which permits or denies the `root` user to login through SSH.
  - By default, `PermitRootLogin` is set to without-password. This is done to permit ExaC@C Cloud Automation to perform software updates and other actions invoked by the customer.
  - It is recommended to preserve the default configuration to permit Cloud Automation functionality. Changing the `PermitRootLogin` to `no` will disable Cloud Automation functionality that requires `root` SSH access to the customer VM.

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - Customers may choose to set `PermitRootLogin` to `no` during normal operations. However, when preparing for maintenance activities, it must be set to `without-password` to ensure automation tasks can complete successfully. After maintenance, it can be reverted to `no` to maintain the desired steady-state configuration.
    As with any customization, it is recommended to validate the configuration to ensure the service operates as expected.

- `session-limit`: Sets the `* hard maxlogins` parameter in `/etc/security/limits.conf`, which is the maximum number of logins for all users. This limit does not apply to a user with `uid=0`.
  Defaults to `* hard maxlogins 10` and it is the recommended secure value.
- `ssh-macs`: Specifies the available Message Authentication Code (MAC) algorithms.
  - The MAC algorithm is used in protocol version 2 for data integrity protection.
  - Defaults to `hmac-sha1`, `hmac-sha2-256`, `hmac-sha2-512` for both server and client.
  - Secure recommended values: `hmac-sha2-256`, `hmac-sha2-512` for both server and client.
- `password-aging`: Sets or displays the current password aging for interactive user accounts.
  - `-M`: Maximum number of days a password may be used.
  - `-m`: Minimum number of days allowed between password changes.
  - `-W`: Number of days warning given before a password expires.
  - Defaults to `-M 99999`, `-m 0`, `-W 7`
  - `--strict_compliance_only -M 60`, `-m 1`, `-W 7`
  - Secure recommended values: `-M 60`, `-m 1`, `-W 7`

Related Topics

- Security Features of Oracle Exadata Database Machine

Guest VM Default Processes

- Exadata Cloud@Customer Guest VM agent: Cloud agent for handling database lifecycle operations
  - Runs as `opc` user.
  - Process table shows it running as a Java process with `jar` names, `dbcs-agent-VersionNumber-SNAPSHOT.jar` and `dbcs-admin-VersionNumver-SNAPSHOT.jar`.
- Oracle Trace File Analyzer agent: Oracle Trace File Analyzer (TFA) provides a number of diagnostic tools in a single bundle, making it easy to gather diagnostic information about the Oracle Database and Oracle Clusterware, which in turn helps with problem resolution when dealing with Oracle Support.
  - Runs as `root` user.
  - Runs as `initd` demon, `/etc/init.d/init.tfa`.
  - Process tables show a Java application, `oracle.rat.tfa.TFAMain`.
- `ExaWatcher`:
  - Runs as `root` and `exawatch` users.

Copyright © 2008, 2026, Oracle and/or its affiliates.

  - Runs as backgroud script, `ExaWatcher.sh` and all its child process run as a Perl process.
  - Process table shows as multiple Perl applications.
- Oracle Database and Oracle Grid Infrastructure (Oracle Clusterware):
  - Runs as `dbmsvc` and `grid` users.
  - Process table shows following applications:
    * `oraagent.bin`, `apx_*`, and `ams_*` as `grid` user.
    * `dbrsMain`, and Java applications, `derbyclient.jar` and `weblogic.Server` as `oracle` user.
- Management Server (MS): Part of Exadata image software for managing and monitoring the image functions.
  - Runs as `dbmadmin` user.
  - Process table shows it running as a Java process.

Guest VM Network Security

Table 7-27 Default Port Matrix for Guest VM Services

Type of interface	Name of interface	Port	Process running
Bridge on client VLAN	`bondeth0`	22	sshd
		1521	Oracle TNS listener
		Optionally, customers can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. Default is 1521.	
		5000	Oracle Trace File Analyzer Collector
		7879	Jetty Management Server
	`bondeth0:1`	1521	Oracle TNS listener
		Optionally, customers can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. Default is 1521.	
	`bondeth0:2`	1521	Oracle TNS listener
		Optionally, customers can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. Default is 1521.	
Bridge on backup VLAN	`bondeth1`	7879	Jetty Management Server

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-27 (Cont.) Default Port Matrix for Guest VM Services

Type of interface	Name of interface	Port	Process running
Oracle Clusterware running on each cluster node communicates through these interfaces.	clib0/clre0	1525	Oracle TNS listener
		3260	Synology DSM iSCSI
		5054	Oracle Grid Interprocess Communication
		7879	Jetty Management Server
		Dynamic Port: 9000-65500 Ports are controlled by the configured ephemeral range in the operating system and are dynamic.	System Monitor service (osysmond)
		Dynamic Port: 9000-65500 Ports are controlled by the configured ephemeral range in the operating system and are dynamic.	Cluster Logger service (ologgerd)
	clib1/clre1	5054	Oracle Grid Interprocess communication
		7879	Jetty Management Server
Cluster nodes use these interfaces to access storage cells (ASM disks). However, the IP/ports 7060/7070 attached to the storage interfaces are used to access DBCS agent from the Control Plane server.	stib0/stre0	7060	dbcs-admin
		7070	dbcs-agent
	stib1/stre1	7060	dbcs-admin
		7070	dbcs-agent
Control Plane server to Guest VM	eth0	22	sshd
Loopback	lo	22	sshd
		2016	Oracle Grid Infrastructure
		6100	Oracle Notification Service (ONS), part of Oracle Grid Infrastructure
		7879	Jetty Management Server
		Dynamic Port 9000-65500	Oracle Trace File Analyzer

Copyright © 2008, 2026, Oracle and/or its affiliates.

> **Note**
>
> TNS listener opens dynamic ports after initial contact to well known ports (1521, 1525).

Default iptables rules for Guest VM:

The default iptables are setup to ACCEPT connections on input, forward, and output chains.

```
#iptables -L -n -v
Chain INPUT (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source
destination

Chain FORWARD (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source
destination

Chain OUTPUT (policy ACCEPT 0 packets, 0 bytes)
 pkts bytes target     prot opt in     out     source
destination
```

OS/VM backups

Oracle does a full backup of guest VM weekly and maintains one or more backup copies. These backups are full disk snapshots of the guest VM (local OS filesystems, not ASM disk groups which reside on Exadata storage). This backup is triggered at a preset time every week. The backups are stored locally in the dom0. Customers can request Oracle to restore the guest VM image from the most recent backup by filing a My Oracle Support (MOS) Service Request (SR). Oracle cannot restore specific files from the image backup. Customers should perform file level backups in the guest VM if they require the ability to perform single-file restore.

Additional Procedures for Updating Security Posture

- Customer Responsibilities
- Enabling Additional Security Capabilities

Customer Responsibilities

Table 7-28 Resposibilities

	Oracle Cloud Platform		Customer/Tenant Instances	
Monitoring	Oracle Cloud Ops	Customer	Oracle Cloud Ops	Customer

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-28 (Cont.) Resposibilities

	Oracle Cloud Platform		Customer/Tenant Instances	
	Infrastructure, Control Plane, hardware faults, availability, capacity	Provide network access to support Oracle infrastructure log collection and monitoring	Infrastructure availability to support customer monitoring of customer services	Monitoring of customer operating system, databases, apps
Incident Management and Resolution	Incident Management and Remediation Spare parts and field dispatch	Onsite diagnostic assistance, for example, network troubleshooting	Support for any incidents related to the underlying platform	Incident Management and resolution for Customer's apps
Patch Management	Proactive patching of hardware, IaaS/ PaaS control stack	Provide network access to support patch delivery	Staging of available patches, for example, Oracle Database patch set	Patching of tenant instances Testing
Backup and Restoration	Infrastructure and Control Plane backup and recovery, receate customer VMs	Provide network access to support cloud automation delivery	Provide running and customer accessible VM	Snapshots / backup and recovey of customer's IaaS and PaaS data using Oracle native or third-party capability
Cloud Support	Response and resolution of SR related to infrastructure or subscription issues	Submit SR through MOS	Response and resolution of SR	Submit SR through Suppot portal

Enabling Additional Security Capabilities

Using Oracle Key Vault as an External TDE Key Store for Databases on Exadata Cloud@Customer

Oracle supports customers using Oracle Key Vault (OKV) as an external key store for databases running on Exadata Cloud@Customer. Instructions for migrating TDE Master Keys to OKV are published in My Oracle Support Document 2823650.1 (*Migration of File based TDE to OKV for Exadata Database Service on Cloud at Customer Gen2*). Oracle does not support third-party hardware security modules (HSM) with Exadata Cloud@Customer.

Modifying Password Complexity Requirements Using host_access_control

Table 7-29 host_access_control password-aging

Option	Description
`-s, --status`	Displays current user password aging.
`-u USER, --user=USER`	A valid interactive user's username.
`--defaults`	Sets all password-aging values to *Exadata factory defaults for all interactive users.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-29 (Cont.) host_access_control password-aging

Option	Description
`--secdefaults`	Sets all password-aging values to **Exadata secure defaults for all interactive users.
`--policy`	Sets all password-aging values to the aging policy as defind by the password-policy command (or `/etc/login.defs`) for all interactive users.
`-M int, --maxdays=int`	Maximum number of days a password may be used. Input limited to from 1 to 99999.
`-m int, --mindays=int`	Minimum number of days allowed between password changes. Input limited to from 0 to 99999, 0 for anytime.
`-W int, --warndays=int`	Number of days warning given before a password expires. Input limited to from 0 to 99999.

`host_access_control password-policy`	`--PASS_MAX_DAYS integer (60)*`
	`--PASS_MIN_DAYS integer ( 1)*`
	`--PASS_MIN_LEN integer ( 8)*`
	`--PASS_WARN_AGE integer ( 7)*`
	`--defaults`
	`--status`

Table 7-30 host_access_control pam-auth

Options	Description
`-h, --help`	Shows this help message and exits.
`-d DENY, --deny=DENY`	Number of failed login attempts before an account will be locked. Input is limited to from 1 to 10. (*Exadata factory default is 5)
`-l LOCK_TIME, --lock=LOCK_TIME`	Number of seconds (integer) an account will be locked due to a single failed login attempt. Input is limited to from 0 to 31557600 (one year) (*Exadata factory default is 600 (10m))
`-p list, --passwdqc=list`	FOR SYSTEMS RUNNING ON LESS THAN OL7 Comma-separated set of 5 values: N0,N1,N2,N3,N4 defining the minimum allowed length for different types of password/passphrases. Each subsequent number is required to be no larger than the preceding one. The keyword "disabled" can be used to disallow passwords of a given kind regardless of their length. (Refer to the pam_passwdqc manpage for an explanation). Passwords must use three character classes. Character classes for passwords are digits, lowercase letters, uppercase letters, and other characters. Minimum password length is 12 characters when using three character classes. Minimum password length is 8 characters when using four character classes. (*Exadata factory default is 5,5,5,5,5) (**Exadata secure default is disabled,disabled,16,12,8)

Copyright © 2008, 2026, Oracle and/or its affiliates.

Table 7-30 (Cont.) host_access_control pam-auth

Options	Description
`-q PWQUALITY, --pwquality=PWQUALITY`	FOR SYSTEMS RUNNING ON OL7 AND GREATER Integer, ranging from 6 to 40, defining the minimum allowed password length. defined by the Exadata secure defaults. All classes will be required for password changes as well as other checks enforced for lengths >7. For lengths <8, class requirements are not used. (*Exadata factory default is: minlen=8 dcredit=-1 ucredit=-1 lcredit=-1 ocredit=-1 difok=8 maxrepeat=3 maxclassrepeat=4) (**Exadata secure default is: minlen=15 dcredit=-1 ucredit=-1 lcredit=-1 ocredit=-1 difok=8 maxrepeat=3 maxclassrepeat=4) (Refer to the pam_pwquality manpage for details)
`-r REMEMBER, --remember=REMEMBER`	The last *n* passwords to remember for password change history. Valid range is an integer from 0 to 1000. (*Exadata factory default is 10)
`--defaults`	Sets all pam-auth values to *Exadata factory defaults.
`--secdefaults`	Sets all pam-auth values to **Exadata secure defaults.
`-s, --status`	Displays current PAM authentication settings.

Implementing or Updating the iptables firewall Configuration in Guest VM

iptables configuration and firewall rules are stored in `/etc/sysconfig/iptables`.

`man iptables` : To get iptables help. Various websites online have many tutorials as well.

`iptables --list` : To get current iptables rules.

Refer to earlier section "Guest VM Network Security" for details on what ports may be required on Guest VM. To configure the firewall manually, create commands like the following example. Note that it is possible to lock yourself out of the system by blocking the ports over which you connect, so it's recommended to consult a test system and engage an experienced iptables administrator if possible.

1. At the command prompt, enter the appropriate command for each port to be opened, for example:

```
# iptables -A INPUT -m state --state NEW -m tcp -p tcp --dport 7002 -j
ACCEPT

# iptables -A INPUT -m state --state NEW -m udp -p udp --dport 123 -j
ACCEPT
```

2. Save the iptables configuration.

```
# service iptables save
```


Copyright © 2008, 2026, Oracle and/or its affiliates.

Changing passwords and Updating Authorized Keys

To change a user password the `password` command is used. Passwords must be changed 7 days prior expiration date. Password policies are described above in the default security settings section.

Default Oracle Exadata Users and Passwords - See My Oracle Support note *https://support.oracle.com/epmos/faces/DocContentDisplay?id=1291766.1*. Other accounts not included in that note are listed below.

Table 7-31 User Accounts

User Name and Password	User Type	Component
`opc` - key-based login only	Operating system user	Oracle Exadata Database Servers
`exawatch` (release 19.1.0 and later) - no logon privileges	Operating system user	Oracle Exadata Database Servers Oracle Exadata Storage Servers
`SYS/We1come$`	Oracle Database user	Oracle Exadata Database Servers
`SYSTEM/We1come$`	Oracle Database user	Oracle Exadata Database Servers
`MSUser` Management Server (MS) uses this account to manage ILOM and reset it if it detects a hang. The `MSUser` password is not persisted anywhere. Each time MS starts up, it deletes the previous `MSUser` account and re-creates the account with a randomly generated password. Do not modify this account. This account is to be used by MS only.	ILOM user	Database server ILOMs Oracle Exadata Storage Server ILOMs

Pay Attention to What Actions May Impact Service-Related Logins for Cloud Automation

For example, procedures will include ensuring that authorized keys used for cloud automation actions remain intact.

For more information about Physical Network access controls including guidelines for Oracle Cloud Automation, see *Oracle Gen2 Exadata Cloud@Customer Security Controls*.

Oracle Cloud Automation access to the customer VM is controlled through token based SSH. Public keys for Oracle Cloud Automation access are stored in the authorized keys files of the `oracle`, `opc`, and `root` users of the customer VM. The private keys used by the automation are stored and protected by the Oracle Cloud Automation software running in the Exadata Cloud@Customer hardware in the customer's data center. The customer's OCI Identity and Access Management (IAM) controls govern if and how a customer can execute Oracle Cloud Automation functionality against the customer VM and databases. The customer may further control access through the management network and Oracle Cloud Automation keys by blocking network access (firewall rules, disabling network interface), and by revoking the credentials used by the Oracle Cloud Automation (remove the public keys from the authorized keys files). Oracle Cloud Automation Access may be temporarily restored by the customer to

Copyright © 2008, 2026, Oracle and/or its affiliates.

permit the subset of functionality required to access the customer VM and customer databases, such as customer VM operating system patching. Oracle Cloud Automation does not need network access the customer VM to perform OCPU scaling, and OCPU scaling functionality will function normally when customers block Oracle Cloud Automation network access to the customer VM.

Configure Encrypted Channels for Database Listener (Oracle Net) Connectivity

For more information, see *Configuring Oracle Database Native Network Encryption and Data Integrity*.

Related Topics

- Managing the Keystore and the Master Encryption Key
- https://support.oracle.com/rs?type=doc&id=1291766.1
- https://support.oracle.com/rs?type=doc&id=2823650.1
- Oracle Gen2 Exadata Cloud@Customer Security Controls
- Configuring Oracle Database Native Network Encryption and Data Integrity

EU Data Act Compliance

Oracle supports the goal of fair and transparent data sharing and access within the European Union.

To learn more about how to export data from Oracle Cloud Infrastructure (OCI) Database as a Service (DBaaS) to on-premises systems or other cloud environments, refer to our compliance page in the base database service documentation.

Related Topics

- Export Data from OCI Database Services

Troubleshooting Oracle Exadata Database Service on Cloud@Customer Systems

These topics cover some common issues you might run into and how to address them.

- Patching Failures on Oracle Exadata Database Service on Cloud@Customer Systems
- Obtaining Further Assistance
- VM Operating System Update Hangs During Database Connection Drain
- Adding a VM to a VM Cluster Fails
- Nodelist is not Updated for Data Guard-Enabled Databases
- CPU Offline Scaling Fails
- Standby Database Fails to Restart After Switchover in Oracle Database 11g Oracle Data Guard Setup
- Using Custom SCAN Listener Port With Data Guard On Disaster Recovery Network Causes Data Guard Association Verification Failures
- PDB Creation Fails After Moving Database to a New DB Home (26ai)
- Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Data Guard Association Creation Fails When Using DR Network

Patching Failures on Oracle Exadata Database Service on Cloud@Customer Systems

Patching operations can fail for various reasons. Typically, an operation fails because a database node is down, there is insufficient space on the file system, or the virtual machine cannot access the object store.

- Determining the Problem
  In the Console, you can identify a failed patching operation by viewing the patch history of an Oracle Exadata Database Service on Cloud@Customer system or an individual database.
- Troubleshooting and Diagnosis
  Diagnose the most common issues that can occur during the patching process of any of the Oracle Exadata Database Service on Cloud@Customer components.

Determining the Problem

In the Console, you can identify a failed patching operation by viewing the patch history of an Oracle Exadata Database Service on Cloud@Customer system or an individual database.

A patch that was not successfully applied displays a status of `Failed` and includes a brief description of the error that caused the failure. If the error message does not contain enough information to point you to a solution, you can use the database CLI and log files to gather more data. Then, refer to the applicable section in this topic for a solution.

Troubleshooting and Diagnosis

Diagnose the most common issues that can occur during the patching process of any of the Oracle Exadata Database Service on Cloud@Customer components.

- Database Server VM Issues
  One or more of the following conditions on the database server VM can cause patching operations to fail.
- Oracle Grid Infrastructure Issues
  One or more of the following conditions on Oracle Grid Infrastructure can cause patching operations to fail.
- Oracle Databases Issues
  An improper database state can lead to patching failures.

Database Server VM Issues

One or more of the following conditions on the database server VM can cause patching operations to fail.

Database Server VM Connectivity Problems

Cloud tooling relies on the proper networking and connectivity configuration between virtual machines of a given VM cluster. If the configuration is not set properly, this may incur in failures on all the operations that require cross-node processing. One example can be not being able to download the required files to apply a given patch.

Given the case, you can perform the following actions:

Copyright © 2008, 2026, Oracle and/or its affiliates.

- Verify that your DNS configuration is correct so that the relevant virtual machine addresses are resolvable within the VM cluster.
- Refer to the relevant Cloud Tooling logs as instructed in the *Obtaining Further Assistance* section and contact Oracle Support for further assistance.

Oracle Grid Infrastructure Issues

One or more of the following conditions on Oracle Grid Infrastructure can cause patching operations to fail.

Oracle Grid Infrastructure is Down

Oracle Clusterware enables servers to communicate with each other so that they can function as a collective unit. The cluster software program must be up and running on the VM Cluster for patching operations to complete. Occasionally you might need to restart the Oracle Clusterware to resolve a patching failure.

In such cases, verify the status of the Oracle Grid Infrastructure as follows:

```
./crsctl check cluster
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
```

If Oracle Grid Infrastructure is down, then restart by running the following commands:

```
crsctl start cluster -all
```

```
crsctl check cluster
```

Oracle Databases Issues

An improper database state can lead to patching failures.

Oracle Database is Down

The database must be active and running on all the active nodes so the patching operations can be completed successfully across the cluster.

Use the following command to check the state of your database, and ensure that any problems that might have put the database in an improper state are resolved:

```
srvctl status database -d db_unique_name -verbose
```

The system returns a message including the database instance status. The instance status must be **Open** for the patching operation to succeed.

If the database is not running, use the following command to start it:

```
srvctl start database -d db_unique_name -o open
```

Obtaining Further Assistance

Copyright © 2008, 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

VM Operating System Update Hangs During Database Connection Drain

Description: This is an intermittent issue. During virtual machine operating system update with 19c Grid Infrastructure and running databases, `dbnodeupdate.sh` waits for `RHPhelper` to drain the connections, which will not progress because of a known bug "DBNODEUPDATE.SH HANGS IN RHPHELPER TO DRAIN SESSIONS AND SHUTDOWN INSTANCE".

Symptoms: There are two possible outcomes due to this bug:

1. VM operating system update hangs in `rhphelper`
   - Hangs the automation
   - Some or none of the database connections will have drained, and some or all of the database instances will remain running.
2. VM operating system update does not drain database connections because `rhphelper` crashed
   - Does not hang automation
   - Some or none of the database connection draining completes

Copyright © 2008, 2026, Oracle and/or its affiliates.

`/var/log/cellos/dbnodeupdate.trc` will show this as the last line:

```
(ACTION:) Executing RHPhelper to drain sessions and shutdown instances.
(trace:/u01/app/grid/crsdata/scaqak04dv0201/rhp//executeRHPDrain.150721125206.trc)
```

Action:

1. Upgrade Grid Infrastructure version to 19.11 or above.
   (OR)

   Disable `rhphelper` before updating and enable it back after updating.

   To disable before updating is started:

   ```
   /u01/app/19.0.0.0/grid/srvm/admin/rhphelper /u01/app/19.0.0.0/grid
   19.10.0.0.0 -setDrainAttributes ENABLE=false
   ```

   To enable after updating is completed:

   ```
   /u01/app/19.0.0.0/grid/srvm/admin/rhphelper /u01/app/19.0.0.0/grid oracle-
   home-current-version -setDrainAttributes ENABLE=true
   ```

   If you disable `rhphelper`, then there will be no database connection draining before database services and instances are shutdown on a node before the operating system is updated.

2. If you missed disabling RHPhelper and upgrade is not progressing and hung, then it is observed that the draining of services is taking time:

   a. Inspect the `/var/log/cellos/dbnodeupdate.trc` trace file, which contains a paragraph similar to the following:

      ```
      (ACTION:) Executing RHPhelper to drain sessions and shutdown instances.
      (trace: /u01/app/grid/crsdata/<nodename>/rhp//executeRHPDrain.150721125206.trc)
      ```

   b. Open the `/var/log/cellos/dbnodeupdate.trc` trace file.
      If `rhphelper` fails, then the trace file contains the message as follows:

      ```
      "Failed execution of RHPhelper"
      ```

      If `rhphelper` hangs, then the trace file contains the message as follows:

      ```
      (ACTION:) Executing RHPhelper to drain sessions and shutdown instances.
      ```

   c. Identify the `rhphelper` processes running at the operating system level and kill them. There are two commands that will have the string “rhphelper” in the name – a Bash shell, and the underlying Java program, which is really `rhphelper` executing.

      `rhphelper` runs as `root`, so must be killed as `root` (`sudo` from `opc`).

      For example:

      ```
      [opc@<HOST> ~] pgrep –lf rhphelper
      191032 rhphelper
      191038 java


      [opc@<HOST> ~] sudo kill –KILL 191032 191038
      ```

   d. Verify that the `dbnodeupdate.trc` file moves and the Grid Infrastructure stack on the node is shutdown.

Copyright © 2008, 2026, Oracle and/or its affiliates.

For more information about RHPhelper, see *Using RHPhelper to Minimize Downtime During Planned Maintenance on Exadata (Doc ID 2385790.1)*.

Related Topics

- Using RHPhelper to Minimize Downtime During Planned Maintenance on Exadata (Doc ID 2385790.1)

Adding a VM to a VM Cluster Fails

Description: When adding a VM to a VM cluster, you might encounter the following issue:

```
[FATAL] [INS-32156] Installer has detected that there are non-readable files in oracle
home.
CAUSE: Following files are non-readable, due to insufficient permission oracle.ahf/data/
scaqak03dv0104/diag/tfa/tfactl/user_root/tfa_client.trc
ACTION: Ensure the above files are readable by grid.
```

Cause: Installer has detected a non-readable trace file, `oracle.ahf/data/scaqak03dv0104/diag/tfa/tfactl/user_root/tfa_client.trc` created by Autonomous Health Framework (AHF) in Oracle home that causes adding a cluster VM to fail.

AHF ran as `root` created a `trc` file with `root` ownership, which the `grid` user is not able to read.

Action: Ensure that the AHF trace files are readable by the `grid` user before you add VMs to a VM cluster. To fix the permission issue, run the following commands as `root` on all the existing VM cluster VMs:

```
chown grid:oinstall /u01/app/19.0.0.0/grid/srvm/admin/logging.properties


chown -R grid:oinstall /u01/app/19.0.0.0/grid/oracle.ahf*


chown -R grid:oinstall /u01/app/grid/oracle.ahf*
```

Nodelist is not Updated for Data Guard-Enabled Databases

Description: Adding a VM to a VM cluster completes successfully, however, for Data Guard-enabled databases, the new VM is not added to the nodelist in the `/var/opt/oracle/creg/<db>.ini` file.

Cause: Data Guard-enabled databases will not be extended to the newly added VM. And therefore, the `<db>.ini` file will also not be updated because the database instance is not configured in the new VM.

Action: To add an instance to primary and standby databases and to the new VMs (Non-Data Guard), and to remove an instance from a Data Guard environment, see *My Oracle Support note 2811352.1*.

Related Topics

- https://support.oracle.com/epmos/faces/DocContentDisplay?id=2811352.1

Copyright © 2008, 2026, Oracle and/or its affiliates.

CPU Offline Scaling Fails

Description: CPU offline scaling fails with the following error:

```
** CPU Scale Update **An error occurred during module execution. Please refer to the log file for more information
```

Cause: After provisioning a VM cluster, the `/var/opt/oracle/cprops/cprops.ini` file, which is automatically generated by the database as a service (DBaaS) is not updated with the `common_dcs_agent_bindHost` and `common_dcs_agent_port` parameters and this causes CPU offline scaling to fail.

Action: As the `root` user, manually add the following entries in the `/var/opt/oracle/cprops/cprops.ini` file.

```
common_dcs_agent_bindHost=<IP_Address>
common_dcs_agent_port=7070
```

> **Note**
>
> The `common_dcs_agent_port` value is **7070** always.

Run the following command to get the IP address:

```
netstat -tunlp | grep 7070
```

For example:

```
netstat -tunlp | grep 7070
tcp 0 0 <IP address 1>:7070 0.0.0.0:* LISTEN 42092/java
tcp 0 0 <IP address 2>:7070 0.0.0.0:* LISTEN 42092/java
```

You can specify either of the two IP addresses, *<IP address 1>* or *<IP address 2>* for the `common_dcs_agent_bindHost` parameter.

Standby Database Fails to Restart After Switchover in Oracle Database 11g Oracle Data Guard Setup

Description: After performing the switchover, the new standby (old primary) database remains shut down and fails to restart.

Action: After performing switchover, do the following:

1. Restart the standby database using the `srvctl start database -db <standby dbname>` command.
2. Reload the listener as `grid` user on all primary and standby cluster nodes.
   - To reload the listener using high availability, download and apply patch **25075940** to the Grid home, and then run `lsnrctl reload -with_ha`.

Copyright © 2008, 2026, Oracle and/or its affiliates.

- To reload the listener, run `lsrnctl reload`.

After reloading the listener, verify that the <*dbname*>_DGMGRL services are loaded into the listener using the `lsnrctl status` command.

To download patch 25075940

1. Log in to My Oracle Support.
2. Click **Patches & Updates**.
3. Select **Bug Number** from the **Number/Name or Bug Number (Simple)** drop-down list.
4. Enter the bug number **34741066**, and then click **Search**.
5. From the search results, click the name of the latest patch.
   You will be redirected to the **Patch 34741066: LSNRCTL RELOAD -WITH_HA FAILED TO READ THE STATIC ENTRY IN LISTENER.ORA** page.
6. Click **Download**.

Using Custom SCAN Listener Port With Data Guard On Disaster Recovery Network Causes Data Guard Association Verification Failures

Description: If the SCAN listener port for the client network and disaster recovery network (DR network) are different, then Data Guard (DG) configuration fails during verification phase of create data guard association.

Action: Use the same SCAN listener ports (or default port) on all networks. To fix the listener port after the cluster has been configured, run the *GI home*`/bin/srvctl modify listener -listener` *listener_name* `-endpoints` *endpoints* command. For more information, see srvctl modify listener in the Oracle Real Application Clusters Administration and Deployment Guide.

PDB Creation Fails After Moving Database to a New DB Home (26ai)

Description: After relocating a database to a different DB Home, creating a new Pluggable Database (PDB) fails. The PDB service fails to start, resulting in the following error:

```
[FATAL] [DBAAS-60022] Command '/u02/app/oracle/product/23.0.0.0/dbhome_3/bin/srvctl
'start' 'service' '-db' 'db23ano' '-service' 'db23ano_PDBJULY242.paas.oracle.com'' has
failed on nodes [localnode].
```

Action: If the Grid Infrastructure version is 23.4.0.24.05, upgrade to version 23.5.0.24.07 to resolve this issue.

Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel

Description: PDB creation might fail for subset of PDBs when multiple PDBs are getting created in parallel.

Cause: PDB creation might notice following error intermittently.

```
ORA-03113: end-of-file on communication channel
```

Action: Retry the failed PDB creation.

Copyright © 2008, 2026, Oracle and/or its affiliates.

Data Guard Association Creation Fails When Using DR Network

Description:

A workflow failure may occur when creating a Data Guard association over the Disaster Recovery (DR) network (third NIC installed through a PM exception). Although the system is configured with a DR network, the workflow still requires connectivity over the client network by default. In environments where client-network communication between the two data centers is restricted, the association creation will fail.

Cause:

By default, the Data Guard association workflow attempts to use the client network, even when a DR network exists. If the database instances are not registered with the DR-network listener, the workflow will fall back to the client network, resulting in a failure when client-network paths are blocked or unavailable.

Action:

Register the database instances with the listener on the DR network by updating the listener_networks parameter. This ensures the workflow uses the DR network exclusively for the Data Guard association.

Example configuration

- Instance 1

```
(( NAME=net2)
 (LOCAL_LISTENER=
   (DESCRIPTION=
     (ADDRESS=(PROTOCOL=TCP)(HOST=VIPDR-1)(PORT=1521))
   )
 )
),
(( NAME=net2)(REMOTE_LISTENER=SCAN-DR-NAME:1521))
```

- Instance 2

```
(( NAME=net2)
 (LOCAL_LISTENER=
   (DESCRIPTION=
     (ADDRESS=(PROTOCOL=TCP)(HOST=VIPDR-2)(PORT=1521))
   )
 )
),
(( NAME=net2)(REMOTE_LISTENER=SCAN-DR-NAME:1521))
```

Summary:

- Although a DR network is available, the workflow still requires client-network connectivity by default during Data Guard association creation.
- To use only the DR network, you must manually register the service with the DR-network listener using the `listener_networks` parameter.
- After the listener is updated and the service registers correctly, the Data Guard association workflow proceeds over the DR network instead of the client network.

Copyright © 2008, 2026, Oracle and/or its affiliates.