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



F84646-42 May 2024

ORACLE

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

F84646-42

Copyright © 2008, 2024, 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, 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, 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 Exadata Database Service on Cloud@Customer Overview

Oracle Exadata Database Service on Cloud@Customer Service Description	1-1
About Oracle Exadata Database Service on Cloud@Customer	1-2
Licensing Considerations for Oracle Exadata Database Service on Cloud@Customer	1-3
Per-Second Billing for OCPU Usage	1-3
Supported Database Edition and Versions	1-3
System and Shape Configuration Options	1-4
System Configuration Options for Oracle Exadata Cloud@Customer	1-4
Oracle Exadata X10M System Model Specifications	1-4
Oracle Exadata X9M-2 System Model Specifications	1-5
Oracle Exadata X8M-2 System Model Specifications	1-6
Oracle Exadata X8-2 System Model Specifications	1-7
Oracle Exadata X7-2 System Model Specifications	1-8
Exadata Cloud Management Interfaces	1-9
Introduction to Exadata Cloud Management Interfaces	1-9
OCI Control Plane Interfaces	1-10
Local VM Command-Line Interfaces	1-12

2 What's New in Oracle Exadata Database Service on Cloud@Customer Gen2

Oracle Database 23ai on Oracle Exadata Database Service on Cloud@Customer	2-3
Change the Home Region of ExaDB-C@C Infrastructure	2-3
Enhancements to Serial Console Functionality	2-4
Provision a VM Cluster with Either an OL7 or OL8-Based Image	2-5
VM Cluster on a Single VM	2-5
Enhancement to Pluggable Database (PDB) Management	2-5
Manage Administrator (SYS User) and TDE Wallet Passwords	2-6
Update Guest VM (domU) Operating System to Oracle Linux 8	2-6
Exadata Fleet Update	2-7
Oracle Key Vault (OKV) Integration with ExaDB-C@C to Manage Transparent Data	
Encryption (TDE) Keys	2-7
X10M System Support	2-7



Manage Serial Console Access to ExaDB-C@C Systems	2-8
Interim Software Updates	2-9
Link Aggregation Control Protocol (LACP) Support for Client and Backup Networks	2-10
VM Cluster and Database Health and Performance Metrics in the OCI Console	2-10
Enable Data Guard with Database Software Image (Custom Image)	2-11
Use Identity and Access Management (IAM) Authentication with Oracle Exadata Database	2 11
Service on Cloud@Customer	2-12
Create Data Guard Associations Across Regions	2-12
Multi-Rack Elastic Compute and Storage Expansion	2-13
Automatic Diagnostic Collection	2-15
Elastic Compute Expansion	2-15
Oracle Standard Tagging for Resources on Oracle Exadata Database Service on Cloud@Customer	2-17
Exadata Infrastructure Maintenance History	2-17
Monthly ExaDB-C@C Infrastructure Security Maintenance	2-18
Enhanced Control to Rollback or Retry Failed Guest VM Operating System Update	2-19
Concurrently Create or Terminate Oracle Databases in a VM Cluster	2-19
Support for Rack Serial Number as a System Tag	2-19
Support for DB Home Minor Version Selection (N-3)	2-19
VM Guest Exadata OS Image Major Version Update	2-20
Database Service Events	2-20
Control Plane Server (CPS) Offline Diagnostic Report	2-21
Enhanced Infrastructure Maintenance Controls	2-22
Manage Pluggable Databases on Exadata Cloud@Customer	2-23
Allow Customers to Choose Data Guard Type	2-23
Specify the Same SID for Primary and Standby Databases in a Data Guard Association	2-24
Specify db_unique_name and SID for Primary and Standby Databases in Data Guard Association	2-24
VM Cluster Node Subsetting	2-24
X9M-2 System Support	2-25
Customize SCAN Listener Port	2-26
Creating DG Association/Standby Database Using Existing Database Home	2-26
Upgrading Oracle Grid Infrastructure on an Exadata Cloud@Customer VM Cluster	2-27
Updating Guest VM Operating System	2-27
Upgrading Oracle Databases	2-27
Download Network Validation Report	2-27
Elastic Storage Expansion	2-28
Oracle Database Software Images	2-29
Exadata Cloud@Customer Infrastructure Patching Automation	2-29
Customer Maintenance Contacts	2-29
X8M-2 System Support	2-30
Enable and Manage Data Guard Associations	2-31



Oracle Exadata Cloud@Customer Deployment Assistant	2-31
Oracle Grid Infrastructure and Oracle Database Patching	2-32
Per-Second Billing for OCPU Usage	2-32
Shared Database Home Resource Tags	2-32
Create and Manage Multiple Virtual Machines per Exadata System (MultiVM)	2-33
Scale OCPUs Without Cloud Connectivity	2-33
Configure Oracle Database Character Set and National Character Set	2-34
Specify a Time Zone While Provisioning Oracle Exadata Database Service on	
Cloud@Customer Infrastructure	2-35
Shared Database Homes for Oracle Exadata Database Service on Cloud@Customer	2-35
X7-2 System Support	2-36

3 Preparing for Exadata Database Service on Cloud@Customer

Oracle Cloud Infrastructure (OCI) Requirements for Oracle Exadata Database Service on Cloud@Customer	3-1
Required IAM Policy for Exadata Database Service on Cloud@Customer	3-2
Site Requirements for Oracle Exadata Database Service on Cloud@Customer	3-3
Space Requirements for Oracle Exadata Database Service on Cloud@Customer	3-4
Weight of Oracle Exadata Database Service on Cloud@Customer Racks	3-4
Receiving, Unpacking, and Access for Oracle Exadata Database Service on Cloud@Customer Racks	3-4
Flooring for Oracle Exadata Database Service on Cloud@Customer Racks	3-5
Electrical Power for Oracle Exadata Database Service on Cloud@Customer Racks	3-5
Temperature and Humidity Ranges for Oracle Exadata Database Service on Cloud@Customer	3-7
Ventilation for Oracle Exadata Database Service on Cloud@Customer Racks	3-8
Network Requirements for Oracle Exadata Database Service on Cloud@Customer	3-9
Network Requirements for Oracle Exadata Database Service on Cloud@Customer	3-9
Data Center Network Services for Exadata Database Service on Cloud@Customer	3-16
IP Addresses and Subnets for Exadata Database Service on Cloud@Customer	3-16
Uplinks for Exadata Database Service on Cloud@Customer	3-18
Network Cabling for Exadata Database Service on Cloud@Customer	3-19
Establish a Secure Connection Between Your CPS and OCI Using OCI's FastConnect Service	3-19
Storage Configuration Requirements for Oracle Exadata Database Service on Cloud@Customer	3-22
About Storage Configuration for Oracle Exadata Database Service on Cloud@Customer	3-23
Allocation of Storage Space Options on Oracle Exadata Storage Servers	3-23
Allocation Proportions for DATA, RECO and SPARSE Disk Groups	3-24
Virtual Machine File System Structure for Exadata Cloud@Customer	3-25
Checklists for Exadata Database Service on Cloud@Customer Deployments	3-25



System Components Checklist for Exadata Database Service on Cloud@Customer	3-26
Data Center Room Checklist for Exadata Database Service on Cloud@Customer	3-26
Data Center Environment Checklist for Oracle Exadata Database Service on Cloud@Customer	3-27
Access Route Checklist for Oracle Exadata Database Service on Cloud@Customer	3-27
Facility Power Checklist for Oracle Exadata Database Service on Cloud@Customer	3-28
Safety Checklist for Oracle Exadata Database Service on Cloud@Customer	3-28
Logistics Checklist for Exadata Database Service on Cloud@Customer	3-28
Network Configuration Checklist for Oracle Exadata Database Service on	
Cloud@Customer	3-30
Reracking Checklist for Oracle Exadata Database Service on Cloud@Customer	3-30

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

About Provisioning Oracle Exadata Database Service on Cloud@Customer Systems	4-2
Tagging Oracle Exadata Database Service on Cloud@Customer Resources	4-2
Oracle Exadata Cloud@Customer Deployment Assistant	4-8
Using Oracle Exadata Cloud@Customer Deployment Assistant	4-8
Accessing Oracle Exadata Cloud@Customer Deployment Assistant	4-9
Step 1: Pre-Installation	4-9
Step 2: Onsite Installation	4-10
Step 3: Post-Installation	4-10
Overview of Elastic Storage Expansion	4-10
Using the Console to Provision Exadata Database Service on Cloud@Customer	4-14
Using the Console to Create Exadata Cloud@Customer Infrastructure	4-15
Enabling or Disabling the Control Plane Server Diagnostic Offline Report	4-22
Viewing the Control Plane Server Diagnostic Offline Report	4-23
Using the Console to View Exadata Infrastructure Network Configuration Details	4-23
Using the Console to Edit Oracle Exadata Database Service on Cloud@Customer Infrastructure Networking Configuration	4-24
Using the Console to Download a File Containing Configuration Data	4-26
Using the Console to Activate Exadata Database Service on Cloud@Customer	
Infrastructure	4-26
Using the Console to Check the Status of Exadata Database Service on Cloud@Customer Infrastructure	4-27
Configuring Oracle-Managed Infrastructure Maintenance	4-28
About Oracle Managed Exadata Cloud@Customer Infrastructure Maintenance Updates	4-28
Overview of the Quarterly Infrastructure Maintenance Process	4-29
Infrastructure Maintenance Contacts	4-31
Using the Console to Configure Oracle-Managed Infrastructure Updates	4-32
View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure	4-32



View or Edit a Scheduled Quarterly Maintenance for Exadata Cloud@Customer Infrastructure	4-34
View the Maintenance History of Exadata Cloud@Customer Infrastructure	4-37
Monitor Infrastructure Maintenance Using Lifecycle State Information	4-37
Receive Notifications about Your Infrastructure Maintenance Updates	4-38
Creating First VM Cluster Network on Exadata Cloud@Customer	4-38
About Managing VM Cluster Networks on Exadata Database Service on	
Cloud@Customer	4-39
Using the Console to Create a VM Cluster Network	4-39
Using the Console to Edit a VM Cluster Network	4-43
Using the Console to Download a File Containing the VM Cluster Network Configuration Details	4-44
Using the Console to Validate a VM Cluster Network	4-45
Provisioning the First VM Cluster on an Exadata Cloud@Customer System	4-46
About Managing VM Clusters on Exadata Database Service on Cloud@Customer	4-46
Prerequisites for VM Clusters on Exadata Database Service on Cloud@Customer	4-46
Using the Console to Create a VM Cluster	4-47
Creating Database Backup Destinations for Exadata Cloud@Customer	4-53
About Managing Backup Destinations for Exadata Database Service on Cloud@Customer	4-53
Prerequisites for Backup Destinations for Exadata Database Service on Cloud@Customer	4-54
Using the Console for Backup Destinations for Exadata Database Service on Cloud@Customer	4-55
Using the Console to Create a Backup Destination	4-56
Using the Console to Edit a Backup Destination	4-58
Using the Console to Move a Backup Destination to Another Compartment	4-59
Using the Console to Delete a Backup Destination	4-59
Creating First Database Home on an Exadata Database Service on Cloud@Customer System	4-60
About Creating Oracle Database Homes on an Exadata Database Service on	
Cloud@Customer System	4-60
Using the Console to Create Oracle Database Home on Exadata Database Service on Cloud@Customer	4-61
Using the API to Create Oracle Database Home on Exadata Cloud@Customer	4-62
Creating First Database on an Exadata Database Service on Cloud@Customer System	4-63
Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer	4-63
Oracle Database Releases Supported by Oracle Exadata Database Service on Cloud@Customer	4-64
About Provisioning and Configuring Oracle Databases on Oracle Exadata Database Service on Cloud@Customer	4-64
Using the Console to Create a Database	4-65
Connecting to an Exadata Database Service on Cloud@Customer System	4-70
-	



Connecting to a Virtual Machine with SSH	4-70
Prerequisites for Connecting to an Exadata Database Service on Cloud@Customer System	4-71
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY	4-71
Accessing a Database After You Connect to the Virtual Machine	4-72
Connecting from a Unix-Style System	4-74
Connecting to a Database with Oracle Net Services	4-74
Using Oracle Net Services to Connect to a Database	4-74
Prerequisites for Connecting to a Database with Oracle Net Services	4-75
Connecting to a Database Using SCAN	4-75
Connecting to a Database Using a Node Listener	4-77

5 How-to Guides

Manage Database Security with Oracle Data Safe	5-2
About Oracle Data Safe	5-2
Get Started	5-3
Using Oracle Data Safe	5-3
Connect to the Exadata Database Service on Cloud@Customer Service	5-5
Connecting to a Virtual Machine with SSH	5-5
Prerequisites for Connecting to an Exadata Database Service on Cloud@Customer System	5-6
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY	5-6
Connecting from a Unix-Style System	5-7
Accessing a Database After You Connect to the Virtual Machine	5-8
Connecting to a Database with Oracle Net Services	5-9
Using Oracle Net Services to Connect to a Database	5-10
Prerequisites for Connecting to a Database with Oracle Net Services	5-10
Connecting to a Database Using SCAN	5-10
Connecting to a Database Using a Node Listener	5-12
Manage Exadata Database Service on Cloud@Customer Infrastructure	5-13
About Provisioning Oracle Exadata Database Service on Cloud@Customer Systems	5-14
About Changing the Home Region of ExaDB-C@C Infrastructure	5-14
Overview of Elastic Storage Expansion	5-16
Using the Console to Provision Exadata Database Service on Cloud@Customer	5-20
Using the Console to Create Exadata Cloud@Customer Infrastructure	5-21
Using the Console to Edit Network Bonding	5-29
Using the Console to View Exadata Infrastructure Network Configuration Details	5-30
Using the Console to Edit Oracle Exadata Database Service on Cloud@Customer Infrastructure Networking Configuration	5-31
Using the Console to Download a File Containing Configuration Data	5-33

Using the Console to Activate Exadata Database Service on Cloud@Customer Infrastructure	5-33
Using the Console to Check the Status of Exadata Database Service on Cloud@Customer Infrastructure	5-34
Using the Console to Scale Infrastructure Compute and Storage	5-35
Using the Console to Download Scale Infrastructure Compute and Storage Configuration File	5-37
Using the Console to Activate New Database and Storage Servers	5-38
Using the Console to Make Storage Capacity from New Server Available for VM Clusters Consumption	5-38
Using the Console to View Details of Exadata Cloud@Customer Infrastructure with Scaled Compute and Storage Capacity	5-39
Using the Console to Move Exadata Database Service on Cloud@Customer Infrastructure	5-40
Using the Console to Delete Exadata Database Service on Cloud@Customer Infrastructure	5-40
Using the Console to Manage Tags for Your Exadata Cloud@Customer Resources	5-41
Managing Infrastructure Maintenance Contacts	5-41
Using the API to Manage Exadata Cloud@Customer Infrastructure	5-44
Configure Oracle-Managed Infrastructure Maintenance	5-45
About Oracle Managed Exadata Cloud@Customer Infrastructure Maintenance Updates	5-45
Overview of the Quarterly Infrastructure Maintenance Process	5-46
Overview of Monthly Security Maintenance	5-48
Understanding Monthly and Quarterly Maintenance in the Same Month	5-49
Infrastructure Maintenance Contacts	5-52
Using the Console to Configure Oracle-Managed Infrastructure Updates	5-52
View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure	5-53
View or Edit a Scheduled Quarterly Maintenance for Exadata Cloud@Customer Infrastructure	5-55
View or Edit a Scheduled Security Maintenance for Exadata Cloud@Customer Infrastructure	5-57
View the Maintenance History of Exadata Cloud@Customer Infrastructure	5-58
View and Edit Quarterly Maintenance While Maintenance is In Progress or Waiting for Custom Action	5-59
Monitor Infrastructure Maintenance Using Lifecycle State Information	5-61
Receive Notifications about Your Infrastructure Maintenance Updates	5-61
Using the API to Manage Exadata Cloud@Customer Infrastructure Maintenance Controls	5-62
Manage VM Cluster Networks	5-62
About Managing VM Cluster Networks on Exadata Database Service on Cloud@Customer	5-63
Using the Console to Create a VM Cluster Network	5-64
Using the Console to View VM Cluster Network Details	5-67



Using the Console to View SCAN Listener Port Configured	5-67
Using the Console to Edit a VM Cluster Network	5-68
Using the Console to Download a File Containing the VM Cluster Network	F 60
Configuration Details	5-69
Using the Console to Validate a VM Cluster Network	5-70
Using the Console to Download Network Validation Report	5-71
Using the Console to Terminate a VM Cluster Network	5-72
Using the Console to Add Another DB Server to the VM Cluster Network	5-73
Using the Console to Remove DB Servers from a VM Cluster Network	5-73
Manage VM Clusters	5-74
About Managing VM Clusters on Exadata Database Service on Cloud@Customer	5-74
Overview of VM Cluster Node Subsetting	5-75
Overview of Automatic Diagnostic Collection	5-76
Incident Logs and Trace Files	5-77
Health Metrics	5-82
Introduction to Scale Up or Scale Down Operations	5-88
Scaling Up or Scaling Down the VM Cluster Resources	5-88
Resizing Memory and Large Pages	5-89
Calculating the ASM Storage	5-90
Estimating How Much Local Storage You Can Provision to Your VMs	5-91
Scaling Local Storage	5-93
Using the Console to Manage VM Clusters on Exadata Cloud@Customer	5-94
Using the Console to Create a VM Cluster	5-95
Using the Console to Enable, Partially Enable, or Disable Diagnostics Collection	5-101
Using the Console to Add VMs to a Provisioned Cluster	5-103
Using the Console to View a List of DB Servers on an Exadata Infrastructure	5-104
Using the Console to Remove a VM from a VM Cluster	5-105
Using the Console to Update the License Type on a VM Cluster	5-105
Using the Console to Add SSH Keys After Creating a VM Cluster	5-106
Using the Console to Scale the Resources on a VM Cluster	5-106
Using the Console to Stop, Start, or Reboot a VM Cluster Virtual Machine	5-108
Using the Console to Check the Status of a VM Cluster Virtual Machine	5-109
Using the Console to Move a VM Cluster to Another Compartment	5-109
Using the Console to Terminate a VM Cluster	5-110
Using the API to Manage Exadata Cloud@Customer VM Clusters	5-110
Troubleshooting Virtual Machines Using Console Connections	5-111
Required IAM Policies	5-113
Prerequisites	5-113
Create the Virtual Machine Serial Console Connection	5-116
Make an SSH Connection to the Serial Console	5-117
Using Cloud Shell to Connect to the Serial Console	5-122



Displaying the Console History for a Virtual Machine	5-123
Troubleshooting Virtual Machines from Guest VM Console Connections on Linux	
Operating Systems	5-126
Exiting the Virtual Machine Serial Console Connection	5-129
Manage Oracle Database Software Images	5-129
Creation and Storage of Database Software Images	5-130
Using a Database Software Image with an Exadata Cloud@Customer System	5-131
Using the Console to Create a Database Software Image	5-132
Using the OPatch Isinventory Command to Verify the Patches Applied to an Oracle Home	5-133
Using the Console to Delete a Database Software Image	5-133
Using the Console to View the Patch Information of a Database Software Image	5-133
Using the Console to Move Database Software Image to a Different Compartment	5-134
Using the API for Managing Oracle Database Software Images	5-134
Create Oracle Database Homes on an Exadata Database Service on Cloud@Customer	
System	5-135
About Creating Oracle Database Homes on an Exadata Database Service on	
Cloud@Customer System	5-135
Using the Console to Create Oracle Database Home on Exadata Database Service on Cloud@Customer	5-136
Using the API to Create Oracle Database Home on Exadata Cloud@Customer	5-137
Manage Oracle Database Homes on Exadata Database Service on Cloud@Customer Systems	5-138
About Managing Oracle Database Homes on Exadata Database Service on	3-130
Cloud@Customer Systems	5-138
Manage Database Home Using the Console	5-138
Using the Console to View Information About an Oracle Database Home	5-139
Using the Console to Delete an Oracle Database Home	5-139
Using the API to Manage Oracle Database Home on Exadata Database Service on	E 140
Cloud@Customer	5-140 5-140
Differences Between Managing Resources with dbaascli and the Database API Manage Databases on Exadata Database Service on Cloud@Customer	5-140 5-141
	J-141
Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Cloud@Customer	5-142
Oracle Database Releases Supported by Oracle Exadata Database Service on Cloud@Customer	5-142
About Provisioning and Configuring Oracle Databases on Oracle Exadata Database	J-142
Service on Cloud@Customer	5-143
Using the Console to Manage Databases on Oracle Exadata Database Service on	
Cloud@Customer	5-144
Using the Console to Create a Database	5-144
Using the Console to Manage SYS User and TDE Wallet Passwords	5-149
Using the Console to Move a Database to Another Database Home	5-149
Using the Console to Terminate a Database	5-150



Using the API to Manage Oracle Database Components	5-150
Changing the Database Passwords	5-152
To Change the SYS Password for an Exadata Database Service on Cloud@Customer Database	5-152
To Change Database Passwords in a Data Guard Environment	5-152
To Change the TDE Wallet Password for an Exadata Database Service on Cloud@Customer Database	5-153
Manage Pluggable Databases on Exadata Database Service on Cloud@Customer	5-153
Pluggable Database Operations	5-153
Connect to an Oracle Database using Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users	5-175
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database	5-176
Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database	5-179
Enable, Disable, and Re-enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database	5-182
Manage Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies, Users, Roles, and Database Passwords	5-184
Connect to Oracle Database with Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication	5-189
Database Links in an Oracle DBaaS-to-IAM Integration	5-195
Configuring Authorization for IAM Users and Oracle Cloud Infrastructure Applications	5-196
Configure Proxy Authentication	5-203
Manage Database Backup and Recovery on Oracle Exadata Database Service on Cloud@Customer	5-206
Backup Destinations	5-206
About Managing Backup Destinations for Exadata Database Service on Cloud@Customer	5-207
Prerequisites for Backup Destinations for Exadata Database Service on Cloud@Customer	5-208
Using the Console for Backup Destinations for Exadata Database Service on Cloud@Customer	5-209
Using the API to Manage Exadata Cloud@Customer Backup Destinations	5-213
Oracle Database Backup Methods in Exadata Cloud	5-213
Oracle Managed Backup	5-214
User Configured Backup	5-227
Configuring and Customizing Backups with bkup_api	5-235
Customizing Backup Settings by Using a Generated Configuration File	5-236
Customizing Which System Files Are Backed Up	5-242
Customizing Which Database Configuration Files Are Backed Up	5-243
Creating an On-Demand Backup by Using the bkup_api Utility	5-244
Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management	5-246



	Customizing Real Time Redo Transport (RTRT) Behavior for Recovery Appliance	
	Backups	5-249
	Alternative Backup Methods	5-249
	Recovering a Database Using Oracle Recovery Manager (RMAN)	5-250
Pa	tch and Update an Exadata Database Service on Cloud@Customer System	5-251
	Perform User Managed Maintenance Updates	5-251
	Patching and Updating an Exadata Database Service on Cloud@Customer System	5-252
	Patching and Updating VM Clusters and Database Homes	5-252
	Updating Guest VM Operating System	5-258
	Upgrading Oracle Grid Infrastructure on an Exadata Cloud@Customer VM Cluster	5-263
	Upgrading Oracle Databases	5-265
	Patching and Updating an Exadata Database Service on Cloud@Customer System Manually	5-270
	Updating Software Manually	5-271
	Updating the Guest VM Operating System Manually	5-271
Inte	erim Software Updates	5-281
	Create an Interim Software Update	5-281
	Download an Interim Software Update	5-281
	Delete an Interim Software Update	5-282
	Move an Interim Software Update Resource to Another Compartment	5-282
	Using the API to Manage Interim Software Updates	5-282
Us	e Oracle Data Guard with Exadata Database Service on Cloud@Customer	5-283
	About Using Oracle Data Guard with Exadata Database Service on Cloud@Customer	5-284
	Prerequisites for Using Oracle Data Guard with Exadata Database Service on	
	Cloud@Customer	5-284
	VM Clusters	5-284
	Password	5-285
	Adding a Node to a VM Cluster	5-285
	Removing a Node from a VM Cluster	5-285
	Working with Data Guard	5-286
	Switchover	5-286
	Failover	5-286
	Reinstate	5-286
	Using the Console to Manage Oracle Data Guard Associations	5-287
	Using the Console to Enable Data Guard on an Exadata Database Service on Cloud@Customer System	5-288
	Using the Console to View Data Guard Associations of Databases in an Exadata VM Cluster	5-291
	Using the Console To View and Edit Data Guard Associations	5-291
	Using the Console To Perform a Database Switchover	5-292
	Using the Console To Perform a Database Failover	5-292
	Using the Console To Reinstate a Database	5-293
	-	



Using the Console To Terminate a Data Guard Association on an Exadata Database Service on Cloud@Customer System	5-293
Using the API to Manage Data Guard Associations on an Exadata Database Service on Cloud@Customer System	5-294
Managing Encryption Keys on External Devices	5-294
Customer-Managed Keys in Exadata Database Service on Cloud@Customer	5-295
About Oracle Key Vault	5-296
Overview of Key Store	5-296
Required IAM Policy for Managing OKV on Oracle Exadata Database Service on Cloud@Customer	5-297
Tagging Resources	5-298
Moving Resources to a Different Compartment	5-298
Setting Up Your Exadata Cloud@Customer to Work With Oracle Key Vault	5-298
Step 1: Create a Vault in OCI Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password	5-299
Step 2: Create a Dynamic Group and a Policy Statement for Key Store to Access Secret in OCI Vault	5-299
Step 3: Create a Dynamic Group and a Policy Statement for Exadata Infrastructure to Key Store	5-300
Step 4: Create a Policy Statement for Database Service to Use Secret from OCI	
Vault Service	5-301
Step 5: Create Key Store	5-301
Managing Your Key Store	5-302
View Key Store Details	5-302
Edit Key Store Details	5-303
Move a Key Store to Another Compartment	5-303
Delete a Key Store	5-303
View Key Store Associated Container Database Details	5-303
Using the API to Manage Key Store	5-304
Administer Transparent Data Encryption (TDE) Keys	5-304
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	E 206
Key Vault (OKV)	5-306
Source CDB and Target CDB are Encrypted with MEK in the Same OKV Server	5-307
Source CDB and Target CDB are Encrypted with MEK in a Different OKV Server	5-317
Migrate to Exadata Database Service on Cloud@Customer	5-328
Moving to Oracle Cloud Using Zero Downtime Migration	5-328
Overview of Exadata Cloud@Customer Gen1 to Out-of-Place Cloud Upgrade to Exadata Database Service on Cloud@Customer Gen2 Infrastructure	5-329
Scope for Exadata Cloud@Customer Gen1 to Gen2 Out-of-Place Cloud Upgrade	5-329
Hardware and Software Required for Out-of-Place Cloud Upgrade to New Exadata Cloud@Customer Gen2 Infrastructure	5-331
Using Oracle Zero Downtime Migration (ZDM) to Migrate Oracle Databases	5-332

During Out-of-Place Cloud Upgrade to New Exadata Cloud@Customer Gen2 Infrastructure	5-336
Post Out-of-Place Cloud Upgrade to New Exadata Cloud@Customer Gen2 Infrastructure	5-337
Best Practices for Out-of-Place Cloud Upgrade to New Exadata Cloud@Customer Gen2 Infrastructure	5-338

6 Autonomous Database on Exadata Cloud@Customer

What's New in ADB-D on Exadata Cloud@Customer	6-2
Provision for Free Developer Database Instances	6-3
Disable and Enable Automatic Backups	6-4
Ability to Scale Up or Scale Down an Autonomous Exadata VM Cluster	6-4
DST Time-Zone File Update for an Autonomous Container Database (AG	CD) 6-5
Rolling and Non-Rolling Maintenance Methods to Update an Autonomou Database	s Container 6-5
Interactive Visualization of Resource Usage Data	6-5
ECPU Billing Model and CPU Allocation	6-6
X10M System Support	6-7
In-Memory Column Store Support	6-7
Create Autonomous Data Guard Associations and Clone an Autonomous	
Across Control Plane Regions	6-9
Set Compartment Quotas Using Policy Statements	6-9
Base Shape support on Oracle Autonomous Database on Dedicated Exa Infrastructure	adata 6-10
Bring Your Own Certificates (BYOC)	6-10
Use Kerberos to Authenticate Autonomous Database Users	6-10
Long-Term Backup	6-11
Save as Stack	6-11
Autonomous Exadata VM Cluster Node Subsetting	6-12
Choose a Specific Oracle Database Version While Provisioning Autonom Databases (ACDs)	nous Container 6-13
Autonomous Data Guard Enhancements	6-13
Configure Non-Default SCAN Listener Ports	6-14
Clone an Autonomous Database Across Exadata Infrastructure Within th	
Region	6-15
Character Set Selection	6-15
Enhanced Resource Tracking for Autonomous Database	6-15
Longer Database Name	6-18
Create a New Autonomous Database Instance from Backup	6-18
Multiple Autonomous VM Cluster Support	6-19
Automatic Failover with a Standby Autonomous Container Database	6-20
X9M-2 System Support	6-20



Fractional OCPU and GB Storage	6-20
Autonomous Data Guard Enabled Autonomous Database and Oracle Key Vault (OKV)	0 _0
Integration	6-21
Infrastructure Patching	6-21
Access Control List (ACL) to Restrict Access to Autonomous Data Guard Enabled	
Autonomous Databases	6-21
ADB-D on Exadata Cloud@Customer: Monitor Performance with Autonomous Database Metrics	6-22
ADB-D on Exadata Cloud@Customer: Autonomous Data Guard	6-22
Access Control List (ACL) to Restrict Access to Autonomous Databases	6-22
Oracle Key Vault (OKV) Integration	6-23
X8M-2 System Support	6-23
Per-Second Billing for Autonomous Database OCPU Usage	6-23
Oracle Autonomous Database on Oracle Exadata Database Service on	
Cloud@Customer	6-23
Introduction to ADB-D on Exadata Cloud@Customer	6-24
Database System Architecture Overview	6-24
Resource Types	6-24
Deployment Order	6-25
User Roles	6-25
Available Exadata Infrastructure Hardware Shapes	6-26
Managing Autonomous Exadata VM Clusters	6-35
About Autonomous Exadata VM Clusters	6-36
Resource Terminology	6-36
Create an Autonomous Exadata VM Cluster	6-37
View a List of DB Servers on an Exadata Infrastructure	6-41
View a List of Autonomous Exadata VM Clusters	6-42
View Details of an Autonomous Exadata VM Cluster	6-42
Scale Autonomous Exadata VM Cluster Resources	6-45
Schedule Oracle-Managed Infrastructure Updates	6-46
Set the Automatic Maintenance Schedule for Autonomous Exadata VM Cluster	6-46
View or Edit the Time of the Next Scheduled Maintenance for Autonomous	
Exadata VM Cluster	6-47
View the Maintenance History of Autonomous Exadata VM Cluster	6-47
Change the License Type on an Autonomous VM Cluster	6-48
Move an Autonomous Exadata VM Cluster to Another Compartment	6-48
Terminate an Autonomous Exadata VM Cluster	6-48
Using the API to Manage Autonomous Exadata VM Clusters	6-49
Rotate Oracle Database TLS Certificate and Oracle REST Data Services (ORDS) TLS Certificate	6-49
Managing Encryption Keys on External Devices	6-51
About Oracle Key Vault	6-52
Overview of Key Store	6-52



Required IAM Policy for Managing OKV on Oracle Exadata Database Service on Cloud@Customer	6-53
Tagging Resources	6-53
Moving Resources to a Different Compartment	6-54
Setting Up Your Exadata Cloud@Customer to Work With Oracle Key Vault	6-54
Step 1: Create a Vault in OCI Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password	6-55
Step 2: Create a Dynamic Group and a Policy Statement for Key Store to Access Secret in OCI Vault	6-55
Step 3: Create a Dynamic Group and a Policy Statement for Exadata Infrastructure to Key Store	e 6-56
Step 4: Create a Policy Statement for Database Service to Use Secret from OCI Vault Service	6-56
Step 5: Create Key Store	6-57
Managing Your Key Store	6-58
View Key Store Details	6-58
Edit Key Store Details	6-58
Move a Key Store to Another Compartment	6-59
Delete a Key Store	6-59
View Key Store Associated Container Database Details	6-59
Using the API to Manage Key Store	6-59
Managing Autonomous Container Databases	6-60
Create an Autonomous Container Database	6-61
View a List of Autonomous Container Databases	6-65
View the List of Autonomous Container Databases in an Autonomous Exadata VM Cluster	6-65
View the List of Autonomous Container Databases in a Compartment	6-65
View Details of an Autonomous Container Database	6-65
Update Autonomous Container Database Maintenance Preferences	6-66
Schedule a Quarterly Maintenance Update	6-67
Rotate CDB Encryption Key	6-68
Change the Backup Retention Policy of an Autonomous Container Database	6-69
Change the Maintenance Schedule of an Autonomous Container Database	6-69
Restart an Autonomous Container Database	6-70
Move an Autonomous Container Database to Another Compartment	6-71
Terminate an Autonomous Container Database	6-71
Using the API to Manage Autonomous Container Databases	6-72
Managing Autonomous Databases	6-72
Oracle Autonomous Database for Developers	6-73
Create an Autonomous Database	6-75
Manage Access Control List of an Autonomous Database	6-79
View a List of Autonomous Databases	6-80
View Details of an Autonomous Database	6-81



Rotate ADB Encryption Key	6-81
Set the Password of an Autonomous Database's ADMIN User	6-82
Scale the CPU Core Count or Storage of an Autonomous Database, or Enable/Disable or Alter the Percentage of System Global Area (SGA) for IM Column Store	6-82
Enable or Disable Auto Scaling for an Autonomous Database	6-84
Move an Autonomous Database to Another Compartment	6-85
Stop or Start an Autonomous Database	6-85
Restart an Autonomous Database	6-86
Back Up an Autonomous Database Manually	6-86
Create a Long-Term Backup	6-87
View Details and Edit Retention Period of a Long-Term Backup	6-88
Delete a Long-Term Backup	6-88
Restore an Autonomous Database	6-89
Restore from a Backup	6-89
Restore to a Point in Time	6-90
Clone an Autonomous Database	6-90
Clone an Autonomous Database Backup	6-94
Clone a Standby Database	6-98
Clone a Standby Database Backup	6-102
Terminate an Autonomous Database	6-106
API to Manage Autonomous Databases	6-107
Monitor Performance with Autonomous Database Metrics	6-107
View Top Six Metrics for an Autonomous Database	6-108
View Aggregated Metrics for Autonomous Databases in a Compartment	6-108
Autonomous Database Metrics and Dimensions	6-109
Connecting to Autonomous Databases	6-109
Download the Wallet for an Autonomous Database	6-110
Get the APEX and SQL Developer Web URLs for an Autonomous Database	6-111
Authenticate Autonomous Database Users with Kerberos	6-111
About Kerberos Authentication	6-112
Components of the Kerberos Authentication System	6-112
Notes about Kerberos Authentication on Autonomous Database	6-113
Enable Kerberos Authentication on Autonomous Database	6-114
Disable Kerberos Authentication on Autonomous Database	6-115
Patching ADB on Exadata Cloud@Customer Infrastructure	6-116
Overview of ADB on Exadata Cloud@Customer Infrastructure Patching	6-116
Specifying When Maintenance Can Occur	6-116
Specifying What Kind of Patches to Apply	6-117
Migrating On-Premises Databases to Oracle Autonomous Database on Exadata Cloud@Customer	6-117
Using Autonomous Data Guard with Autonomous Database on Exadata Cloud@Customer	6-118
Enabling Autonomous Data Guard on an Autonomous Container Database	6-118



Create an Autonomous Data Guard Enabled Autonomous Container Database	6-120
View Details of a Data Guard Enabled Primary or Standby Autonomous Container	
Database	6-125
Edit Autonomous Container Database Backup Settings	6-125
Convert a Physical Standby ACD to Snapshot Standby ACD	6-127
Convert a Snapshot Standby ACD to Physical Standby ACD	6-128
Rotate CDB Encryption Key	6-129
Managing a Standby Autonomous Container Database	6-129
Perform a Failover to Standby Autonomous Container Database	6-131
Perform a Switchover to Standby or Primary Autonomous Container Database	6-131
Reinstate Data Guard Enabled Standby Autonomous Container Database	6-132
Terminate a Data Guard Enabled Primary Autonomous Container Database	6-133
Terminate a Data Guard Enabled Standby Autonomous Container Database	6-133
Operations Performed Using the APIs	6-133
Enabling Autonomous Data Guard on an Autonomous Database	6-135
View Autonomous Data Guard Enablement	6-135
Create an Autonomous Data Guard Enabled Autonomous Database	6-135
View Details of a Data Guard Enabled Primary or Standby Autonomous Database	6-138
Rotate ADB Encryption Key	6-138
Maintenance Scheduling and Patching Data Guard Enabled Autonomous Container	
Database	6-139
Configure Automatic Maintenance Schedule for a Data Guard Enabled Autonomous Container Database	6-139
View the Next Scheduled Maintenance Run of a Data Guard Enabled Autonomous	
Container Database	6-141
View the Maintenance History of a Data Guard Enabled Autonomous Container	
Database	6-141
Immediately Patch a Data Guard Enabled Autonomous Container Database	6-142
Reschedule or Skip scheduled Maintenance for Data Guard Enabled Autonomous Container Database	6-142
Using Performance Hub	6-143

7 Reference Guides for Exadata Database Service on Cloud@Customer

Using the dbaascli Utility with Exadata Database Service on Cloud@Customer	7-1
About Using the dbaascli Utility on Exadata Database Service on Cloud@Customer	7-2
Creating Oracle Database Using dbaascli	7-3
Listing Available Software Images and Versions for Database	7-3
Creating Oracle Database Home	7-4
Creating Oracle Database In the Specified Oracle Database Home	7-5
Changing the Database Passwords	7-8
To Change the SYS Password for an Exadata Database Service on	
Cloud@Customer Database	7-9



To Change Database Passwords in a Data Guard Environment	7-9
To Change the TDE Wallet Password for an Exadata Database Service on Cloud@Customer Database	7-9
Managing Exadata Database Service on Cloud@Customer Software Images Using the Dbaascli Utility	7-9
Listing Available Software Images and Versions for Database and Grid Infrastructure	7-10
To download a software image	7-11
Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli	7-12
Patching Databases using dbaascli	7-13
Patching Oracle Grid Infrastructure	7-15
Listing Available Software Images and Versions for Database and Grid Infrastructure	7-16
Performing a Precheck Before Patching Databases and Grid Infrastructure	7-17
Resuming or Rolling Back a Patching Operation	7-20
Updating Cloud Tooling Using dbaascli	7-21
Creating a Duplicate Database	7-22
Using dbaascli to Duplicate a Cloud Database	7-22
Duplicate an On-Premises Database	7-23
Release Notes	7-25
Release 24.1.2.0.0 (240327)	7-26
Release 24.1.1.0.0 (240118, 240219)	7-26
Release 23.4.1.0.0 (231219)	7-26
Release 23.3.2.0.0 (231115)	7-26
Release 23.3.1.0.0 (230817, 231020)	7-27
Release 23.2.1.0.0 (230708, 230724)	7-27
Release 23.1.2.0.0 (230411, 230616)	7-28
Release 23.1.1.0.1 (230302)	7-28
Release 22.4.1.0.1 (221214)	7-28
Release 22.3.1.1.0 (221003)	7-29
Release 22.3.1.0.1 (220831)	7-29
Release 22.2.1.1.0 (220713)	7-30
Release 22.2.1.0.1 (220504)	7-31
Release 22.1.1.1.0 (220301)	7-32
Release 22.1.1.0.1 (220223)	7-32
Release 21.4.1.1.0	7-33
Release 21.3.1.2.0	7-34
Release 21.3.1.1.0	7-34
Release 21.3.1.0.1	7-34
Release 21.2.1.x.x	7-35
dbaascli Command Reference	7-36
dbaascli admin updateStack	7-41



dbaascli cpuscale get_status	7-42
dbaascli cpuscale update	7-42
dbaascli cswlib deleteLocal	7-43
dbaascli cswlib download	7-44
dbaascli cswlib listLocal	7-44
dbaascli cswlib showImages	7-45
dbaascli database addInstance	7-46
dbaascli database backup	7-46
dbaascli database bounce	7-48
dbaascli database changepassword	7-49
dbaascli database convertToPDB	7-50
dbaascli database create	7-51
dbaascli database createTemplate	7-52
dbaascli database delete	7-53
dbaascli database deleteInstance	7-54
dbaascli database duplicate	7-55
dbaascli database getPDBs	7-57
dbaascli database modifyParameters	7-57
dbaascli database move	7-58
dbaascli database recover	7-59
dbaascli database runDatapatch	7-60
dbaascli database start	7-61
dbaascli database stop	7-61
dbaascli database upgrade	7-62
dbaascli dataguard prepareStandbyBlob	7-63
dbaascli dataguard updateDGConfigAttributes	7-64
dbaascli dbhome create	7-64
dbaascli dbHome delete	7-65
dbaascli dbhome getDatabases	7-65
dbaascli dbHome getDetails	7-65
dbaascli dbHome patch	7-66
dbaascli diag collect	7-67
dbaascli diag healthCheck	7-68
dbaascli gridHome create	7-69
dbaascli grid configureTCPS	7-69
dbaascli grid patch	7-71
dbaascli grid removeTCPSCert	7-72
dbaascli grid rotateTCPSCert	7-72
dbaascli grid upgrade	7-73
dbaascli job getStatus	7-74
dbaascli patch db apply	7-75



dbaascli patch db prereq	7-76
dbaascli pdb backup	7-76
dbaascli pdb bounce	7-77
dbaascli pdb close	7-78
dbaascli pdb getConnectString	7-79
dbaascli pdb create	7-80
dbaascli pdb delete	7-81
dbaascli pdb getDetails	7-82
dbaascli pdb list	7-83
dbaascli pdb localClone	7-83
dbaascli pdb open	7-84
dbaascli pdb recover	7-85
dbaascli pdb refresh	7-86
dbaascli pdb relocate	7-87
dbaascli pdb remoteClone	7-88
dbaascli system getDBHomes	7-90
dbaascli system getGridHomes	7-91
dbaascli tde addSecondaryHsmKey	7-91
dbaascli tde changePassword	7-92
dbaascli tde enableWalletRoot	7-92
dbaascli tde encryptTablespacesInPDB	7-93
dbaascli tde fileToHsm	7-94
dbaascli tde getHsmKeys	7-95
dbaascli tde getMkidForKeyVersionOCID	7-96
dbaascli tde getPrimaryHsmKey	7-96
dbaascli tde hsmToFile	7-97
dbaascli tde listKeys	7-98
dbaascli tde removeSecondaryHsmKey	7-99
dbaascli tde rotateMasterKey	7-100
dbaascli tde setKeyVersion	7-100
dbaascli tde setPrimaryHsmKey	7-101
dbaascli tde status	7-102
Monitoring and Managing Exadata Storage Servers with ExaCLI	7-103
About the ExaCLI Command	7-103
Exadata Storage Server Username and Password	7-103
ExaCLI Command	7-104
Connecting to a Storage Server with ExaCLI	7-110
Monitor Metrics for VM Cluster Resources	7-110
Prerequisites for Using Metrics	7-111
View Metrics for a VM Cluster	7-112
View Metrics for a Database	7-113



View Metrics for VM Clusters in a Compartment	7-114
View Metrics for Databases in a Compartment	7-115
Manage Oracle Trace File Analyzer	7-116
Metrics for Exadata Database Service on Cloud@Customer in the Monitoring Service	7-117
Oracle Exadata Database Service on Cloud@Customer Events	7-122
About Event Types on Exadata Cloud@Customer	7-123
Exadata Infrastructure Event Types	7-123
VM Cluster Network Event Types	7-126
VM Cluster Event Types	7-127
Backup Destination Event Types	7-129
Database Node Event Types (Cloud@Customer)	7-130
Database Home Event Types (Cloud@Customer)	7-131
Database Event Types (Cloud@Customer)	7-132
Database and Grid Infrastructure Patching Event Types	7-133
Autonomous VM Cluster Event Types	7-136
Autonomous Container Database Event Types	7-140
Autonomous Database Event Types	7-141
Data Guard Event Types	7-142
Autonomous Data Guard Association Event Types	7-145
Key Store Event Types	7-149
Exadata Cloud@Customer Infrastructure Maintenance Event Types	7-150
Storage Expansion Event Types	7-168
Database Software Images Event Types	7-170
Database Upgrade Event Types	7-179
Pluggable Database Event Types	7-186
VM Node Subsetting Event Types	7-197
Database Service Events	7-201
Overview of Database Service Events	7-201
Receive Notifications about Database Service Events	7-204
Database Service Event Types	7-205
Temporarily Restrict Automatic Diagnostic Collections for Specific Events	7-217
Remediation	7-221
Compute Expansion Event Types	7-238
Serial Console Connection Event Types	7-242
Viewing Audit Log Events	7-247
Serial Console History Event Types	7-248
Interim Software Updates Event Types	7-253
Policy Details for Exadata Database Service on Cloud@Customer	7-259
About Resource-Types	7-260
Resource-Types for Exadata Database Service on Cloud@Customer	7-260
Supported Variables	7-261



Details for Verb + Resource-Type Combinations	7-262
Database-Family Resource Types	7-263
exadata-infrastructures	7-263
vmcluster-networks	7-265
vmclusters	7-266
backup-destinations	7-267
db-nodes	7-268
db-homes	7-269
databases	7-270
backups	7-271
database-software-image	7-271
autonomous-databases	7-272
autonomous-backups	7-273
autonomous-container-databases	7-274
autonomous-vmclusters	7-275
autonomousContainerDatabaseDataguardAssociations	7-276
AutonomousDatabaseDataguardAssociation	7-277
autonomous-virtual-machine	7-277
key-stores	7-278
pluggable-databases (PDBs)	7-279
dbServers	7-281
dbnode-console-connection	7-281
dbnode-console-history	7-282
oneoffPatch	7-282
Permissions Required for Each API Operation	7-283
Managing Exadata Resources with Oracle Enterprise Manager Cloud Control	7-290
Overview of Oracle Enterprise Manager Cloud Control	7-290
Features of Enterprise Manager Cloud Control	7-290
Security Guide for Exadata Database Service on Cloud@Customer Systems	7-291
Security Configurations and Default Enabled Features	7-291
Responsibilities	7-292
Guiding Principles Followed for Security Configuration Defaults	7-292
Security Features	7-293
Guest VM Default Fixed Users	7-295
Guest VM Default Security Settings	7-297
Guest VM Default Processes	7-298
Guest VM Network Security	7-298
Additional Procedures for Updating Security Posture	7-301
Customer Responsibilities	7-301
Enabling Additional Security Capabilities	7-302
Troubleshooting Exadata Database Service on Cloud@Customer Systems	7-306



Patching Failures on Exadata Database Service on Cloud@Customer Systems	7-306
Determining the Problem	7-307
Troubleshooting and Diagnosis	7-307
Obtaining Further Assistance	7-308
Collecting Cloud Tooling Logs	7-309
Collecting Oracle Diagnostics	7-309
VM Operating System Update Hangs During Database Connection Drain	7-309
Adding a VM to a VM Cluster Fails	7-311
Nodelist is not Updated for Data Guard-Enabled Databases	7-311
CPU Offline Scaling Fails	7-311
Standby Database Fails to Restart After Switchover in Oracle Database 11g Oracle Data Guard Setup	7-312
Using Custom SCAN Listener Port With Data Guard On Disaster Recovery Network	
Causes Data Guard Association Verification Failures	7-313



1 Exadata Database Service on Cloud@Customer Overview

This topic is an overview of the 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 Exadata Database Service on Cloud@Customer provides a variety of management interfaces to fit your use case and automation needs.

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 related to the ADB-D service on Exadata Cloud@Customer please refer to Autonomous 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



Related Topics

Autonomous 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.

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

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.

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

- 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

- System Configuration Options for Oracle Exadata Cloud@Customer Exadata Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X10M.
- 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 Cloud@Customer

Exadata Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X10M.

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

- **X10M-Standard:** Starts with two standard memory (1390 GB) database servers and three Oracle Exadata storage servers.
- **X10M-Large:** Starts with two large memory (2090 GB) database servers and three Oracle Exadata storage servers.
- X10M-Extra Large: Starts with two 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.

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

Oracle Exadata X10M System Model Specifications

Review the technical specifications of X10M system model.

Table 1-1 Oracle Exadata X10M System Model Specifications

Property	Quarter Rack X10M (min config)
Number of DB servers per System	2



Property	Quarter Rack X10M (min config)
Number of Storage Servers per System	3
Total Usable Cores in DB Servers per System	380
Total Memory Available for VMs (GB)	• X10M-Standard: 2780
	• X10M-Large: 4180
	• X10M-Extra Large: 5600
Max Usable Local Storage Per DB Server (GB)	2243
Max Usable Local Storage Per VM for /u02 (GB)	900
VM Image capacity per VM (GB)	184
Max Number of VM Clusters per System*	12
Maximum number of VMs per DB server	8
Total Flash Capacity (TB)	81
Total Usable Disk Storage Capacity (TB)	240

Table 1-1	(Cont.) Oracle Exadat	a X10M System Model Specifications
-----------	-----------------------	------------------------------------

* A maximum of 24 VM Clusters can be created on a single system that contains greater than 2 DB servers. For more information, see Estimating How Much Local Storage You Can Provision to 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.
- 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 1-2 Oracle Exadata X9M-2 System Model Specifications

Property	Base Rack	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



Property	Base Rack	Quarter Rack	Half Rack	Full Rack
Total Memory Available for VMs (GB)	656	2780	5560	11120
Max Usable Local Storage Per DB Server (GB)	1077	2243	2243	2243
Max Usable Local Storage Per VM for /u02 (GB)	892	900	900	900
VM Image capacity per VM (GB)	184	184	184	184
Max Number of VM Clusters per System	4	12	24	24
Max number of VMs per DB server	4	8	8	8
Total Flash Capacity (TB)	38	76	153	307
Total Usable Disk Storage Capacity (TB)	73	190	381	763

Table 1-2 (Cont.) Oracle Exadata X9M-2 System Model Specifications

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

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.

Property	Base Rack	Quarter Rack	Half Rack	Full Rack
Number of DB servers per System	2	2	4	8

Table 1-3 Oracle Exadata X8M-2 System Model Specifications



Property	Base Rack	Quarter Rack	Half Rack	Full Rack
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 Local Storage Per VM for /u02 (GB)	900	900	900	900
VM Image capacity per VM (GB)	160	160	160	160
Max Number of VM Clusters per System	4	8	16	16
Max number of VMs per DB server	4	8	8	8
Total Flash Capacity (TB)	38	76	153	307
Total Usable Disk Storage Capacity (TB)	73	148	297	595

Table 1-3 (Cont.) Oracle Exadata X8M-2 System Model Specifications

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

Oracle Exadata X8-2 System Model Specifications

Review the technical specifications of available Exadata System Shapes.

Table 1-4	Oracle Exadata X	(8-2 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



Property	Base System	Quarter Rack	Half Rack	Full Rack
Total Memory Available for VMs (GB)	720	1440	2880	5760
Max Usable Local Storage Per DB Server (GB)	1237	1037	1037	1037
Max Usable Local Storage Per VM for /u02 (GB)	1100	900	900	900
VM Image capacity per VM (GB)	137	137	137	137
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	76	153	307
Total Usable Disk Storage Capacity (TB)	73	148	297	595

Table 1-4 (Cont.) Oracle Exadata X8-2 System Model Specifications

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

Oracle Exadata X7-2 System Model Specifications

Review the technical specifications of available Exadata System Shapes.

Table 1-5 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



Property	Base System	Quarter Rack	Half Rack	Full Rack
Max Usable Local Storage Per VM for /u02 (GB	1100	1100	1100	1100
VM Image capacity per VM (GB)	137	137	137	137
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	76	153	307
Total Usable Disk Storage Capacity (TB)	42	105	211	423

Table 1-5 (Cont.) Oracle Exadata X7-2 System Model Specifications

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

Exadata Cloud Management Interfaces

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.

- **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 Exadata Cloud@Customer Infrastructure
 Oracle Exadata Database Service on Cloud@Customer uses the same API as Oracle
 Cloud Infrastructure.
- Using the API to Manage Exadata Cloud@Customer Backup Destinations Review the list of API calls to manage your Exadata Database Service on Cloud@Customer backup destinations.
- Using the API to Manage Exadata Cloud@Customer VM Clusters Review the list of API calls to manage your Exadata Database Service on Cloud@Customer VM cluster networks and VM clusters.
- Using the API to Create Oracle Database Home on Exadata Cloud@Customer To create an Oracle Database home, review the list of API calls.
- Using the API to Manage Oracle Database Home on Exadata Database Service on Cloud@Customer
 Review the list of API calls to manage Oracle Database home

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 Exadata Database Service on Cloud@Customer System
 Learn how to use the API to manage Data Guard associations on an 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
- **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 Exadata Database Service on Cloud@Customer

You can use the dbaascli utility to perform various database lifecycle and administration operations on Exadata 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 the bkup_api Utility You can use the bkup_api utility 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.



2

What's New in Oracle Exadata Database Service on Cloud@Customer Gen2

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.

Note:

For information on new features for Autonomous Database on Oracle Exadata Cloud@Customer see What's New in ADB-D on Oracle Exadata Cloud@Customer

- Oracle Database 23ai 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 (domU) 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



- 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
- 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

Oracle Database 23ai on Oracle Exadata Database Service on Cloud@Customer

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

Oracle Database 23ai 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 23ai databases.

Note:

Currently, Grid Infrastructure upgrade from 19c to 23ai is not supported for single node VM clusters.

Limitations:

- When planning to upgrade your Grid Infrastructure to 23ai, make sure that for each ASM diskgroup, compatible.rdbms has a value set to 19.0.0.0 and later.
- You can only restore a 23ai database in the same DB home (in-place).
- When you configure the Data Guard association for 23ai databases, the primary and standby databases must be on the same major release version while the standby database can be on a higher minor version.

Related Topics

 Using the Console to Create Oracle Database Home on Exadata Database Service on Cloud@Customer
 To create on Oracle Database home in an existing VM eluster with the Console home

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 VM Cluster
 To create your VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Change the Home Region of ExaDB-C@C Infrastructure

- Services: Database
- 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 ExaDB-C@C 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.

- 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 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.
- Resource-Types for Exadata Database Service on Cloud@Customer Review the list of resource-types specific to Exadata Database Service on Cloud@Customer.
- dbnode-console-history Review the list of permissions and API operations for dbnode-console-history resourcetype.

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 a VM Cluster
 To create your VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

VM Cluster on a Single VM

- Service: 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 Exadata Database Service on
 Cloud@Customer.

Enhancement to Pluggable Database (PDB) Management

- Service: Database
- Release Date: October 11, 2023

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

- Manage Pluggable Databases on Exadata Database Service on Cloud@Customer Learn to manage pluggable databases on Exadata Cloud@Customer.
- Permissions Required for Each API Operation Review the list of API operations for Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.



- pluggable-databases (PDBs) Review the list of permissions and API operations for pluggable-databases resource-type.
- Pluggable Database Event Types
 Review the list of event types that Pluggable Databases emit.

Manage Administrator (SYS User) and TDE Wallet Passwords

- Service: 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 (domU) Operating System to Oracle Linux 8

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

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

- Updating Guest VM Operating System
 Learn to update the operating system image on Exadata 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.

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 Cloud@Customer Exadata Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X10M.



- Oracle Exadata X10M System Model Specifications Review the technical specifications of X10M system model.
- Estimating How Much Local Storage You Can Provision to 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 Exadata Cloud@Customer Infrastructure To create your Exadata 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.

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 Exadata Database Service on Cloud@Customer Review the list of resource-types specific to Exadata Database Service on Cloud@Customer.
- db-nodes Review the list of permissions and API operations for db-nodes resource-type.
- Permissions Required for Each API Operation Review the list of API operations for Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- dbnode-console-connection Review the list of permissions and API operations for dbnode-console-connection resource-type.
- 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.

- 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 Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- oneoffPatch
- 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.

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 Exadata Cloud@Customer Infrastructure To create your Exadata 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.
- Metrics for Exadata Database Service on Cloud@Customer in the Monitoring Service Review the metrics emitted by the 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 Exadata Database Service on Cloud@Customer System Learn to enable Data Guard association 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.

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 Exadata Database Service on Cloud@Customer System Learn to enable Data Guard association 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
- 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).

- Using the Console to Create Exadata Cloud@Customer Infrastructure To create your Exadata 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 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.
- 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 Exadata 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.
- 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 a VM Cluster
 To create your 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 Exadata Cloud@Customer Infrastructure To create your Exadata 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 Exadata Database Service on Cloud@Customer Infrastructure
 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.
- 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 Exadata 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.

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 of Exadata Cloud@Customer Infrastructure Learn how to view the maintenance history for an Exadata 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 with CVSS scores greater than or equal to 7.
- View or Edit a Scheduled Security Maintenance for Exadata Cloud@Customer Infrastructure

Learn how to view and edit the next scheduled security maintenance.

- Exadata Cloud@Customer Infrastructure Maintenance Event Types Review the list of event types that Exadata Cloud@Customer 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 optout 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 Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure

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 a Scheduled Quarterly Maintenance for Exadata Cloud@Customer Infrastructure

Learn how to view and edit the time of the next scheduled maintenance.

- View the Maintenance History of Exadata Cloud@Customer Infrastructure Learn how to view the maintenance history for an Exadata Cloud@Customer Infrastructure.
- View and Edit Quarterly Maintenance While Maintenance is In Progress or 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.

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
- 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 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

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

- 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 tfact1 blackout command to temporarily suppress automatic diagnostic collections.
- Using the Console to Create a VM Cluster
 To create your 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



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 Exadata Cloud@Customer Infrastructure To create your Exadata 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 Exadata Cloud@Customer Infrastructure To create your Exadata Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.



- View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure 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
- Using the API to Manage Exadata Cloud@Customer Infrastructure Maintenance Controls Oracle Exadata 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
 Exadata 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 Cloud@Customer Infrastructure Maintenance Event Types Review the list of event types that Exadata Cloud@Customer Infrastructure Maintenance emits.

Manage Pluggable Databases on Exadata Cloud@Customer

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

those already scheduled.

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

Related Topics

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

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



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 Exadata Database Service on Cloud@Customer System
 Learn to enable Data Guard association 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 Exadata Database Service on Cloud@Customer System Learn to enable Data Guard association 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.

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 a VM Cluster
 To create your 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 Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- dbServers
 Review the list of permissions and API operations for dbServers resource-type.
- 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 Cloud@Customer Exadata Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X10M.
- Oracle Exadata X9M-2 System Model Specifications
 Review the technical specifications of available Exadata System Shapes.
- Estimating How Much Local Storage You Can Provision to Your VMs
- 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 Exadata Database Service on Cloud@Customer System
 Learn to enable Data Guard association 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*.

Related Topics

 Upgrading Oracle Grid Infrastructure on an Exadata Cloud@Customer VM Cluster Learn to upgrade Oracle Grid Infrastructure on an Exadata 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 Exadata 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 Oracle Database 19c (Long Term Release) 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

- 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 Exadata 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 Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.
- exadata-infrastructures Review the list of permissions and API operations for exadata-infrastructures resource-type.
- Storage Expansion Event Types Review the list of event types that storage expansion emits.

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 Oracle Database Software Images Learn about Database Software Image resource type and how you can use it to create Oracle Databases and Oracle Database Homes and to patch databases.
- Using the Console to Create Oracle Database Home on 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 a Patch 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
- Release Notes: Exadata Cloud@Customer: Customer 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
- Managing Infrastructure Maintenance Contacts

Related Topics

- Using the Console to Create Exadata Cloud@Customer Infrastructure To create your Exadata Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Managing Infrastructure Maintenance Contacts Learn to manage your Exadata infrastructure maintenance contacts.

X8M-2 System Support

- Services: Database
- Release Date: August 18, 2020
- Release Notes: Exadata Cloud@Customer: X8M-2 System Support

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

- System Configuration Options for Oracle Exadata Cloud@Customer Exadata Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X10M.
- 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
- Release Notes: Exadata Cloud@Customer: Enable and Manage Data Guard
 Associations

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
- 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 Exadata Database Service on Cloud@Customer Learn to configure and manage Data Guard associations in your VM cluster.
- Using the API to Manage Data Guard Associations on an Exadata Database Service on Cloud@Customer System
 Learn how to use the API to manage Data Guard associations on an Exadata Database Service on Cloud@Customer system.
- Permissions Required for Each API Operation Review the list of API operations for 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 Guard associations emit.

Oracle Exadata Cloud@Customer Deployment Assistant

- Services: Database
- Release Date: August 08, 2020
- Release Notes: Exadata Cloud@Customer: Deployment Assistant

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.

For more information, see Oracle Exadata Cloud@Customer Deployment Assistant



Oracle Grid Infrastructure and Oracle Database Patching

- Services: Database
- Release Date: July 28, 2020
- Release Notes: Exadata Cloud@Customer: Oracle Grid Infrastructure and Oracle
 Database Patching

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 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 Exadata Database Service on Cloud@Customer Systems These topics cover some common issues you might run into and how to address them.
- 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
- Release Notes: Exadata Cloud@Customer: Per-Second Billing for OCPU Usage

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
- Release Notes: Exadata Cloud@Customer: Shared Database Home Resource Tags



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 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
- Release Notes: Exadata Cloud@Customer: Create and Manage Multiple Virtual Machines per Exadata System (MultiVM)

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

Related Topics

- Using the Console to Create a VM Cluster To create your 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 Exadata Database Service on Cloud@Customer Gen2, 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



Release Notes: Exadata Cloud@Customer: Scale OCPUs Without Cloud Connectivity

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 Exadata Database Service on Cloud@Customer

You can use the dbaascli utility to perform various database lifecycle and administration operations on Exadata 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

- Services: Database
- Release Date: May 7, 2020
- Release Notes: Exadata Cloud@Customer: Configure Oracle Database Character Set and National Character Set

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
- Release Notes: Exadata Cloud@Customer: Specify a Time Zone While Provisioning Oracle Exadata Cloud@Customer Infrastructure

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 Exadata Cloud@Customer Infrastructure To create your Exadata 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
- Release Notes: Exadata Cloud@Customer: Shared Database Homes

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

X7-2 System Support

- Services: Database
- Release Date: March 19, 2020
- Release Notes: Exadata Cloud@Customer: X7-2 System Support

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.



3

Preparing for Exadata Database Service on Cloud@Customer

Review OCI as well as the site, network and storage requirements to prepare and deploy 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
 Review the storage requirements for the ASM disk groups and VM file systems to plan the best options for your enterprise needs.
- Checklists for Exadata Database Service on Cloud@Customer Deployments To determine your readiness for an 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. After Oracle creates your tenancy in OCI, an administrator at your company will need to perform some set up tasks and establish an organization plan for your cloud resources and users. The information in this topic will help you get started.



- Managing Regions
 This topic describes the basics of managing your region subscriptions.
- Managing Compartments This topic describes the basics of working with compartments.
- Managing Users This topic describes the basics of working with users.
- Managing Groups
 This topic describes the basics of working with groups.
- Required IAM Policy for Exadata 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 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.



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 Exadata Database Service on Cloud@Customer Learn to write policies to control access to 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.



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)	
X9M Base System	436.7 kg	962.7 lbs	
X9M Quarter Rack	449.7 kg	991.4 lbs	
X9M Half Rack	592.3 kg	1305.7 lbs	
X9M Full Rack	886.5 kg	1954.3 lbs	
X9M Expansion Rack min	374.8 kg	826.4 lbs	
X9M Expansion Rack max	904.7 kg	1994.5 lbs	
X10M Elastic Rack min	469.8 kg	1035.8 lbs	
X10M Elastic Rack max	900.1 kg	1984.3 lbs	
X10M Expansion Rack min	416.1 kg	917.3 lbs	
X10M Expansion Rack max	963.7 kg	2124.5 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 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.

Model	Minimum PDU Rating (kVA)
Base System	15 kVA
Quarter Rack	15 kVA
Half Rack	15 kVA
Full Rack	22 kVA

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.

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.

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 Ranges



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.

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 Exadata Database Service on Cloud@Customer Before you deploy Exadata Database Service on Cloud@Customer, ensure that your data center network meets requirements.
- IP Addresses and Subnets for 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 Exadata Database Service on Cloud@Customer Ensure that your Exadata Database Service on Cloud@Customer system meets control plane server and database server uplink requirements.
- Network Cabling for 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:

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. <i>oci_regi on</i> .oci.oraclecloud .com
Secure Tunnel Service for remote Oracle operator access	443 outbound	Use these URL formats, replacing <i>oci_region</i> with your region:
		https:// mgmthel.exacc. <i>oci_ region</i> .oci.oraclec loud.com
		https:// mgmthe2.exacc. <i>oci_</i> <i>region</i> .oci.oraclec loud.com

Table 3-2	Ports to O	pen for Control	I Plane Connectivity	1
-----------	------------	-----------------	----------------------	---



Description / Purpose	Open Port	Location
Object Storage Service to retrieve system updates, infrastructure monitoring, and log collection	443 outbound	Use this URL format, replacing <i>oci_region</i> with your region:
		https:// objectstorage. <i>oci_</i> <i>region</i> .oraclecloud .com
		<pre>https:// swiftobjectstorage .oci_region.oracle cloud.com</pre>
		<pre>https:// *.objectstorage.oc i_region.oci.custo mer-oci.com</pre>
Monitoring Service to record and process Infrastructure Monitoring Metrics (IMM)	443 outbound	Use this URL format, replacing <i>oci_region</i> with your region:
		<pre>https://telemetry- ingestion.oci_regi on.oraclecloud.com</pre>
Identity Service for Authorization and Authentication	443 outbound	Use this URL format, replacing <i>oci_region</i> with your region:
		https:// identity. <i>oci_regio</i> n.oraclecloud.com
		https:// auth. <i>oci_region</i> .or aclecloud.com

Table 3-2 (Cont.) 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 <i>oci_region</i> with your region:
		https:// wsshe.adbd-
		exacc.oci region.o
		ci.oraclecloud.com
Logging Service	443 outbound	Use this URL format replacing <i>oci_region</i> with your region:
		https://
		frontend.logging.a
		dl. <i>oci_region</i> .orac leiaas.com
		https://
		frontend.logging.a d2. <i>oci region</i> .orac
		leiaas.com
		https://
		frontend.logging.a d3. <i>oci region</i> .orac
		leiaas.com
		https://
		controlplane.loggi ng.adl. <i>oci region</i> .
		oracleiaas.com
		https://
		controlplane.loggi ng.ad2. <i>oci region</i> .
		oracleiaas.com
		https://
		<pre>controlplane.loggi ng.ad3.oci_region.</pre>
		oracleiaas.com

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



Description / Purpose	Open Port	Location
Resource Principal based authentication and Autonomous Database service delivery	443 outbound	Use this URL format, replacing <i>oci_region</i> with your region:
		https://
		database. <i>oci_regio</i> n.oraclecloud.com
VM console	443 outbound	Use this URL format, replacing <i>oci_region</i> with your region:
		<pre>https:// console1.exacc.oci _region.oci.oracle cloud.com</pre>
		https:// console2.exacc. <i>oci</i> _ <i>region</i> .oci.oracle cloud.com

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

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. Customer HTTP proxy, for example, passive/ corporate proxy supported for the Control Plane Server connection to OCI. Customer HTTPS, challenge proxy, and traffic inspection 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 (NETO). 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 NETO 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 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:

Description / Purpose	Open Port	Location
Object Storage Service for cloud-based database backups (Optional)	443 outbound	Use this URL format, replacing <i>oci_region</i> with your region:
		https:// objectstorage. <i>oci_reg ion</i> .oraclecloud.com
		<pre>https:// swiftobjectstorage.oc i_region.oraclecloud. com</pre>

Table 3-3 Ports to Open for Backup Network

Disaster Recovery Network (authorized customers only)

Note:

Only certain authorized customers will have this additional network due to special circumstances, your system may not be so-equipped.

This network is configured similarly to the client access network, however it transmits Data Guard traffic only. If your system is so equipped, rather than sending Data Guard traffic over the client network, this Data Guard traffic will instead go over this Disaster Recovery network using Single Client Access Name (SCAN) and Oracle Real Application Clusters (Oracle RAC) Virtual IP (VIP) interfaces.

The Disaster Recovery network uses a pair of network interfaces on each database server, which are connected to the customer network.

Related Topics

Object Storage Service API



- Monitoring API
- Identity and Access Management Service API

Data Center Network Services for Exadata Database Service on Cloud@Customer

Before you deploy 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 Exadata Database Service on Cloud@Customer network interfaces. Oracle requires that you register the host names and IP addresses for the 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 Exadata Database Service on Cloud@Customer to ensure coverage in case a server is unavailable.

Network Time Protocol (NTP) Services Requirements

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 Exadata Database Service on Cloud@Customer to ensure coverage in case a server is unavailable.

IP Addresses and Subnets for 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 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 Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth.



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 Exadata Cloud@Customer Infrastructure.

Host Name and IP Address Requirements for Oracle Cloud@Customer

To connect to your corporate network, 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 Exadata Database Service on Cloud@Customer, work with your network team to reserve enough IP addresses to accommodate any future growth.



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:	Maximum CIDR block prefix length:
	/28	/27
	Recommended CIDR block prefix length:	Recommended CIDR block prefix length:
	/27	/26
Backup network	Maximum CIDR block prefix length:	Maximum CIDR block prefix length:
	/29	/28
	Recommended CIDR block prefix length:	Recommended CIDR block prefix length:
	/28	/27
(For authorized customers only) Disaster Recovery network	Maximum CIDR block prefix length:	Maximum CIDR block prefix length:
	/28	/27
	Recommended CIDR block prefix length:	Recommended CIDR block prefix length:
	/27	/26

Uplinks for Exadata Database Service on Cloud@Customer

Ensure that your 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 the Exadata Cloud@Customer service.

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.



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 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 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 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.







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.







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

Review the storage requirements for the ASM disk groups and VM file systems to plan the best options for 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 to Your VMs

 About Storage Configuration for Oracle Exadata Database Service on Cloud@Customer

As part of configuring each 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 Exadata Cloud@Customer

About Storage Configuration for Oracle Exadata Database Service on Cloud@Customer

As part of configuring each 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 *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 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%

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%

ORACLE[°]

- RECO Disk Group: 50%
- SPARSE **Disk Group 15%**

Virtual Machine File System Structure for Exadata Cloud@Customer

Exadata Cloud@Customer X8M systems 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	1
/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 to Your VMs

Checklists for Exadata Database Service on Cloud@Customer Deployments

To determine your readiness for an Exadata Database Service on Cloud@Customer deployment, review the deployment checklists.

- System Components Checklist for Exadata Database Service on Cloud@Customer Use this checklist to ensure that the system component considerations are addressed.
- Data Center Room Checklist for 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 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 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 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)?



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?



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 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?
- 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?

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)?
- 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 InfiniBand 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 IOA 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?

4

Getting Started with Exadata Database Service on Cloud@Customer Deployment

After completing the preparation tasks in Preparing for Exadata Database Service on Cloud@Customer, get started with deploying your 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.
- Oracle Exadata Cloud@Customer Deployment Assistant
 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.
- 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 Exadata Database Service on Cloud@Customer Learn how to provision an Exadata Database Service on Cloud@Customer system.
- Configuring Oracle-Managed Infrastructure Maintenance
 Oracle performs the updates to all of the Oracle-managed infrastructure components on
 Exadata Cloud@Customer.
- Creating First VM Cluster Network on Exadata Cloud@Customer Follow the steps in the following sections to create your first VM Cluster Network.
- Provisioning the First VM Cluster on an Exadata Cloud@Customer System Learn how to create the first VM cluster on your Exadata Cloud@Customer system.
- Creating Database Backup Destinations for Exadata Cloud@Customer Exadata Cloud@Customer provides a backup facility, which you can configure individually on each database.
- Creating First Database Home on an 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 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 Exadata Cloud@Customer.
- Connecting to an Exadata Database Service on Cloud@Customer System
 After deploying your Exadata 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 Exadata Cloud@Customer database using Oracle Net Services (SQL*Net).

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.

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 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 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 Database, and VM Cluster Networks are some of them. Tags can either be applied while creating the resources or modified later. For example, you can apply tags to an Autonomous Container Database (ACD) while provisioning the ACD or add them later from its **Details** page.

See *How Tagging Works* for more details on using tags. Tagging integrates with Oracle Cloud Infrastructure authorization system. You can use IAM policy controls to enable delegation or restriction of tag manipulation. See *Authentication and Authorization* to learn about the permissions required to work with defined and free-form tags. (Required) Enter introductory text here, including the definition and purpose of the concept.



Tip:

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

Your tenancies come with a library of standard tags that would apply to most resources. These tags are currently available as a set of Tag Namespaces that your governance administrators can deploy. OCI best practices recommend applying these tags to all resources a standard tag can be applied to. Besides reporting and governance, OCI service automation can deliver workload-specific optimizations based on standard tag values.

For example, database deployments for the PeopleSoft application require a specific configuration. By setting the appropriate application tag key in the Oracle-ApplicationName tag namespace while deploying an Autonomous Database, can ensure that the database is configured ready for the particular application, for example, PeopleSoft out of the box.

Figure 4-1 Tagging Example

Encryption Key	Tags					
Optional tags to orga	anize and trac	k resources in yo	our tenancy. <u>How do I u</u>	ise tags?		
Tag Namespace			Tag Key		Tag Value	
Oracle-Application	Name	\$	PeopleSoft	\$	9.2	≎ ×
					9.2	Tag
					9.1	lag

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.



Tag Key	Tag Value Options	Description
OracleStandard.Criticalit Y	ExtremeHighMediumLow	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.Criticalit y set to Extreme or High may require the highest availability SLA and may need to be configured with Autonomous Data Guard.
OracleStandard.Environmen t	 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.Sensitivit Y	 Public Internal Sensitive Highly Sensitive Extremely Sensitive 	An application or database classification tag. OracleStandard.Sensitivit y set to Highly Sensitive may indicate that an access control list or certain Network Security Group (NSG) enforcement is mandatory to restrict access.
OracleStandard.Regulation	Refer to <i>List of Compliance</i> <i>Regulations</i> for values.	 Denotes one or more compliance regulations that a resource must adhere to. Tag administrators may add values to the list from the OCI Governance and Administration console. Refer to Using Predefined Values for more details.
OracleStandard.TargetUser s	 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.

Table 4-1	Oracle St	tandard Tags
-----------	-----------	--------------



Tag Key	Tag Value Options	Description
OracleStandard.EndUserCou nt	 1 10 100 1000 10000 100000 1000000 1000000 1000000 	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.comor 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	12345WebMarketing	Freeform field for cost center.
OracleStandard.RecoveryTi meObjectiveMinutes	0-10080	Time in minutes. Denotes the maximum time within which the resource is required to recover from a failure.
OracleStandard.RecoveryPo intObjectiveMinutes	0-1440	Time in minutes. Maximum data loss tolerance for a data store resource such as a database or a storage device.

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

List of Compliance Regulations

Table 4-2	List of Compliance Regulations
-----------	--------------------------------

Regulation	Description	
PCIDSS	Payment Card Industry Data Security Standard	
НІРАА	Health Insurance Portability and Accountability Act	
ISO	International Standards Organization	
SOC1	System and Organization Controls 1	
SOC 2	System and Organization Controls 2	



Regulation	Description
FedRamp	Federal Risk and Authorization Management Program
GLBA	Gramm–Leach–Bliley Act
ССРА	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
СОРРА	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

 Table 4-2
 (Cont.) List of Compliance Regulations

Oracle Application Name Tags

Table 4-3 Oracle Application Name Tags

Tag Key	Tag Value Options	Description
Hyperion	• 11.2	Denotes the version of the
	• 11.1	Hyperion application.
JD Edwards	• 9.2	Denotes the version of the JD
	• 9.1	Edwards application.
	• 9.0	
Oracle_E-Business_Suite	• 12.2	Denotes the version of the
	• 12.1	Oracle E-Business Suite
	• 12.1	application.
	• 11i	
PeopleSoft	• 9.2	Denotes the version of the
	• 9.1	PeopleSoft application.
Siebel	• 8.2	Denotes the version of the Siebel
	• 8.1	application.



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.

Table 4-3 (Cont.) Oracle Application Name Tags

Related Topics

- Cloud Advisor Overview
- Oracle Autonomous Database Dedicated for Fleet Administrators Workshop
- How Tagging Works
- Authentication and Authorization
- Managing Tag Namespaces
- Using Predefined Values

Oracle Exadata Cloud@Customer Deployment Assistant

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.

- Using Oracle Exadata Cloud@Customer Deployment Assistant
 - Oracle Exadata Cloud@Customer Deployment Assistant gathers your configuration details and creates the Oracle Exadata Cloud@Customer infrastructure configuration file. The configuration file drives the automated installation and configuration processes for Oracle Exadata Cloud@Customer infrastructure.
- Accessing Oracle Exadata Cloud@Customer Deployment Assistant Follow these steps to run the Deployment Assistant.

• Step 1: Pre-Installation

Create Exadata Cloud@Customer infrastructure, VM cluster network, and download the infrastructure configuration file before your engineered system arrives at your premises.

Step 2: Onsite Installation

Ensure that you activate the Exadata Cloud@Customer infrastructure and validate the VM cluster network.

• Step 3: Post-Installation Create a VM cluster, install Oracle Database, and validate your installation before performing any administrative tasks.

Using Oracle Exadata Cloud@Customer Deployment Assistant

Oracle Exadata Cloud@Customer Deployment Assistant gathers your configuration details and creates the Oracle Exadata Cloud@Customer infrastructure configuration



file. The configuration file drives the automated installation and configuration processes for Oracle Exadata Cloud@Customer infrastructure.

Before your engineered system arrives, do the following:

- Work with your network and database administrators to evaluate the current network settings, such as current IP address use and network configuration.
- Define the settings for the rack, such as network configuration and backup method.
- Download the Oracle Exadata Cloud@Customer infrastructure configuration file.

During deployment, if you find any discrepancies at any stage, then click **Close and complete later** to exit Oracle Exadata Cloud@Customer Deployment Assistant. You will lose all of your settings and you will have to start afresh the next time.

Accessing Oracle Exadata Cloud@Customer Deployment Assistant

Follow these steps to run the Deployment Assistant.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Click Exadata Infrastructure.
- 3. Click Deployment Assistant.
- 4. Choose a deployment type.
 - New Deployment: Creates an Exadata Cloud@Customer infrastructure and all the resources needed to create your first Oracle Database.
 Select a Compartment, and then click Continue.
 - **Existing Deployment**: Uses an existing Exadata Cloud@Customer infrastructure and guides you through completing the deployment.
 - a. Select a Compartment.
 - b. Select an Exadata Cloud@Customer infrastructure.
 - c. Select a VM Cluster Network or create one.
 - d. Click Continue.
- 5. If you have created an Exadata Cloud@Customer infrastructure and if it is in **Active** state, then do the following:
 - a. Go the Exadata Cloud@Customer infrastructure details page.
 - b. Click Deployment Assistant.
- 6. If you have created an Exadata Cloud@Customer infrastructure and if it is in **Requires** Activation state, then do the following:
 - a. Go the Exadata Cloud@Customer infrastructure details page.
 - b. Click More Actions and then select Deployment Assistant.

Step 1: Pre-Installation

Create Exadata Cloud@Customer infrastructure, VM cluster network, and download the infrastructure configuration file before your engineered system arrives at your premises.

1. Create Oracle Exadata Database Service on Cloud@Customer infrastructure.



For more information and instructions, see Using the Console to Create Exadata Cloud@Customer Infrastructure.

2. Create VM cluster network.

For more information and instructions, see Using the Console to Create a VM Cluster Network.

 Download the Oracle Exadata Database Service on Cloud@Customer configuration file.

For more information and instructions, see Using the Console to Download a File Containing Configuration Data.

Step 2: Onsite Installation

Ensure that you activate the Exadata Cloud@Customer infrastructure and validate the VM cluster network.

1. Add Infrastructure Contacts.

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

You must add a primary contact to activate your infrastructure. Ensure that you provide the details of the contact that you used while registering the Customer Support Identifier (CSI) associated with this infrastructure, as a primary contact.

For more information and instructions, see Infrastructure Maintenance Contacts

2. Activate Oracle Exadata Database Service on Cloud@Customer infrastructure.

For more information and instructions, see Using the Console to Activate Exadata Database Service on Cloud@Customer Infrastructure.

3. Validate the 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.

For more information and instructions, see Using the Console to Validate a VM Cluster Network.

Step 3: Post-Installation

Create a VM cluster, install Oracle Database, and validate your installation before performing any administrative tasks.

1. Create VM cluster.

For more information and instructions, see Using the Console to Create a VM Cluster.

2. Create Oracle Database.

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

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.

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

Table 4-4 Key Additional Resources

Table 4-5 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 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



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 4-7 Key Additional Resources

Table 4-8 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

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: 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 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.




Figure 4-2 Overview of Order and Deployment Process

Multi-Rack Expansion

- Rack Minimum Configuration
 - 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.
- 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 Exadata 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 Exadata Database Service on Cloud@Customer

Learn how to provision an Exadata Database Service on Cloud@Customer system.

- Using the Console to Create Exadata Cloud@Customer Infrastructure To create your Exadata 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 Exadata Database Service on Cloud@Customer Infrastructure
 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 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 Exadata Cloud@Customer Infrastructure

To create your Exadata Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.

- 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.

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:
 - **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.

• **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.

- **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.
- 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 servers:
 - * **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:

- **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 checkbox is automatically selected. You can also check 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 select this check box. 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.

• Configure the cloud 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.

- 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 /22. The maximum CIDR block /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 4-9X10M (Quarter Rack X10M, Quarter Rack X10M-L, and Quarter RackX10M-XL) CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
6	64	/23	/22
7	45	_	



Table 4-9 (Cont.) X10M (Quarter Rack X10M, Quarter Rack X10M-L, andQuarter Rack X10M-XL) CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
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 4-10 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



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

Table 4-11 X8M CIDR Requirements

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.
- 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 browserdetected 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**.

Network bonding mode

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
 - Configure automatic maintenance Click Edit Maintenance Preferences.

Edit Maintenance Preferences dialog is displayed.

In the Edit Maintenance Preferences dialog, configure the following:



- * 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.
 - * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

Default: 30 minutes

Maximum: 120 minutes

Maintenance schedule:

*

- * **No preference:** The system assigns a date and start time for infrastructure maintenance.
- * **Specify a schedule:** Choose your preferred month, week, weekday, start time, and lead time for infrastructure maintenance.
- * **Lead Time:** Specify the minimum number of weeks ahead of the maintenance event you would like to receive a notification message.

Click Save Changes.

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.

Provide maintenance contacts

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

You can skip adding maintenance contacts while creating your infrastructure. However, you must add a primary contact prior to activating your infrastructure. Ensure that you provide the details of the contact that you used while registering the Customer Support Identifier (CSI) associated with this infrastructure, as a primary contact. Optionally, you can 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. Note that 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.

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 Exadata Infrastructure.

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 Cloud@Customer Exadata Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X10M.
- ExaCC gen2: Troubleshooting VPN/WSS connection from Customer Side
- Resource Tags

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 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.

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 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 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. 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.

b. 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.

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.

c. 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.
- **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 browserdetected 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**.

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 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.

Using the Console to Activate 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 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 Activate on the banner displayed.

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 dialog to upload the activation file, and then click Activate Now.

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 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 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.

- 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.

Configuring Oracle-Managed Infrastructure Maintenance

Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud@Customer.

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 the infrastructure maintenance process and configuring the maintenance controls refer to the following:

 About Oracle Managed Exadata Cloud@Customer Infrastructure Maintenance Updates

Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud@Customer.

- Infrastructure Maintenance Contacts
 Maintenance contacts are required for service request based communications for
 hardware replacement and other maintenance events.
- 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 optout of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal and allows scheduling flexibility to help you plan for them.
- 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.

About Oracle Managed Exadata Cloud@Customer Infrastructure Maintenance Updates

Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud@Customer.

Oracle patches and updates include the physical database server hosts, Exadata Storage Servers, Network Fabric Switches, management switch, power distribution units (PDUs), integrated lights-out management (ILOM) interfaces, and Control Plane Servers. This is referred to as Exadata Cloud@Customer infrastructure maintenance.

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.

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 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.

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.

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 impact older applications not written to handle a rolling instance outage. 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 do not impact VM cluster VM's availability. However, the rolling storage server maintenance can result in reduced IO performance as storage servers are taken offline (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 not under maintenance will minimize (or eliminate) any performance impact.

Note that while databases are expected to be available during the rolling maintenance process, the automated maintenance verifies Oracle Clusterware is running but does not verify that all database services and pluggable databases (PDBs) are available after a server is brought back online. The availability of database services and PDBs after maintenance can depend on the application service definition. For example, a database service, configured with certain preferred and available nodes, may be relocated during the maintenance and wouldn't automatically be relocated back to its original node after the maintenance completes. Oracle recommends reviewing the documentation on *Achieving Continuous Availability for Your Applications* on Exadata Cloud Systems to reduce the potential for impact to your applications. By following the documentation's guidelines, the impact of infrastructure



maintenance will be only minor service degradation as database servers are sequentially updated.

Oracle recommends that you follow the *Maximum Availability Architecture (MAA)* best practices and use Data Guard to ensure the highest availability for your critical applications. For databases with Data Guard enabled, Oracle recommends that you separate the maintenance windows for the infrastructure instances running the primary and standby databases. You may also perform a switchover prior to the maintenance operations for the infrastructure instance hosting the primary database. This allows you to avoid any impact on your primary database during infrastructure maintenance.

Prechecks are performed on the Exadata Cloud@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 10 days prior to the maintenance start and another precheck is run approximately 24 hours prior to maintenance start. If the prechecks identify an issue that requires rescheduling the maintenance notification is sent to the maintenance contacts.

The time taken to update infrastructure components varies depending on the number of database servers and storage servers in the Exadata infrastructure, the maintenance method, and whether custom action has been enabled. The approximate times provided are estimates. Time for custom action, if configured, is not included in the estimates below. Database server maintenance time may vary depending on the time required to shutdown each VM before the update and then start each VM and associated resources after the update of each node before proceeding to the next node. The storage server maintenance time will vary depending on the time required for the ASM rebalance, which is not included in the estimates below. If issues are encountered during maintenance this may also delay completion beyond the approximate time listed. In such a situation, if Oracle cloud operations determine resolution would extend beyond the expected window, they will send a notification and may reschedule the maintenance.

Note:

The timeframes mentioned below can change if Oracle cloud operations determine that additional maintenance work is needed. If additional time is necessary, Oracle will send a customer notification in advance to inform customers that additional time will be required for the next quarterly maintenance window.

Rolling:

- Each database server takes 90 minutes on average.
- Each storage server takes 60 minutes on average.
- Each InfiniBand or RoCE fabric switch takes 30 minutes on average.
- The approximate total time for infrastructure maintenance is as follows:
 - * Base and Quarter Rack (2 Database Servers/3 Storage Servers): Approximately 7 hours
 - 2 Database Servers X 90 = 180 minutes
 - 3 Storage Servers X 60 = 180 minutes



2 InfiniBand or RoCE Fabric Switch X 30 = 60 minutes

- * Half Rack (4 Database Servers/6 Storage Servers): Approximately 13 hours
 4 Database Servers X 90 = 360 minutes
 - 6 Storage Servers X 60 = 360 minutes

2 InfiniBand or RoCE Fabric Switch X 30 = 60 minutes

* Full Rack (8 Database Servers/12 Storage Servers): Approximately 26 hours

8 Database Servers X 90 = 720 minutes

12 Storage Servers X 60 = 720 minutes

2 InfiniBand or RoCE Fabric Switch X 30 = 60 minutes

- Non-Rolling:
 - All database servers take 180 minutes on average.
 - All storage servers take 60 minutes on average.
 - Storage Servers and Database servers are brought back online prior to starting fabric switch maintenance.
 - Network fabric switches are still updated in a rolling method and take 30 minutes each on average.
 - The approximate total time for infrastructure maintenance is 5 hours regardless of shape:
 - * All Database Servers = 180 minutes
 - * All Storage Servers = 60 minutes
 - * 2 InfiniBand or RoCE Fabric Switch = 60 minutes

Related Topics

- Achieving Continuous Availability For Your Applications
- Maximum Availability Architecture (MAA) Best Practices

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.

Related Topics

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



 Managing Infrastructure Maintenance Contacts Learn to manage your Exadata infrastructure maintenance contacts.

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.

For more information, see the following:

- View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure
- View or Edit a Scheduled Quarterly Maintenance for Exadata Cloud@Customer Infrastructure
- View or Edit a Scheduled Security Maintenance for Exadata Cloud@Customer Infrastructure
- View the Maintenance History of Exadata Cloud@Customer Infrastructure
- View and Edit Quarterly Maintenance While Maintenance is In Progress or Waiting for Custom Action
- View or Edit a Scheduled Security Maintenance for Exadata Cloud@Customer Infrastructure
- View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure

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 a Scheduled Quarterly Maintenance for Exadata Cloud@Customer Infrastructure

Learn how to view and edit the time of the next scheduled maintenance.

• View the Maintenance History of Exadata Cloud@Customer Infrastructure Learn how to view the maintenance history for an Exadata Cloud@Customer Infrastructure.

View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure

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 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.
- 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 Maintenance Preferences.

Edit Maintenance Preferences page is displayed.

Note:

Changes made to maintenance preferences apply only to future maintenance, not the maintenance that has already been scheduled. To modify scheduled maintenance, see *View or Edit a Scheduled Maintenance for Exadata Cloud@Customer Infrastructure*.

- 6. On the Edit Maintenance Preferences page, configure the following:
 - 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 nonrolling 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.
 - Default: 30 minutes

Maximum: 120 minutes

- Maintenance schedule:
 - No preference: The system assigns a date and start time for infrastructure maintenance.
 - **Specify a schedule:** Choose your preferred month, week, weekday, start time, and lead time for infrastructure maintenance.



- * Under **Maintenance months**, specify at least one month for each quarter during which Exadata infrastructure maintenance will take place. You can select more than one month per quarter. If you specify a long lead time for advanced notification (for example, 4 weeks), you may wish to specify 2 or 3 months per quarter during which maintenance runs can occur. This will ensure that your maintenance updates are applied in a timely manner after accounting for your required lead time. Lead time is discussed in the following steps.
- * 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 Start hour**, 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 **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.

7. Click Save Changes.

If you switch from rolling to non-rolling maintenance method, then Confirm Non-rolling Maintenance Method dialog is displayed.

- a. Enter the name of the infrastructure in the field provided to confirm the changes.
- b. Click Save Changes.

View or Edit a Scheduled Quarterly Maintenance for Exadata Cloud@Customer Infrastructure

Learn how to view and edit the time of the next scheduled maintenance.

- 1. Open the navigation menu. Under Oracle 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, under Maintenance, click the view link in the Next Quarterly Maintenance field.

The Exadata Infrastructure Maintenance page is displayed.

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 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.
- 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. You cannot reschedule the maintenance if the current time is within 2 hours of the scheduled maintenance start time.
- 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.
- 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.

Related Topics

https://support.oracle.com/epmos/faces/DocContentDisplay?id=888828.1



View the Maintenance History of Exadata Cloud@Customer Infrastructure

Learn how to view the maintenance history for an Exadata Cloud@Customer Infrastructure.

- 1. Open the navigation menu. Under Oracle 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, under Maintenance, click the view link in the Next Quarterly Maintenance field.

The Exadata Infrastructure Maintenance page is displayed.

6. On the **Exadata Infrastructure Maintenance** 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

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.
- 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 Exadata Cloud@Customer resources emit events, which are structured messages that indicate changes in resources.
- Overview of Events
- Notifications Overview
- Managing Infrastructure Maintenance Contacts
 Learn to manage your Exadata infrastructure maintenance contacts.

Creating First VM Cluster Network on Exadata Cloud@Customer

Follow the steps in the following sections to create your first VM Cluster Network.



 About Managing VM Cluster Networks on Exadata Database Service on Cloud@Customer
 The VM cluster provides a link between your 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 Exadata Database Service on Cloud@Customer

The VM cluster provides a link between your 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 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.
- 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 DB Server Selection. Edit DB Server Selection dialog is displayed.
- ii. Clear the checkboxes to omit one or more DB servers.
- iii. Click Save Changes. 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.
- **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.

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.
- 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 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.)

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.
- **Hostname:** Specify the hostname for each address in the client network.
- IP Address: Specify the IP address for each address in the client network.

b. 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.
- Hostname: Specify the hostname for each address in the backup network.
- **IP Address:** Specify the IP address for each address in the backup network.

c. 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.
- Hostname: Specify the hostname for each address in the backup network.
- **IP Address:** Specify the IP address for each address in the backup network.

d. 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.

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 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 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.

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 Exadata Cloud@Customer System

Learn how to create the first VM cluster on your Exadata Cloud@Customer system.

- About Managing VM Clusters on Exadata Database Service on Cloud@Customer The VM cluster provides a link between your Exadata Database Service on Cloud@Customer infrastructure and Oracle Databases you deploy.
- Prerequisites for VM Clusters on 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 a VM Cluster
 To create your VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

About Managing VM Clusters on Exadata Database Service on Cloud@Customer

The VM cluster provides a link between your 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.

Prerequisites for VM Clusters on 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/
Hc26biw3TXWGEakrK10Q== rsa-key-20160304
```



Related Topics

Managing Key Pairs on Linux Instances

Using the Console to Create a VM Cluster

To create your VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

To create a VM cluster, ensure that you have:

- Active Exadata infrastructure is available to host the VM cluster.
- A validated VM cluster network is available for the VM cluster to use.
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Choose the **Region** that contains your Exadata infrastructure.
- 3. Click VM Clusters.
- 4. Click Create VM Cluster.
- 5. Provide the requested information on the Create VM Cluster page:
 - a. **Choose 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 Database Service on Cloud@Customer 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. Choose the Oracle Grid Infrastructure version: From the list, choose the Oracle Grid Infrastructure release (19c and 23ai) that you want to install on the VM cluster.

The Oracle Grid Infrastructure release determines the Oracle Database releases that can be supported on the VM cluster. You cannot run an Oracle Database release that is later than the Oracle Grid Infrastructure software release.

- f. Choose an Exadata image 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.
 - 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.
- g. Configure VM Cluster:
 - Click Select DB Servers for VM placement to allocate VM resources.
 - On the **Select DB Servers** dialog, 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.

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 based on hardware generation 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 Changes.
- h. Specify the OCPU count per VM: Specify the OCPU count for each individual VM. The minimum value is 2 OCPUs per VM (for a live VM condition), unless you are specifying zero OCPUs (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 resources. See *Using the Console to Scale the Resources on a VM Cluster*.



The value for OCPU count for the whole VM Cluster will be calculated automatically based upon the per VM OCPU count you have specified and the number of physical Database Servers configured for the system. There is one VM created on each physical Database Server available.

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.

- i. Requested OCPU 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 count per VM field. This field is not editable.
- j. **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.
- k. 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.
- I. Specify the local file system size per VM (GB): Specify the local file system 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 X8-2 and X7-2 infrastructures.

Note that the minimum size of local system storage must be 60 GB. In addition to the 60 GB, each node of the VM must have at least 137 GB free for miscellaneous VM files. Each time when you create a new VM cluster, the space remaining out of the total available space is utilized for the new VM cluster.

For more information and instructions to specify the size for each individual VM, see *Introduction to Scale Up or Scale Down Operations*.

- m. Reserved local storage per VM (GB): Displays the local storage size reserved internally for root file systems, Oracle Grid Infrastructure Homes, and diagnostic logs. This field is not editable.
- n. Configure the Exadata Storage: The following settings define how the Exadata storage is configured for use with the VM cluster. These settings cannot be changed after creating the VM cluster.
 - **Specify Usable Exadata Storage:** Specify the size for each individual VM. The minimum recommended size is 2 TB.
 - 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.



Storage Allocation	DATA Disk Group	RECO Disk Group	SPARSE Disk Group
Exadata Snapshots: No	80%	20%	0% (The SPARSE disk group is not
Enable Backups on Local Exadata Storage: No			created.)
Exadata Snapshots: No	40%	60%	0% (The SPARSE disk group is not
Enable Backups on Local Exadata Storage: Yes			created.)
Allocate Storage for Exadata Snapshots: Yes	60%	20%	20%
Enable Backups on Local Exadata Storage: No			
Allocate Storage for Exadata Snapshots: Yes	35%	50%	15%
Enable Backups on Local Exadata Storage: Yes			

Table 4-12 Storage Allocation

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.

- p. 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.
 - License Included: Select this option to subscribe to Oracle Database software licenses as part of Exadata Database Service on Cloud@Customer.

q. 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).

r. 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**.

- **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**.

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 Exadata Database Service on Cloud@Customer Gen2, 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 to 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

Creating Database Backup Destinations for Exadata Cloud@Customer

Exadata 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 Exadata Database Service on Cloud@Customer

For backups, you can either use the 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 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 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 Exadata Database Service on Cloud@Customer

For backups, you can either use the Exadata Database Service on Cloud@Customer backup facility, or you can configure a backup location on a location you manage.



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 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 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 Exadata Database Service on Cloud@Customer Systems Learn to manage Oracle Database homes on 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

Prerequisites for Backup Destinations for Exadata Database Service on Cloud@Customer

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.

- 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 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:
 - 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.
 - Autonomous Database on Cloud@Customer:

Note:

Oracle Autonomous 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.

Using the Console for Backup Destinations for 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 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.



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



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 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

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 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 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.

Creating First Database Home on an 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 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 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 Exadata Cloud@Customer To create an Oracle Database home, review the list of API calls.

About Creating Oracle Database Homes on an 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 <code>dbaascli</code> utility.

Related Topics

• Using the dbaascli Utility with Exadata Database Service on Cloud@Customer Learn to use the dbaascli utility on Exadata Cloud@Customer.



Using the Console to Create Oracle Database Home on 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.

1. Open the navigation menu. Under Oracle 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.
 - 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 select an Image Type.

Oracle Provided Database Software Images: These images contain generally available versions of Oracle Database software.

Custom Database Software Images: These images are created by your organization and contain customized configurations of software updates and patches. Use the **Select a compartment** and **Select a Database version** selectors to limit the list of custom Database Software Images to a specific compartment or Oracle Database software major release version.

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 $\ensuremath{\texttt{Provisioning}}\xspace$ to $\ensuremath{\texttt{Available}}\xspace$.

Related Topics

Resource Tags

Using the API to Create Oracle Database Home on Exadata 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 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 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 Exadata 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.

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.



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 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

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.
 - 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.

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.

- Open the navigation menu Under Oracle Database, and 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. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
 VM Clusters is selected by default.
- b. Choose your **Compartment**. A list of VM Clusters is displayed for the chosen Compartment.
- c. 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.
- d. Click Database Homes.
- e. Click the name of the Database Home where you want to create the database.
- f. 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

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.
- **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 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.

- **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:
 - **None**: Select to not define a backup configuration for the database.
 - Local: Select to store backups locally in the Oracle Exadata Storage Servers on your Oracle Exadata Cloud at Customer system.
 This option is available only if you enabled backups on local Oracle Exadata storage in the VM cluster that you want to host the database.
 - 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. However, if you select **None** now, then you can select a backup destination after the database is created.

 Enable automatic backups: Select this option to enable daily backups using the policy for automatic backups.

This option is only enabled when you select a **Backup Destination Type** other than **None**. You can change this setting after database creation.

- (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 <code>instance_name</code> 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
- **Backup retention period**: From the list, you can choose the length of time that you want automatic backups to be 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 Oracle Zero Data Loss 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.

- 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 Database.

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.

Connecting to an Exadata Database Service on Cloud@Customer System

After deploying your Exadata 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 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 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 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 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 Exadata Database Service on Cloud@Customer System
 To access a virtual machine in an 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 Cloud@Customer system from a Unixstyle system using SSH, use this procedure.

Related Topics

• http://www.putty.org/



Prerequisites for Connecting to an Exadata Database Service on Cloud@Customer System

To access a virtual machine in an 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 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
- The SSH private key file that matches the public key associated with the deployment. This private key file must be in the PuTTY .ppk format. If the private key file was originally created on the Linux platform, you can use the PuTTYgen program to convert it to the .ppk format.

To connect to a virtual machine using the PuTTY program on Windows:

1. Download and install PuTTY.

To download PuTTY, go to http://www.putty.org/ and click the You can download PuTTY here link.

2. Run the PuTTY program (putty.exe).

The PuTTY Configuration window is displayed, showing the **Session** panel.

- 3. In the Host Name (or IP address) field, enter the host name or IP address of the virtual machine that you want to access.
- 4. Confirm that the **Connection type** option is set to **SSH**.



5. In the **Category** tree, expand **Connection** if necessary and then click **Data**.

The **Data** panel is displayed.

- 6. In the **Auto-login username** field, enter the operating system user you want to connect as:
 - Connect as the user opc to perform operations that require root or oracle access to the virtual machine, such as backing up or patching; this user can use the sudo command to gain root or oracle access to the VM.
- 7. Confirm that the When username is not specified option is set to Prompt.
- 8. In the Category tree, expand SSH and then click Auth.

The **Auth** panel is displayed.

- 9. Click the Browse button next to the Private key file for authentication field. Then, in the Select private key file window, navigate to and open the private key file that matches the public key that is associated with the deployment.
- 10. In the Category tree, click Session.

The **Session** panel is displayed.

- **11.** In the **Saved Sessions** field, enter a name for the connection configuration. Then, click **Save**.
- **12.** Click **Open** to open the connection.

The PuTTY Configuration window closes and the PuTTY terminal window displays.

If this is the first time you are connecting to the VM, the PuTTY Security Alert window is displayed, prompting you to confirm the public key. Click **Yes** to continue connecting.

Accessing a Database After You Connect to the Virtual Machine

After you connect to a virtual machine, you can use the following series of commands to identify a database and connect to it.

- 1. SSH in as the opc user.
- 2. sudo su oracle
- 3. Use the srvct1 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
```

 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
```

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 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 Exadata Database Service on Cloud@Customer system using Oracle Net Services.

- Using Oracle Net Services to Connect to a Database
 Oracle Database 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 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 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 Exadata Database Service on Cloud@Customer supports remote database access by using Oracle Net Services.

Because 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 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 Cloud@Customer using Oracle Net Services.

To connect to an Oracle Database on 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.

 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)))
```

Where:

alias-name is the name you use to identify the alias.

scan-name is the custom SCAN name.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- SID=*sid*-*name*. For example: SID=S12C1.
- SERVICE_NAME=service-name. For example: SERVICE NAME=PDB1.example.yourcloud.com.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

scan-name:1521/sid-or-service-entry

For example:

exalscan.example.com:1521/S12C1

Or

exalscan.example.com:1521/PDB1.example.yourcloud.com

Connecting to a Database Using a Node Listener

To connect to an Oracle Database instance on Exadata 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.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- SID=*sid*-name. For example, SID=S12C1.
- SERVICE_NAME=service-name. For example, SERVICE_NAME=PDB1.example.oraclecloudatcust.com.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

node:1521/sid-or-service-entry

For example:

exalnode01.example.com:1521/S12C1

Or

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



5 How-to Guides

A collection of tasks and procedures for managing Exadata Database Service on Cloud@Customer.

- Manage Database Security with Oracle Data Safe
- Connect to the Exadata Database Service on Cloud@Customer Service Learn how to connect to an Exadata Cloud@Customer system using SSH, and how to connect to an Exadata Cloud@Customer database using Oracle Net Services (SQL*Net).
- Manage 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
 Exadata 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 Exadata Database Service on
 Cloud@Customer.
- Manage Oracle Database Software Images Learn about Database Software Image resource type and how you can use it to create Oracle Databases and Oracle Database Homes and to patch databases.
- Create Oracle Database Homes on an Exadata Database Service on Cloud@Customer System

Learn to create Oracle Database Homes on Exadata Database Service on Cloud@Customer.

 Manage Oracle Database Homes on Exadata Database Service on Cloud@Customer Systems
 Learn to manage Oracle Database homes on Exadata Database Service on

Learn to manage Oracle Database homes on Exadata Database Service on Cloud@Customer.

- Manage Databases on 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 Exadata Database Service on Cloud@Customer System Learn to update and patch the 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 Exadata Database Service on Cloud@Customer Learn to configure and manage Data Guard associations in your VM cluster.



- Managing Encryption Keys on External Devices Learn how to store and manage database encryption keys.
- Migrate to Exadata Database Service on Cloud@Customer
 For general guidance on methods and tools to migrate databases to Oracle Cloud
 Infrastructure database services, including Exadata Cloud@Customer see
 "Migrating Databases to the Cloud".
- Overview of Exadata Cloud@Customer Gen1 to Out-of-Place Cloud Upgrade to 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. 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 Cloud@Customer database*.

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 at Cloud@Customer, on-premises, in the Oracle Cloud Infrastructure or in any other cloud.

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

Data Safe provides the following features:

- Security Assessment: Configuration errors and configuration drift are significant contributors to data breaches. Use security assessment to evaluate your database's configuration and compare it to Oracle and industry best practices. Security assessment reports on areas of risk and notifies you when configurations change.
- **User Assessment**: Many breaches start with a compromised user account. User Assessment helps you spot the riskiest database accounts those accounts which, if compromised, could cause the most damage and take proactive steps to secure them. User Assessment Baselines make it easy to know when new



accounts are added, or an account's privileges are modified. Use OCI events to receive proactive notifications when a database deviates from its baseline.

- Activity Auditing: Understanding and reporting on user activity, data access, and changes to database structures supports regulatory compliance requirements and can aid in post-incident investigations. Activity auditing collects audit records from databases and helps you manage audit policies. Audit insights make it easy to identify inefficient audit policies, while alerts based on audit data proactively notify you of risky activity.
- Sensitive Data Discovery: Knowing what sensitive data is managed in your applications is critical for security and privacy. Data discovery scans your database for over 150 different types of sensitive data, helping you understand what types and how much sensitive data you are storing. Use these reports to formulate audit policies, develop data masking templates, and create effective access control policies.
- Data Masking: Minimizing the amount of sensitive data your organization maintains helps you meet compliance requirements and satisfy data privacy regulations. Data masking helps you remove risk from your non-production databases by replacing sensitive information with masked data. With reusable masking templates, over 50 included masking formats, and the ability to easily create custom formats for your organization's unique requirements, data masking can streamline your application development and testing operations.

*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

Using Oracle Data Safe

Once your database is registered in Data Safe, you can leverage all features.



Security Assessment

Security Assessments are automatically scheduled once a week in Data Safe and provide an overall picture of your database security posture. It analyzes your database configurations, users and user entitlements, as well as security policies to uncover security risks and improve the security posture of Oracle Databases within your organization. A security assessment provides findings with recommendations for remediation activities that follow best practices to reduce or mitigate risk.

Start by reviewing the security assessment report for your database: View the latest assessment for a target database

You can find more details on Security Assessment under Security Assessment Overview.

User Assessment

User Assessments are automatically scheduled once a week in Data Safe and help you to identify highly privileged user accounts that could pose a threat if misused or compromised. User Assessment reviews information about your users in the data dictionaries on your target databases and then calculates a potential risk for each user, based on system privileges and role grants.

Start by reviewing the user assessment report for your database: View the latest user assessment for a target database

You can find more details on User Assessment under User Assessment Overview.

Data Discovery

Data Discovery searches for sensitive columns in your database. It comes with over 150 pre-defined sensitive types and you can also create your own sensitive types. You tell Data Discovery if you want to scan your entire database or just certain schemas and what type of sensitive information to look for, and it finds the sensitive columns that meet your criteria and stores them in a sensitive data model (SDM).

Start by discovering sensitive data in your database: Create Sensitive Data Models

You can find more details on Data Discovery under Data Discovery Overview.

Data Masking

Data masking, also known as static data masking helps you to replace sensitive or confidential information in your non-production databases with realistic and fully functional data with similar characteristics as the original data. Data Safe comes with pre-defined masking formats for each of the pre-defined sensitive types that can also be leveraged for your own sensitive types.

Once you know where sensitive data is stored in your database (for instance after running Data Discovery in Data Safe), you can start by creating a masking policy: Create Masking Policies

After you created a masking policy and copied your production database, you can mask your non-production copy: Mask Sensitive Data on a Target Database

You can find more details on Data Masking under Data Masking Overview.



Activity Auditing

Activity Auditing in Oracle Data Safe helps to ensure accountability and improve regulatory compliance. With Activity Auditing, you can collect and retain audit records per industry and regulatory compliance requirements and monitor user activities on Oracle databases with pre-defined reports and alerts. For example, you can audit access to sensitive data, security-relevant events, administrator and user activities, activities recommended by compliance regulations like the Center for Internet Security (CIS), and activities defined by your own organization.

If you are using the audit collection in Data Safe, up to 1 million audit records per target database per month are included for your Cloud@Customer database.

To use activity auditing, start the audit trail for your target database in Data Safe: Start an Audit Trail

Once the audit trail is started, you can monitor and analyze your audit data with pre-defined audit reports: View a Predefined or Custom Audit Report

You can find more details on Activity Auditing under Activity Auditing Overview.

Connect to the Exadata Database Service on Cloud@Customer Service

Learn how to connect to an Exadata 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 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 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 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 Exadata Database Service on Cloud@Customer System

To access a virtual machine in an 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 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 Exadata Database Service on Cloud@Customer System

To access a virtual machine in an 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 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
- 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.

Connecting from a Unix-Style System

To access a virtual machine on an Oracle 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 srvct1 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
```

 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

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 Exadata Database Service on Cloud@Customer system using Oracle Net Services.

- Using Oracle Net Services to Connect to a Database
 Oracle Database 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 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 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 Exadata Database Service on Cloud@Customer supports remote database access by using Oracle Net Services.

Because 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 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 Cloud@Customer using Oracle Net Services.

To connect to an Oracle Database on 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.

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)))
```

Where:

alias-name is the name you use to identify the alias.

scan-name is the custom SCAN name.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- SID=*sid*-name. For example: SID=S12C1.
- SERVICE_NAME=service-name. For example: SERVICE NAME=PDB1.example.yourcloud.com.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

scan-name:1521/sid-or-service-entry

For example:

exalscan.example.com:1521/S12C1

Or

exalscan.example.com:1521/PDB1.example.yourcloud.com

Connecting to a Database Using a Node Listener

To connect to an Oracle Database instance on Exadata 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 5-1 Defining a Net Service Alias That Directly References the Node

```
alias-name = (DESCRIPTION=
  (CONNECT_TIMEOUT=timeout)
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=node)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```



Where:

alias-name is the name you use to identify the alias.

timeout specifies a timeout period (in seconds), which enables you to terminate a connection attempt without having to wait for a TCP timeout. The (CONNECT_TIMEOUT=timeout) parameter is optional.

node is the hostname or IP address for the virtual machine that you want to use.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- SID=*sid-name*. For example, SID=S12C1.
- SERVICE_NAME=service-name. For example, SERVICE NAME=PDB1.example.oraclecloudatcust.com.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

node:1521/sid-or-service-entry

For example:

exalnode01.example.com:1521/S12C1

Or

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

Manage 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.
- About Changing the Home Region of ExaDB-C@C Infrastructure
- 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 Exadata Database Service on Cloud@Customer Learn how to provision an Exadata Database Service on Cloud@Customer system.
- Using the API to Manage Exadata 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.

About Changing the Home Region of ExaDB-C@C Infrastructure

ExaDB-C@C 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.
- 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
- 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.

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

Table 5-1 Key Additional Resources

Table 5-2 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 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



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 5-4 Key Additional Resources

Table 5-5 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

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: 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 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.




Figure 5-1 Overview of Order and Deployment Process

Multi-Rack Expansion

- Rack Minimum Configuration
 - 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.
- 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 Exadata 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 Exadata Database Service on Cloud@Customer

Learn how to provision an Exadata Database Service on Cloud@Customer system.

- Using the Console to Create Exadata Cloud@Customer Infrastructure To create your Exadata 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 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 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.
- 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 Exadata 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 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 Delete Exadata Database Service on Cloud@Customer
 Infrastructure

To delete Oracle Exadata Database Service on Cloud@Customer infrastructure, complete the prerequisites, and then complete this procedure.

- Using the Console to Manage Tags for Your Exadata Cloud@Customer Resources
- Managing Infrastructure Maintenance Contacts Learn to manage your Exadata infrastructure maintenance contacts.

Using the Console to Create Exadata Cloud@Customer Infrastructure

To create your Exadata Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.

- 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.

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:



• **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.

• **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.

- **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.
- 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 servers:
 - * **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:
 - * **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 checkbox is automatically selected. You can also check 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 select this check box. 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.

• Configure the cloud 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.

- 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:
 - 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 /22. The maximum CIDR block /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.

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 5-6X10M (Quarter Rack X10M, Quarter Rack X10M-L, and QuarterRack X10M-XL) CIDR Requirements

Table 5-7 X9M CIDR Requirements

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
6	64	/23	/22

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
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 5-7 (Cont.) X9M CIDR Requirements

Table 5-8 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	—	



Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
32	64	/21	/20

Table 5-8 (Cont.) X8M CIDR Requirements

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.
- 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 browserdetected 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**.

Network bonding mode



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
 - Configure automatic maintenance Click Edit Maintenance Preferences.

Edit Maintenance Preferences dialog is displayed.

In the Edit Maintenance Preferences dialog, configure the following:

- 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.

- * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.
 - Default: 30 minutes

Maximum: 120 minutes

- * Maintenance schedule:
 - * **No preference:** The system assigns a date and start time for infrastructure maintenance.
 - * **Specify a schedule:** Choose your preferred month, week, weekday, start time, and lead time for infrastructure maintenance.
 - * **Lead Time:** Specify the minimum number of weeks ahead of the maintenance event you would like to receive a notification message.

Click Save Changes.

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.

Provide maintenance contacts

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

You can skip adding maintenance contacts while creating your infrastructure. However, you must add a primary contact prior to activating your infrastructure. Ensure that you provide the details of the contact that you used while registering the Customer Support Identifier (CSI) associated with this infrastructure, as a primary contact.

Optionally, you can 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. Note that 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.

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 Exadata Infrastructure.

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 Cloud@Customer Exadata Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X10M.
- ExaCC gen2: Troubleshooting VPN/WSS connection from Customer Side
- Resource Tags

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 Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Exadata Infrastructure.
- 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 Actions and then select 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.

 Click Apply 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**.

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 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 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. 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.

b. 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.

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.

c. 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.
- **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**.

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 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.

Using the Console to Activate 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 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 **Activate** on the banner displayed.

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 dialog to upload the activation file, and then click **Activate Now**.

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 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 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.

- 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.

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 5-9 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	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 5-10 X9M CIDR Requirements

Table 5-11 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	_	

Maximum DB Servers	Maximum Storage Servers	Required Administration Network CIDR	Required Secure Fabric Network CIDR
28	35		
29	18		
32	64	/21	/20

Table 5-11	(Cont.)	X8M CIDR	Requirements
------------	---------	----------	--------------

- 1. Open the navigation menu. Under Oracle 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.
- 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 Infrastructure.

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 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.
- 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 New 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 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. 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 Databases, 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.
- 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 Add Storage Capacity.
- 7. Review the advisory on the Add Storage Capacity dialog, and then click Add Storage Capacity.

Note:

While Add Storage Capacity 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 Add Storage Capacity operation is in progress.

Using the Console to View Details of Exadata 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 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 a VM Cluster To create your VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Using the Console to Move 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 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.

The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.

- 5. Click Move Resource.
- 6. In the resulting dialog, choose the new compartment for the Oracle Exadata infrastructure, and click **Move Resource**.

Using the Console to Delete Exadata Database Service on Cloud@Customer Infrastructure

To delete Oracle Exadata Database Service on Cloud@Customer infrastructure, complete the prerequisites, and then complete this procedure.

Deleting Exadata Database Service on Cloud@Customer infrastructure removes it from the Cloud Control Plane.

If you are deleting Oracle Exadata infrastructure before activation, then if required, you can create replacement Oracle Exadata infrastructure without any input from Oracle.

If you are deleting 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 delete 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 delete the Oracle Exadata infrastructure. In response to the SR, Oracle flags the Oracle Exadata infrastructure as ready for deletion.

After Oracle has flagged the Oracle Exadata infrastructure, delete the Exadata infrastructure by using the following process:

- 1. Open the navigation menu. Under Oracle 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 delete.
- 3. Click Exadata Infrastructure.
- 4. Click the name of the Oracle Exadata infrastructure that you want to delete.

The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.

- 5. Click Delete.
- 6. In the resulting dialog, enter the Oracle Exadata infrastructure name and click **Delete Exadata Infrastructure** to confirm the action.

Using the Console to Manage Tags for Your Exadata Cloud@Customer Resources

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Choose your Compartment.
- 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

Managing Infrastructure Maintenance Contacts

Learn to manage your Exadata infrastructure maintenance contacts.

- View Primary Maintenance Contact You must associate primary contact with the Customer Support Identifier (CSI).
- Add Secondary Contacts You can add up to nine secondary contacts.
- Edit Maintenance Contacts Edit maintenance contacts to update details.
- Promote a Secondary Contact to Primary You can promote a secondary contact to primary. The current primary is automatically demoted to secondary.



Remove a Secondary Contact

You can remove a secondary contact anytime you want to.

View Primary Maintenance Contact

You must associate primary contact with the Customer Support Identifier (CSI).

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 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.
- 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.
- Find the CSI and the primary contact under Maintenance. The operations team sets the infrastructure maintenance Service Level Objective (SLO) to Degraded:
 - If the primary contact CSI verification has failed.
 - If the primary contact is missing.
 - If the primary contact is verified and unresponsive.

Also, a warning message is displayed on the Console as follows:

```
"Ensure that the primary contact associated with your Customer Support
Identifier (CSI) is available for Oracle support to coordinate maintenance-
related activities.
The infrastructure maintenance Service Level Objective (SLO) is set to
degraded status without proper primary contact."
```

Add a primary contact before you activate the infrastructure.

Add Secondary Contacts

You can add up to nine secondary contacts.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 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.
- 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 Add Contacts window, add contact details.



8. Click Add Contacts.

Edit Maintenance Contacts

Edit maintenance contacts to update details.

- 1. Open the navigation menu. Under Oracle 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.
- 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. In the Manage Exadata Infrastructure Contacts window, click the actions button, and then select **Edit Contact**.
- 7. In the Edit Contacts window, edit the details.
- 8. Click Save.

Promote a Secondary Contact to Primary

You can promote a secondary contact to primary. The current primary is automatically demoted to secondary.

- 1. Open the navigation menu. Under Oracle 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 promote a secondary contact to primary.
- 3. Click Exadata Infrastructure.
- Click the name of the Oracle Exadata infrastructure for which you want to promote a secondary contact to primary. 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 the actions button, and then select **Make Primary**.
- 7. In the Promote to Primary Contact dialog, click **Promote**.

Remove a Secondary Contact

You can remove a secondary contact anytime you want to.

1. Open the navigation menu. Under Oracle 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.
- 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. In the Manage Exadata Infrastructure Contacts window, click the actions button, and then select **Remove**.
- 7. In the Remove Infrastructure Contact dialog, click **Remove**.

Using the API to Manage Exadata 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

Configure Oracle-Managed Infrastructure Maintenance

Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud@Customer.

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 the infrastructure maintenance process and configuring the maintenance controls refer to the following:

- About Oracle Managed Exadata Cloud@Customer Infrastructure Maintenance Updates Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud@Customer.
- Infrastructure Maintenance Contacts Maintenance contacts are required for service request based communications for hardware replacement and other maintenance events.
- 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.
- 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 Exadata Cloud@Customer Infrastructure Maintenance Controls Oracle Exadata Cloud@Customer uses the same API as Oracle Cloud Infrastructure to manage infrastructure maintenance controls.

About Oracle Managed Exadata Cloud@Customer Infrastructure Maintenance Updates

Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud@Customer.

Oracle patches and updates include the physical database server hosts, Exadata Storage Servers, Network Fabric Switches, management switch, power distribution units (PDUs), integrated lights-out management (ILOM) interfaces, and Control Plane Servers. This is referred to as Exadata Cloud@Customer infrastructure maintenance.



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.

- 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 with CVSS scores greater than or equal to 7.
- 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.

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.

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 impact older applications not written to handle a rolling instance outage. 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 do not impact VM cluster VM's availability. However, the rolling storage server maintenance can result in reduced IO performance as storage servers are taken offline (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 not under maintenance will minimize (or eliminate) any performance impact.

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 instances. This allows you to avoid any impact on your primary database during infrastructure maintenance.

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 10 days prior to the maintenance start and another precheck is run approximately 24 hours prior to maintenance start. If the prechecks identify an issue that requires rescheduling the maintenance contacts.

The time taken to update infrastructure components varies depending on the number of database servers and storage servers in the Exadata infrastructure, the maintenance method, and whether custom action has been enabled. The approximate times provided are estimates. Time for custom action, if configured, is not included in the estimates below. Database server maintenance time may vary depending on the time required to shutdown each VM before the update and then start each VM and associated resources after the update of each node before proceeding to the next node. The storage server maintenance time will vary depending on the time required for the ASM rebalance, which is not included in the estimates below. If issues are encountered during maintenance this may also delay completion beyond the approximate time listed. In such a situation, if Oracle cloud operations determine resolution would extend beyond the expected window, they will send a notification and may reschedule the maintenance.

Note:

The timeframes mentioned below can change if Oracle cloud operations determine that additional maintenance work is needed. If additional time is necessary, Oracle will send a customer notification in advance to inform customers that additional time will be required for the next quarterly maintenance window.

- Rolling:
 - Each database server takes 90 minutes on average.
 - Each storage server takes 60 minutes on average.
 - Each InfiniBand or RoCE fabric switch takes 30 minutes on average.
 - The approximate total time for infrastructure maintenance is as follows:



- * Base and Quarter Rack (2 Database Servers/3 Storage Servers): Approximately 7 hours
 - 2 Database Servers X 90 = 180 minutes
 - 3 Storage Servers X 60 = 180 minutes
 - 2 InfiniBand or RoCE Fabric Switch X 30 = 60 minutes
- * Half Rack (4 Database Servers/6 Storage Servers): Approximately 13 hours
 - 4 Database Servers X 90 = 360 minutes
 - 6 Storage Servers X 60 = 360 minutes
 - 2 InfiniBand or RoCE Fabric Switch X 30 = 60 minutes
- * Full Rack (8 Database Servers/12 Storage Servers): Approximately 26 hours
 - 8 Database Servers X 90 = 720 minutes
 - 12 Storage Servers X 60 = 720 minutes
 - 2 InfiniBand or RoCE Fabric Switch X 30 = 60 minutes
- Non-Rolling:
 - All database servers take 180 minutes on average.
 - All storage servers take 60 minutes on average.
 - Storage Servers and Database servers are brought back online prior to starting fabric switch maintenance.
 - Network fabric switches are still updated in a rolling method and take 30 minutes each on average.
 - The approximate total time for infrastructure maintenance is 5 hours regardless of shape:
 - * All Database Servers = 180 minutes
 - * All Storage Servers = 60 minutes
 - * 2 InfiniBand or RoCE Fabric Switch = 60 minutes

Related Topics

- Achieving Continuous Availability For Your Applications
- Maximum Availability Architecture (MAA) Best Practices

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 with CVSS scores greater than or equal to 7.



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 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.

Updates to database servers are applied online via Ksplice technology, and have no impact to workloads running on the compute (database) servers, as database server security updates are applied online to the host server while your VM and all processes within the VM, including databases, remain up and running. Servers and VMs are not restarted. Updates to storage servers are applied in a rolling fashion. As with quarterly maintenance, the impact of rebooting storage servers should be minimal to applications.

While updating your services infrastructure, some operations including memory, and storage scaling, operating system and Grid Infrastructure patching (including prechecks), and elastic expansion of compute and storage servers may be blocked.

Note:

Only VM startup and shutdown operations are supported during monthly infrastructure maintenance.

Please plan to defer these operations until after the updates are complete. Application of security updates takes about 15 minutes per DB server host, plus 60 minutes per storage server depending on the I/O activity. 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 for Exadata Cloud@Customer Infrastructure

Learn how to view and edit the next scheduled security maintenance.

 Exadata Cloud@Customer Infrastructure Maintenance Event Types Review the list of event types that Exadata Cloud@Customer Infrastructure Maintenance emits.

Understanding Monthly and Quarterly Maintenance in the Same Month



Special considerations are made when both quarterly and monthly security maintenance are scheduled to run in the same month. Quarterly maintenance will reapply any security fixes already applied by security maintenance, and neither quarterly nor monthly maintenance will apply a storage server update if the existing storage server version is the same or newer than the version contained in the update.

- The contents of the updates applied during quarterly maintenance are determined at the start of the maintenance quarter and use the latest Exadata release from the month prior to the start of the maintenance quarter. If any additional security fixes are available at that time, those updates are included in the quarterly maintenance. That image is then used throughout the quarter. For example, the January release is used for quarterly maintenance in Feb, March, and April.
- When quarterly maintenance is applied it is possible there are security updates
 previously installed on the database servers are not included in the quarterly
 maintenance release to be applied. In that case, the automation will apply the
 same security fixes to new release installed by the quarterly maintenance so there
 will not be any regression in security fixes. If the current image on the storage
 server is the same or newer than that to be applied by the quarterly or monthly
 security maintenance, that maintenance will be skipped for the storage servers.

If quarterly maintenance is scheduled within 24 hours of the time the monthly is scheduled, the scheduled monthly maintenance will be skipped, and the monthly update will instead be applied immediately following the quarterly maintenance.

- When scheduled at the same time, the monthly update is executed immediately following the completion of the quarterly maintenance.
- If monthly maintenance is scheduled to begin 0-24 hours ahead of the quarterly maintenance, then the monthly maintenance will not execute as scheduled, but instead, wait and be executed immediately following the quarterly maintenance. If the quarterly maintenance is subsequently rescheduled, then the monthly security maintenance will begin immediately. Oracle, therefore, recommends scheduling quarterly and monthly maintenance at the same time. As a result, if you reschedule the quarterly at the last moment, the monthly maintenance will run at the scheduled time instead of immediately upon editing the schedule. You can also reschedule the monthly security maintenance when rescheduling the quarterly maintenance 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 guarterly maintenance is completed. The guarterly maintenance will perform rolling maintenance on the database servers, which will impact applications not written to handle a rolling reboot. The guarterly 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 guarterly maintenance. Exceptions to this rule may occur if it is the first month of a maintenance guarter, 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 guarterly maintenance may contain the same storage server release as the security maintenance and the security maintenance storage server updates will be skipped.
- The impact to end users of running the quarterly maintenance before the security maintenance should be roughly the same as running the security maintenance first. The quarterly maintenance will be a disruptive event, but the security maintenance rebooting the storage servers should cause minimal disruption, and the security maintenance is applied to the database servers online. However, if you must coordinate a maintenance window with your end users for the security maintenance, this will require two maintenance windows. You can schedule those two maintenance windows to run back-to-back, to appear as single maintenance window to end users. To do this, reschedule the security maintenance. 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.

Related Topics

- Using the Console to Create Exadata Cloud@Customer Infrastructure To create your Exadata Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Managing Infrastructure Maintenance Contacts
 Learn to manage your Exadata infrastructure maintenance contacts.

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.

For more information, see the following:

- View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure
- View or Edit a Scheduled Quarterly Maintenance for Exadata Cloud@Customer Infrastructure
- View or Edit a Scheduled Security Maintenance for Exadata Cloud@Customer Infrastructure
- View the Maintenance History of Exadata Cloud@Customer Infrastructure
- View and Edit Quarterly Maintenance While Maintenance is In Progress or Waiting for Custom Action
- View or Edit a Scheduled Security Maintenance for Exadata Cloud@Customer Infrastructure
- View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure To edit your Oracle Exadata Database Service on Cloud@Customer infrastructure quarterly maintenance preferences, he prepared to provide values for the 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 a Scheduled Quarterly Maintenance for Exadata Cloud@Customer Infrastructure

Learn how to view and edit the time of the next scheduled maintenance.

 View or Edit a Scheduled Security Maintenance for Exadata Cloud@Customer Infrastructure

Learn how to view and edit the next scheduled security maintenance.

- View the Maintenance History of Exadata Cloud@Customer Infrastructure Learn how to view the maintenance history for an Exadata Cloud@Customer Infrastructure.
- View and Edit Quarterly Maintenance While Maintenance is In Progress or 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.

View or Edit Quarterly Infrastructure Maintenance Preferences for Exadata Cloud@Customer Infrastructure

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 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.
- 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 Maintenance Preferences. Edit Maintenance Preferences page is displayed.

Note:

Changes made to maintenance preferences apply only to future maintenance, not the maintenance that has already been scheduled. To modify scheduled maintenance, see *View or Edit a Scheduled Maintenance for Exadata Cloud@Customer Infrastructure.*

- 6. On the Edit Maintenance Preferences page, configure the following:
 - 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.

Default: 30 minutes

Maximum: 120 minutes

- Maintenance schedule:
 - No preference: The system assigns a date and start time for infrastructure maintenance.
 - **Specify a schedule:** Choose your preferred month, week, weekday, start time, and lead time for infrastructure maintenance.
 - * Under **Maintenance months**, specify at least one month for each quarter during which Exadata infrastructure maintenance will take place. You can select more than one month per quarter. If you specify a long lead time for advanced notification (for example, 4 weeks), you may wish to specify 2 or 3 months per quarter during which

maintenance runs can occur. This will ensure that your maintenance updates are applied in a timely manner after accounting for your required lead time. Lead time is discussed in the following steps.

- * 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 Start hour**, 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 **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.

7. Click Save Changes.

If you switch from rolling to non-rolling maintenance method, then Confirm Non-rolling Maintenance Method dialog is displayed.

- a. Enter the name of the infrastructure in the field provided to confirm the changes.
- b. Click Save Changes.

View or Edit a Scheduled Quarterly Maintenance for Exadata Cloud@Customer Infrastructure

Learn how to view and edit the time of the next scheduled maintenance.

- 1. Open the navigation menu. Under Oracle 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, under **Maintenance**, click the view link in the **Next Quarterly Maintenance** field.

The Exadata Infrastructure Maintenance page is displayed.

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.



- 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. You cannot reschedule the maintenance if the current time is within 2 hours of the scheduled maintenance start time.
- 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.

Related Topics

https://support.oracle.com/epmos/faces/DocContentDisplay?id=888828.1

View or Edit a Scheduled Security Maintenance for Exadata Cloud@Customer Infrastructure

Learn how to view and edit the next scheduled security maintenance.

1. Open the navigation menu. Under Oracle 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. 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, under Maintenance, click the view link in the Next Security Maintenance field.

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 the **Edit** link in the **Scheduled Start Time** field 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.

Edit Maintenance Start Time dialog is displayed.

- 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 of Exadata Cloud@Customer Infrastructure

Learn how to view the maintenance history for an Exadata Cloud@Customer Infrastructure.

- 1. Open the navigation menu. Under Oracle 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, under **Maintenance**, click the **view** link in the **Next Quarterly Maintenance** field.



The Exadata Infrastructure Maintenance page is displayed.

6. On the **Exadata Infrastructure Maintenance** 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

View and Edit Quarterly Maintenance While Maintenance is In Progress or 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 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.

View and Edit Maintenance While Maintenance is In Progress

1. Click the View link in the Next Quarterly Maintenance field.

The Exadata Infrastructure Maintenance page is displayed.

2. 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.



- 3. 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

4. Click Save Changes.

If you have configured the rolling maintenance method, then the **View** link is displayed in the **Total Estimated Maintenance Time** field.

a. Click View.

Estimated Maintenance Time Details page is displayed with details that include:

- Total Estimated Maintenance Time
- Database Servers Estimated Maintenance Time
- Storage Servers Estimated Maintenance Time
- Network Switches Estimated Maintenance Time
- Order in which components are updated. In rolling maintenance, components are updated in the sequence displayed.
- b. Click Close.

View and Edit Maintenance While Maintenance is Waiting for Custom Action

1. Click the View link in the Next Quarterly Maintenance field.

Exadata Infrastructure Maintenance page is displayed.

Note:

- Editing a maintenance run is not available while waiting for custom action.
- While maintenance is waiting for custom action, an information block is displayed. The information block is removed after the maintenance resumes.
- 2. 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.

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.
- 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 Exadata Cloud@Customer resources emit events, which are structured messages that indicate changes in resources.
- Overview of Events
- Notifications Overview
- Managing Infrastructure Maintenance Contacts Learn to manage your Exadata infrastructure maintenance contacts.

Using the API to Manage Exadata Cloud@Customer Infrastructure Maintenance Controls

Oracle Exadata 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
- ListExadataInfrastructures
- UpdateExadataInfrastructure
- UpdateMaintenanceRun
- GetMaintenanceRun
- ListMaintenanceRuns

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 Exadata Database Service on Cloud@Customer The VM cluster provides a link between your 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.

About Managing VM Cluster Networks on Exadata Database Service on Cloud@Customer

The VM cluster provides a link between your 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 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.
- 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 DB Server Selection. Edit DB Server Selection dialog is displayed.
- ii. Clear the checkboxes to omit one or more DB servers.
- iii. Click Save Changes.
 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.
- **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.

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.
- 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 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.

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 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.
- 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 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.)

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.
- Hostname: Specify the hostname for each address in the client network.
- IP Address: Specify the IP address for each address in the client network.

b. 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.
- Hostname: Specify the hostname for each address in the backup network.
- **IP Address:** Specify the IP address for each address in the backup network.

c. 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.
- **Hostname:** Specify the hostname for each address in the backup network.
- **IP Address:** Specify the IP address for each address in the backup network.

d. 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.

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 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.
- 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 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.
- 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.

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 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 Report.

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
<vlanId>. Please ensure vlan Id, IP address and Gateway are correct"
    ],
    "Gateway":[
        "Gateway":[
        "Gateway <G1> is not accessible from source <ipaddr > with vlan id
<vlanId>. Please ensure vlan Id, IP address and Gateway are correct"
    ],
    "Gateway <G1> is not accessible from source <ipaddr > with vlan id
<vlanId>. Please ensure vlan Id, IP address and Gateway are correct",
    ],
    "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"
   1,
   "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."
   1
}
```

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 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.
- 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.

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.
- 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**.

 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 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**.
- 7. Click **Remove** on the confirmation dialog displayed.

Manage VM Clusters

Learn how to manage your VM clusters on Exadata Database Service on Cloud@Customer.

- About Managing VM Clusters on Exadata Database Service on Cloud@Customer The VM cluster provides a link between your 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 Exadata Cloud@Customer Learn how to use the console to create, edit, and manage your VM Clusters on Oracle Exadata Cloud@Customer.
- Using the API to Manage Exadata Cloud@Customer VM Clusters Review the list of API calls to manage your 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 Exadata Database Service on Cloud@Customer

The VM cluster provides a link between your 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.

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.

Related Topics

• Using the Console to Create a VM Cluster

To create your 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).

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 a VM Cluster To create your 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.



Oracle Trace File Analyze (TFA) Component Driven Logs Collections

The directories are generally assigned to a component and that component can then be used to guide TFA to the files it needs to collect, for example, requesting the CRS component would tell TFA to look at directories mapped to the CRS component and find files that match the required collection time frame.

Note:

If have previously opted in for incident log and trace file collection and decide to opt out when Oracle Cloud operations run a log collection job, then the job will run its course and will not cancel. Future log collections won't happen until you opt-in again to the incident logs and trace file collection option.

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



Component	Script Files/Directories	
OS: Operating system logs	oscollect.pl	 /var/log/messages OSWatcher archive Exadata Only: ExaWatcher archive /opt/ oracle.ExaWatcher /archive/

Table 5-12 Oracle Trace File Analyze (TFA) Component Driven Logs Collections

crscollect.pl	
	 /etc/oracle GIHOME/crf/db/ HOSTNAME1 GIHOME/crs/log GIHOME/cvs/log GIHOME/cv/log GIHOME/evm/ admin/log GIHOME/evm/admin/ logger GIHOME/log/-/ client GIHOME/log/ HOSTNAME1 GIHOME/log/ HOSTNAME1/client GIHOME/log/ HOSTNAME1/client GIHOME/log/ HOSTNAME1/crflogd GIHOME/log/ HOSTNAME1/crflogd GIHOME/log/ HOSTNAME1/crfmond GIHOME/log/ HOSTNAME1/crsd GIHOME/log/ HOSTNAME1/crsd GIHOME/log/ HOSTNAME1/crsd GIHOME/log/ HOSTNAME1/ctssd GIHOME/log/ HOSTNAME1/ctssd GIHOME/log/ HOSTNAME1/diskmon GIHOME/log/ HOSTNAME1/diskmon GIHOME/log/ HOSTNAME1/gipcd GIHOME/log/ HOSTNAME1/gipcd GIHOME/log/ HOSTNAME1/gipcd GIHOME/log/ HOSTNAME1/gipcd GIHOME/log/ HOSTNAME1/gipcd GIHOME/log/ HOSTNAME1/gipcd GIHOME/log/ HOSTNAME1/gipcd GIHOME/log/ HOSTNAME1/gipcd GIHOME/log/ HOSTNAME1/gipcd
	HOSTNAME1/gipcd GIHOME/log/ HOSTNAME1/gnsd GIHOME/log/
	 HOSTNAME1/gpnpd GIHOME/log/ HOSTNAME1/mdnsd GIHOME/log/ HOSTNAME1/ohasd GIHOME/log/ HOSTNAME1/racg GIHOME/log/

Table 5-12(Cont.) Oracle Trace File Analyze (TFA) Component Driven LogsCollections



Component	Script	Files/Directories
		• GIHOME/log/
		HOSTNAME1/xag
		 GIHOME/log/diag/ asmtool
		 GIHOME/log/diag/
		clients
		• GIHOME/log/
		procwatcher/ PRW SYS HOSTNAME
		• GIHOME/
		network/log
		 GIHOME/opmn/logs
		 GIHOME/racg/log
		• GIHOME/ Tacg/ Tog • GIHOME/
		scheduler/log
		• GIHOME/srvm/log
		• GRIDBASE/crsdata
		@global/cvu
		• GRIDBASE/crsdata HOSTNAME1/core
		 GRIDBASE/crsdata HOSTNAME1/
		crsconfig
		• GRIDBASE/crsdata HOSTNAME1/crsdia
		• GRIDBASE/crsdata
		HOSTNAME1/cvu
		 GRIDBASE/crsdata HOSTNAME1/evm
		 GRIDBASE/crsdata HOSTNAME1/output
		 GRIDBASE/crsdata
		HOSTNAME1/
		ovmmwallets
		 GRIDBASE/crsdata HOSTNAME1/script
		 GRIDBASE/crsdata
		HOSTNAME1/trace
		• GRIDBASE/
		diag/crs/-/crs/ cdump
		• GRIDBASE/
		diag/crs/
		HOSTNAME1/crs/
		cdump
		• GRIDBASE/
		diag/crs/

Table 5-12(Cont.) Oracle Trace File Analyze (TFA) Component Driven LogsCollections



Component	Script	Files/Directories
		HOSTNAME1/crs/ incident GRIDBASE/ diag/crs/ HOSTNAME1/crs/ trace
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</instance_name></dbname> ORACLE_BASE/diag/ rdbms/<dbname>/ <instance_name>/ trace</instance_name></dbname> ORACLE_BASE/diag/ rdbms/<dbname>/ <instance_name>/ instance_name>/ incident</instance_name></dbname>

Table 5-12(Cont.) Oracle Trace File Analyze (TFA) Component Driven LogsCollections

Cloud Tool Logs

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

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

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

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

DCS Agent Logs

/opt/oracle/dcs/log/



Tooling-Related Grid Infrastructure/Database Logs

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

Related Topics

- Overview of Automatic Diagnostic Collection
 By enabling diagnostics collection and notifications, Oracle Cloud Operations and
 you will be able to identify, investigate, track, and resolve guest VM issues quickly
 and effectively. Subscribe to Events to get notified about resource state changes.
- Health Metrics
 Review the list of database and non-database health metrics collected by Oracle
 Trace File Analyzer.
- Using the Console to Create a VM Cluster
 To create your 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.



Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
CpuUtiliza tion	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.
StorageUti lization	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.
BlockChang es	DB Block Changes	Changes per second	Mean	One minute	Five minutes	The Average number of blocks changed per second.

Table 5-13 Guest VM Health Metrics List - Database Metrics



Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
ExecuteCou nt	Execute Count	Count	Sum	One minute	Five minutes	The number of user and recursive calls that executed SQL statements during the selected interval.
CurrentLog ons	Current Logons	Count	Sum	One minute	Five minutes	The number of successfu logons during the selected interval.
Transactio nCount	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.
StorageUse d	Storage Space Used	GB	Max	One hour	One hour	Total amount of storage space used by the database at the collection time.

Table 5-13 (Cont.) Guest VM Health Metrics List - Database Metrics



Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
StorageAll ocated	Storage Space Allocated	GB	Max	One hour	One hour	Total amount of storage space allocated to the database at the collection time.
StorageUse dByTablesp ace	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.
StorageAll ocatedByTa blespace	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.

Table 5-13 (Cont.) Guest VM Health Metrics List - Database Metrics



Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
StorageUti lizationBy Tablespace	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.

Table 5-13 (Cont.) Guest VM Health Metrics List - Database Metrics

Guest VM Health Metrics List - Non-Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
ASMDiskgrou pUtilizatio n	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 contain database files for recovery such as archives and flashback logs.
FilesystemU tilization	Filesystem Utilization	Percentage	Max	One minute	Percent utilization of provisioned filesystem.



Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
CpuUtilizat ion	CPU Utilization	Percentage	Mean	One minute	Percent CPU utilization.
MemoryUtili zation	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.
SwapUtiliza tion	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.
OcpusAlloca ted	OCPU Allocated	Integer	Max	One minute	The number of OCPUs allocated.

Table 5-14	(Cont.) Guest VM Health Metrics List - Non-Database Metrics
------------	---

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 a VM Cluster To create your 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.
- Resizing Memory and Large Pages
- Calculating the ASM Storage
- Estimating How Much Local Storage You Can Provision to 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.

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.

Scaling up or down of these resources requires thorough auditing of existing usage and capacity management by the customer DB administrator. Review the existing usage to avoid failures during or after a scale down operation. While scaling up, consider how much of these resources are left for the next VM cluster you are planning to create. Exadata Cloud@Customer Cloud tooling calculates the current usage of memory, local disk, and ASM storage in the VM cluster, adds headroom to it, and arrives at a "minimum" value below which you cannot scale down, and expects that you specify the value below this minimum value.



Note:

- When creating or scaling a VM Cluster, setting the number of OCPUs 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 for each VM. These reserved OCPUs 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 when showing maximum available OCPU, so you should account for these reserved OCPU when performing any subsequent scaling operations to ensure the operation can acquire enough OCPUs 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. Productions 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.

Resizing Memory and Large Pages

You can scale the database server memory up and down in a VM Cluster. Scaling memory requires a rolling restart of the database servers to take effect.

Changing the memory in a VM Cluster will affect the large pages (HugePages) settings for the VMs in that cluster. When a VM is initially created, each VM's operating system is configured with 50% of the memory allocated to the VM for large pages, and databases are configured to use that memory for their SGA. Oracle recommends 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 not recommended, you are allowed to modify the large pages configuration. Any changes you make may be modified by automation should you subsequently resize the memory available to the VM. Upon a memory resize operation, the cloud automation will attempt to maintain the same amount of large pages memory as a percentage of the total memory, with a cap of 60%. If you configure large pages to be greater than 60% of total memory, the cloud automation will resize it to 60% of total memory. This ensures sufficient conventional memory for the virtual machine to boot. The automation will perform a precheck to determine the actual large pages memory in use by the running database instances, and ensure after the resize there is enough large pages memory available to support those same databases. If there will not be sufficient memory available after the resize, the precheck will fail and the resize will not continue.



Calculating the ASM Storage

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/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
                   Y
0
                         DATAC5/
                        512
MOUNTED HIGH N
                512
                                   4096 4194304 3135456
3036336 348384
                        895984
                  N RECOC5/
0
ASMCMD>
```

• With SPARSE:

```
/u01/app/19.0.0.0/grid/bin/asmcmd lsdg
ASMCMD>
State Type Rebal Sector Logical Sector Block AU
Total MB Free MB Req mir free MB Usable file MB
Offline disks Voting files Name
               512 512
MOUNTED HIGH N
                                 4096 4194304
12591936 10426224 1399104
                                3009040
                 Y DATAC5/
512 512
0
MOUNTED HIGH N
                                 4096 4194304
3135456 3036336 348384
                                895984
0 N RECOC5/
MOUNTED HIGH N 512 512
                                 4096 4194304
31354560 31354500 3483840
                                8959840
```



Ν SPRC5/ ASMCMD>

Note:

0

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.
- New requested size should pass the following conditions: 4.
 - 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

For RECO: (32.36 * 1.15) <= 1638 GB

For SPR: (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.

Estimating How Much Local Storage You Can Provision to Your VMs

X8-2 and X7-2 Systems

You specify how much space is provisioned from local storage to each VM. This space is mounted at location /u02, and is used primarily for Oracle Database homes. The amount of



local storage available will vary with the number of virtual machines running on each physical node, as each VM requires a fixed amount of storage for the root file systems, GI homes, and diagnostic log space. Refer to the tables below to see the maximum amount of space available to provision to local storage (/u02) across all VMs.

- Total space available for VM images (X7 All Systems): 1237 GB
- Total space available for VM images (X8 All Systems): 1037 GB
- Fixed storage per VM: 137 GB

#VMs	Fixed Storage All VMs (GB)	X8-2 Space for ALL /u02 (GB)	X7-2 Space for ALL /u02 (GB)
1	137	900	1100
2	274	763	963
3	411	626	826
4	548	489	689
5	685	352	552
6	822	N/A	415

Table 5-15Space allocated to VMs

For an X8-2, to get the max space available for the nth VM, take the number in the table above and subtract anything previously allocated for /u02 to the other VMs. So if you allocated 60 GB to VM1, 70 GB to VM2, 80 GB to VM3, 60 GB to VM4 (total 270 GB) in an X8-2, the maximum available for VM 5 would be 352 - 270 = 82 GB.

In ExaC@C Gen 2, we require a minimum of 60 GB per /u02, so with that minimum size there is a maximum of 5 VMs in X8-2 and 6 VMs in X7-2.

X8M-2 Systems

The maximum number of VMs for an X8M-2 will be 8, regardless of whether there is local disk space or other resources available.

For an X8M-2 system, the fixed consumption per VM is 160 GB.

Total space available to all VMs on an ExaC@C X8M databases node is 2500 GB. Although there is 2500 GB per database node, with a single VM, you can allocate a maximum of 900 GB local storage. Similarly, for the second VM, there is 1800 GB local storage available given the max limit of 900 GB per VM. With the third VM, the amount of space available is 2500 - (160Gb * 3) = 2020 GB. And so on for 4 and more VMs.

- Total space available for VM images (X8M Base System): 1237 GB
- Total space available for VM images (X8M Qtr/Half/Full Racks): 2500 GB
- Fixed storage per VM: 160 GB



#VMs	Fixed Storage All VMs (GB)	X8M-2 Base System Space for All /u02 (GB)	X8M-2 Quarter/Half/ Full Rack Space for All /u02 (GB)
1	160	900	900*
2	320	740	1800*
3	480	580	2020
4	640	420	1860
5	800	N/A	1700
6	960	N/A	1540
7	1120	N/A	1380
8	1280	N/A	1220

Table 5-16 Space allocated to VMs

*Space is limited by 900 GB max per VM

For an X8M-2, to get the max space available for the nth VM, take the number in the table above and subtract anything previously allocated for /u02 to the other VMs. So, for a quarter and larger rack, if you allocated 60 GB to VM1, 70 GB to VM2, 80 GB to VM3, 60 GB to VM4 (total 270 GB) in an X8M-2, the maximum available for VM 5 would be 1700 - 270 = 1430 GB. However, the per VM maximum is 900 GB, so that would take precedent and limits VM5 to 900 GB.

X10M and X9M-2 Systems

- Total Available for VM Images (Base System X9M): 1077 GB
- Total Available for VM Images (Qtr/Half/Full Racks): 2243 GB
- Fixed overhead per VM: 184 GB

Table 5-17 Space allocated to VMs

#VMs	Fixed Storage All VMs (GB)	X9M-2 Base System Space All /u02 (GB)	X10M/X9M-2 Qtr/Half/ Full Racks All /u02 (GB)
1	184	892	900*
2	368	708	1800*
3	552	524	1691
4	736	340	1507
5	920	N/A	1323
6	1104	N/A	1139
7	1288	N/A	955
8	1472	N/A	771

*Space is limited by 900 GB max per VM

Scaling Local Storage

Scale Local Space Operation Guidelines

The scaling local storage operation expects you to input a local space value that you want for the /u02 file system on each node.

Maximum size

The maximum supported size of /u02 is 900 GB per VM. Depending on the number of other VMs in the system and the space consumed by those VMs, you may be subject to a lower maximum for /u02 due to limited space. Refer to tables 5-12 through 5-14.

- Minimum size The minimum supported size of /u02 is 60 GB per VM.
- 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 ${\tt df}_{-kh}$ 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).

• A rolling restart of each VM is required to resize /u02.

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.

Using the Console to Manage VM Clusters on Exadata Cloud@Customer

Learn how to use the console to create, edit, and manage your VM Clusters on Oracle Exadata Cloud@Customer.

- Using the Console to Create a VM Cluster
 To create your 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.



- 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 Exadata Database Service on Cloud@Customer Gen2, 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 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 a VM Cluster

To create your VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

To create a VM cluster, ensure that you have:

- Active Exadata infrastructure is available to host the VM cluster.
- A validated VM cluster network is available for the VM cluster to use.
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Choose the **Region** that contains your Exadata infrastructure.
- 3. Click VM Clusters.
- 4. Click Create VM Cluster.
- 5. Provide the requested information on the Create VM Cluster page:
 - a. Choose 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 Database Service on Cloud@Customer 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. Choose the Oracle Grid Infrastructure version: From the list, choose the Oracle Grid Infrastructure release (19c and 23ai) that you want to install on the VM cluster.

The Oracle Grid Infrastructure release determines the Oracle Database releases that can be supported on the VM cluster. You cannot run an Oracle Database release that is later than the Oracle Grid Infrastructure software release.

- f. Choose an Exadata image 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.
 - 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.
- g. Configure VM Cluster:
 - Click Select DB Servers for VM placement to allocate VM resources.
 - On the **Select DB Servers** dialog, 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.



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 based on hardware generation 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 Changes.
- h. Specify the OCPU count per VM: Specify the OCPU count for each individual VM. The minimum value is 2 OCPUs per VM (for a live VM condition), unless you are specifying zero OCPUs (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 resources. See *Using the Console to Scale the Resources on a VM Cluster*.

The value for OCPU count for the whole VM Cluster will be calculated automatically based upon the per VM OCPU count you have specified and the number of physical Database Servers configured for the system. There is one VM created on each physical Database Server available.

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.

- i. Requested OCPU 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 count per VM field. This field is not editable.
- j. **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.
- k. 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.
- I. Specify the local file system size per VM (GB): Specify the local file system 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 X8-2 and X7-2 infrastructures.



Note that the minimum size of local system storage must be 60 GB. In addition to the 60 GB, each node of the VM must have at least 137 GB free for miscellaneous VM files. Each time when you create a new VM cluster, the space remaining out of the total available space is utilized for the new VM cluster.

For more information and instructions to specify the size for each individual VM, see *Introduction to Scale Up or Scale Down Operations*.

- m. Reserved local storage per VM (GB): Displays the local storage size reserved internally for root file systems, Oracle Grid Infrastructure Homes, and diagnostic logs. This field is not editable.
- n. Configure the Exadata Storage: The following settings define how the Exadata storage is configured for use with the VM cluster. These settings cannot be changed after creating the VM cluster.
 - **Specify Usable Exadata Storage:** Specify the size for each individual VM. The minimum recommended size is 2 TB.
 - 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.

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%

Table 5-18 Storage Allocation



 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.

- p. 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.
 - License Included: Select this option to subscribe to Oracle Database software licenses as part of Exadata Database Service on Cloud@Customer.

q. 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).

r. 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 browserdetected 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**.

- **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**.

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 Exadata Database Service on Cloud@Customer Gen2, 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 to 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 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.
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Choose the **Region** that contains your Exadata infrastructure.
- 3. Click VM Clusters.
- 4. Click the name of the VM cluster you want to enable or disable diagnostic data collection.
- 5. On the VM Cluster Details page, under General Information, enable, partially enable, or disable Diagnostics Collection.
- 6. Click Edit.

Edit Diagnostics Collection Settings window is displayed.

7. Select or clear the checkboxes and then click Save Changes.

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 optin 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 a VM Cluster
 To create your 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.
- Adding a VM to a cluster will not automatically extend any database which is part of a Data Guard configuration (either primary or standby) to the newly provisioned VM.
- 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 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**.
- 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, Available to add, and Insufficient resources. You can only add DB servers with the Available to add status.

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** status and then click **Save Changes**.

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.

- 1. Open the navigation menu. Under Oracle 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 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 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 **Remove**.
- 7. On the Terminate Virtual Machine dialog, enter the name of the virtual machine, and then click **Remove**.

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 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.
- 3. Click VM Clusters.
- 4. Click the name of the VM cluster for which you want to update the license type.

The VM Cluster Details page displays information about the selected VM cluster.

- 5. Click Update License Type.
- 6. In the dialog box, choose one of the following license types and then click **Save Changes**.
 - Bring Your Own License (BYOL): Select this option if your organization already owns Oracle Database software licenses that you want to use on the VM cluster.



• License Included: Select this option to subscribe to Oracle Database software licenses as part of Exadata Database Service on Cloud@Customer.

Updating the license type does not change the functionality or interrupt the operation of the VM cluster.

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

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Choose the **Region** that contains your Exadata infrastructure.
- 3. Click 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 Exadata Database Service on Cloud@Customer Gen2, 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.
- 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
- Memory



- Local storage
- Exadata storage

Each scaling operation can take several minutes to complete. The time for each operation will vary based on activity in the system, but as a general rule, most operations should complete within 15 minutes for a quarter rack, 20 minutes for a half rack, and 30 minutes for a full or larger rack. Performing multiple OCPU scaling operations over a short period of time can lengthen the time for completion. Although online, OCPU 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 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 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 Up/Down.
- 6. In the dialog box, adjust any or all of the following:
 - OCPU Count:

The OCPU 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 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*.

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.

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.

Note:

When you downsize, the new size must be at least 15% more than the currently used size.

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 Changes.

Related Topics

 System Configuration Options for Oracle Exadata Cloud@Customer Exadata Cloud@Customer is now shipping with the latest generation Exadata infrastructure, Exadata X10M.

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

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

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on 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 VM Clusters.
- Click the name of the VM cluster that contains the virtual machine that you want to stop, start, or reboot.

The VM Cluster Details page displays information about the selected VM cluster.

5. In the **Resources** list, click **Virtual Machines**.



The list of virtual machines is displayed.

- 6. In the list of nodes, click the **Actions** icon (three dots) for a node, and then click one of the following actions:
 - **a. Start:** Restarts a stopped node. After the node is restarted, the **Stop** action is enabled.
 - b. Stop: Shuts down the node. After the node is stopped, the Start action is enabled.
 - c. **Reboot:** Shuts down the node, and then restarts it.

Using the Console to Check the Status of a VM Cluster Virtual Machine

Review the health status of a VM cluster virtual machine.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on 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 VM Clusters.
- 4. Click the name of the VM cluster that contains the virtual machine that you are interested in.

The VM Cluster Details page displays information about the selected VM cluster.

5. In the **Resources** list, click **Virtual Machines**.

The list of virtual machines displays. For each virtual machine in the VM cluster, the name, state, and client IP address are displayed.

6. In the node list, find the virtual machine that you are interested in and check its state.

The color of the icon and the associated text it indicates its status.

- Available: Green icon. The node is operational.
- **Starting:** Yellow icon. The node is starting because of a start or reboot action in the Console or API.
- **Stopping:** Yellow icon. The node is stopping because of a stop or reboot action in the Console or API.
- **Stopped:** Yellow icon. The node is stopped.
- **Failed:** Red icon. An error condition prevents the continued operation of the virtual machine.

Using the Console to Move a VM Cluster to Another Compartment

To change the compartment that contains your VM cluster on 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 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 VM Clusters.
- 4. Click the name of the VM cluster that you want to move.

The VM Cluster Details page displays information about the selected VM cluster.

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

Using the Console to Terminate a VM Cluster

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

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

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to terminate.
- 3. Click VM Clusters.
- 4. Click the name of the VM cluster that you want to terminate.

The VM Cluster Details page displays information about the selected VM cluster.

- 5. Click Terminate.
- 6. In the resulting dialog, enter the name of the VM cluster, and click **Terminate VM Cluster** to confirm the action.

Using the API to Manage Exadata Cloud@Customer VM Clusters

Review the list of API calls to manage your 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 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
- 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.
- 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 <code>dbnode-console-connection</code>. The resource name for virtual machine is <code>db-nodes</code>. 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)

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>

Description	
Use the RSA algorithm.	
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.	
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.	
A name to identify the key.	
The location where the key pair will be saved and the root name for the files.	

Create the SSH Key Pair for Windows Using PuTTY

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

Note: Ensure that you are using the latest version of PuTTY, see http://www.putty.org.

- 1. Find puttygen.exe in the PuTTY folder on your computer, for example, C:\Program Files (x86)\PuTTY. Double-click puttygen.exe to open it.
- 2. Specify a key type of SSH-2 RSA and a key size of 2048 bits:
 - In the Key menu, confirm that the default value of SSH-2 RSA key is selected.
 - For the Type of key to generate, accept the default key type of RSA.
 - Set the Number of bits in a generated key to 2048 if not already set.
- 3. Click Generate.
- 4. To generate random data in the key, move your mouse around the blank area in the PuTTY window.

When the key is generated, it appears under **Public key for pasting into OpenSSH** authorized_keys file.

- 5. A **Key comment** is generated for you, including the date and timestamp. You can keep the default comment or replace it with your own more descriptive comment.
- 6. Leave the Key passphrase field blank.
- 7. Click **Save private key**, and then click **Yes** in the prompt about saving the key without a passphrase.

The key pair is saved in the PuTTY Private Key (PPK) format, which is a proprietary format that works only with the PuTTY tool set.

You can name the key anything you want, but use the ppk file extension. For example, mykey.ppk.

8. Select all of the generated 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

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

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

Related Topics

Create SSH Key Pairs

Make an SSH Connection to the Serial Console

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

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



- Validate Server Host Keys
- 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 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**.
- 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:

```
\ensuremath{\mathsf{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.



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.
- 5. Paste the modified connection string into the PowerShell window or your OpenSSH client, and then press **Enter** to connect to the console.
- 6. If prompted, validate and accept the fingerprint of the server host key.

If you have previously accepted a fingerprint for the server host key, but the key has been rotated, a warning appears indicating a potential attack. The warning includes a Host key verification failed error and a line number in your .ssh/known_hosts file. Delete the specified line in your .ssh/known_hosts file and then reconnect to the serial console. Validate and accept the new server host key fingerprint.

- 7. Press Enter again to activate the console.
- 8. Reboot your virtual machine.



You do not need to enter a 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) 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:

- 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



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 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.

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 Database, click Exadata Database Service on Cloud@Customer.
- 2. Click the VM Cluster that you're interested in.
- **3.** n 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. In 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 Database, click Exadata Database Service on Cloud@Customer.
- 2. Click the VM Cluster that you're interested in.
- **3.** n 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. In 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 Database, click Exadata Database Service on Cloud@Customer.
- 2. Click the VM Cluster that you're interested in.



3. n 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. In 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 Database, click Exadata Database Service on Cloud@Customer.
- 2. Click the VM Cluster that you're interested in.
- 3. n 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. In 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 Database, click Exadata Database Service on Cloud@Customer.
- 2. Click the VM Cluster that you're interested in.
- 3. n 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. In the console history list, for the console history capture that you want to view, click the **Actions** menu, and then click **Delete**.
- 6. In the confirmation dialog, click **Delete console history**.

Using the API to Manage the Console History Data

Review the list of API calls to manage console history data.

For information about using the API and signing requests, see REST APIs and Security Credentials. For information about SDKs, see Software Development Kits and Command Line Interface.



For the complete list of APIs, see Database Service API.

Use the following API operations to manage the console history data.

- To capture the console history, use the createDbNodeConsoleHistory method.
- To get details of console history metadata, use the **getDbNodeConsoleHistory** method.
- To get the details of console history content, use the getDbNodeConsoleHistoryContent method.
- To edit console history metadata, use the **updateDbNodeConsoleHistory** method.
- To list console history captures, use the listDbNodeConsoleHistories method.
- To delete console history captures, use the deleteDbNodeConsoleHistory method.

Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems

After you are connected with an instance console connection, you can perform various tasks, such as:

- Edit system configuration files.
- Add or reset the SSH keys for the opc user.
- Reset the password for the opc user.

These tasks require you to boot into a Bash shell in maintenance mode.

- To boot into maintenance mode
- 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**.
- 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.



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 <code>root</code> password would be an acceptable alternative to setting an <code>opc</code> 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 $\scriptstyle\sim$ character at the beginning of a new line is used as an escape character.

1. To exit the serial console, enter:

~.

2. To suspend the SSH session, enter:

~^z

The ^ character represents the CTRL key.

3. To see all the SSH escape commands, enter:

~?

To delete the serial console connection for a Virtual Machine

- 1. Open the navigation menu. Under Oracle Database, click 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 Oracle Database Software Images

Learn about Database Software Image resource type and how you can use it to create Oracle Databases and Oracle Database Homes and to patch databases.

Database Software Images enable you to create a customized Oracle Database software configuration that includes your chosen updates (PSU, RU or RUR), and optionally, a list of one-off (or interim) patches or an Oracle Home inventory file. This reduces the time required



to provision and configure your databases, and makes it easy for your organization to create an approved "gold image" for developers and database administrators.

- Creation and Storage of Database Software Images
 Database software images are resources within your tenancy that you create prior
 to provisioning or patching an Exadata Cloud@Customer instance, a Database
 Home, or a database.
- Using a Database Software Image with an Exadata Cloud@Customer System Create, save, and reuse an Oracle Database Software Image.
- Using the Console to Create a Database Software Image To create an Oracle Database Software Image with the console, use this procedure.
- Using the OPatch Isinventory Command to Verify the Patches Applied to an Oracle Home
 OPatch utility enables you to apply the interim patches to Oracle Database

software. You can find opatch utility in the <code>\$ORACLE_HOME/Opatch</code> directory.

- Using the Console to Delete a Database Software Image
- Using the Console to View the Patch Information of a Database Software Image To view the Oracle Database version, update information (PSU/BP/RU level), and included one-off (interim) patches of a database software image, use the following instructions:
- Using the Console to Move Database Software Image to a Different Compartment Follow these steps to move a Database Software Image to a different compartment of your choice:
- Using the API for Managing Oracle Database Software Images Review the list of API calls to manage Oracle Database Software Image.

Related Topics

- Using the Console to Create Oracle Database Home on 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 a Patch Operation on a Database Home Learn to apply patches on a Database Home.

Creation and Storage of Database Software Images

Database software images are resources within your tenancy that you create prior to provisioning or patching an Exadata Cloud@Customer instance, a Database Home, or a database.

There is no limit on the number of database software images you can create in your tenancy, and you can create your images with any Oracle Database software version and update supported in Oracle Cloud Infrastructure.

Database software images are automatically stored in Oracle-managed Object Storage and can be viewed and managed in the Oracle Cloud Infrastructure Console. Note that database software images incur Object Storage usage costs. Database software image are regional-level resources and can be accessed from any **availability domain** within their region. For information on creating an image, see Using the Console to Create a Database Software Image.

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 Database Software Image with an Exadata Cloud@Customer System

Create, save, and reuse an Oracle Database Software Image.

Creating an Oracle Database Software Image enables you to:

- Create custom Database Home images based on Database Software Images, RU/RUR, and one-off patches.
- Save a custom image automatically to Object Storage as a resource.
- Use a Database Software Image to create Database Homes.
- Patch the Database Home created using the Database Software Image.
- Clone Database Software Image to another service in the Data Guard creation process.

Provisioning: After you create a database software image, you can use it to create an Oracle Database Home in an Exadata Cloud@Customer system. For more information, see Using the Console to Create Oracle Database Home on Exadata Cloud@Customer.

Patching: To patch a database in an Exadata Cloud@Customer system using a custom database software image, create the Database Home using the image, and then move the database to that Database Home. For more information, see *Patching an Exadata Cloud* @*Customer System*.

Setting up Data Guard: When creating an Oracle Data Guard association, the custom Database Software Image associated with the primary database will be used to create a Database Home for the new standby. If the Database Software Image used by the primary database is no longer available, then enable Data Guard operation will not proceed. For more information, see *Using Oracle Data Guard with Exadata Cloud@Customer*.

Note:

The Database 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 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 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 Exadata Database Service on Cloud@Customer Learn to configure and manage Data Guard associations in your VM cluster.

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 Database, click Exadata Cloud@Customer.
- 2. Under Resources, click Database Software Images.
- 3. Click Create Database Software Image.
- 4. In the **Display name** field, provide a display name for your image. Avoid entering confidential information.
- 5. Choose your **Compartment**.
- 6. Choose the **Database version** for your image.
- 7. Choose the **patch set update**, **proactive bundle patch**, **or release update**. For information on Oracle Database patching models, see *Release Update Introduction and FAQ (Doc ID 2285040.1)*.
- 8. Optionally, you can enter a comma-delimited list of one-off (interim) patch numbers.
- 9. Optionally, you can upload an Oracle Home inventory file from an existing Oracle Database. For more information, see Using the OPatch Isinventory Command to Verify the Patches Applied to an Oracle Home.
- 10. 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.
- 11. Click Create Database Software Image.

Related Topics

- Release Update Introduction and FAQ (Doc ID 2285040.1)
- Using the OPatch Isinventory Command to Verify the Patches Applied to an Oracle Home

OPatch utility enables you to apply the interim patches to Oracle Database software. You can find opatch utility in the <code>\$ORACLE_HOME/Opatch</code> directory.

Resource Tags



Using the OPatch Isinventory Command to Verify the Patches Applied to an Oracle Home

OPatch utility enables you to apply the interim patches to Oracle Database software. You can find opatch utility in the $\clicklinese{thmmedol}$ directory.

1. Run the opatch lsinventory command to get the list of interim patches applied.

```
$ORACLE_HOME/OPatch/opatch lsinventory
Oracle Interim Patch Installer version 12.2.0.1.21
Copyright (c) 2021, Oracle Corporation. All rights reserved.
Oracle Home : /u02/app/oracle/product/19.0.0.0/dbhome_2
Central Inventory : /u01/app/oraInventory
from : /u02/app/oracle/product/19.0.0.0/dbhome_2/oraInst.loc
OPatch version : 12.2.0.1.21
OUI version : 12.2.0.7.0
Log file location : /u02/app/oracle/product/19.0.0.0/dbhome_2/cfgtoollogs/
opatch/opatch2021-01-21_09-22-45AM_1.log
```

```
Lsinventory Output file location : /u02/app/oracle/product/19.0.0.0/
dbhome 2/cfgtoollogs/opatch/lsinv/lsinventory2021-01-21 09-22-45AM.txt
```

2. Use the lsinventory output file to extract the additional One Off Patches applied to a specific Oracle Home.

Using the Console to Delete a Database Software Image

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Under Resources, click Database Software Images.
- 3. In the list of database 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. Enter the Database Software Image name to confirm the deletion and click **Delete Database Software Image**.

Using the Console to View the Patch Information of a Database Software Image

To view the Oracle Database version, update information (PSU/BP/RU level), and included one-off (interim) patches of a database software image, use the following instructions:

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Under Resources, click Database Software Images.
- 3. In the list of database software images, find the image you want to view and click on the display name of the image.



- 4. On the Database Software Image Details page for your selected image, details about the image are displayed:
 - The Oracle Database version is displayed in the **General Information** section. For example: 19.0.0.0
 - The **PSU/BP/RU** field of the **Patch Information** section displays the update level for the image. For example: 19.5.0.0
 - The **One-Off Patches** field displays the number of one-off patches included in the image if any. The count includes all patches specified when creating the image (excluding patches listed in lsinventory). To view the included patches (if any are included), click the **Copy All** link and paste the list of included patches into a text editor. The copied list of patch numbers is comma-delimited and can be used to create additional Database Software Images.

Using the Console to Move Database Software Image to a Different Compartment

Follow these steps to move a Database Software Image to a different compartment of your choice:

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Under Resources, click Database Software Images.
- 3. In the list of database software images, find the image you want to delete and click the action icon (three dots) at the end of the row.
- 4. Click Move Resource.
- 5. Choose a new compartment in the **Move Resource to a Different Compartment** dialog.
- 6. Click Move Resource.

Using the API for Managing Oracle Database Software Images

Review the list of API calls to manage Oracle Database Software Image.

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 Exadata Database Service on Cloud@Customer System

Learn to create Oracle Database Homes on Exadata Database Service on Cloud@Customer.

- About Creating Oracle Database Homes on an 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 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 Exadata Cloud@Customer To create an Oracle Database home, review the list of API calls.

About Creating Oracle Database Homes on an 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 Exadata Database Service on Cloud@Customer Learn to use the dbaascli utility on Exadata Cloud@Customer.

Using the Console to Create Oracle Database Home on 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.

1. Open the navigation menu. Under Oracle 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.
 - **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 select an **Image Type**.

Oracle Provided Database Software Images: These images contain generally available versions of Oracle Database software.

Custom Database Software Images: These images are created by your organization and contain customized configurations of software updates and patches. Use the **Select a compartment** and **Select a Database version** selectors to limit the list of custom Database Software Images to a specific compartment or Oracle Database software major release version.

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.

Related Topics

Resource Tags

Using the API to Create Oracle Database Home on Exadata 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 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

Manage Oracle Database Homes on Exadata Database Service on Cloud@Customer Systems

Learn to manage Oracle Database homes on Exadata Database Service on Cloud@Customer.

- About Managing Oracle Database Homes on 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 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 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 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 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 Exadata Database Service on Cloud@Customer Learn to use the dbaascli utility on Exadata Cloud@Customer.

Manage Database Home Using the Console

Use the OCI console to manage the various operations needed on a Database Home.



- 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 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 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.



Using the API to Manage Oracle Database Home on 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 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



Operation	Host Tooling Parameters (dbaascli / dbaasapi)	Database Service API Parameters
Updating Database Home	DatabaseVersion, dbHomeLocation	dbVersion, dbHomeLocation
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 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 Exadata Database Service on Cloud@Customer Learn to manage pluggable databases on Exadata 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.

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.

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 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

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.
 - 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.

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 Move a Database to Another Database Home Learn to move a database to another Database Home.
- 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 Database

To create an Oracle Database with the console, use this procedure.

- Open the navigation menu Under Oracle Database, and 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. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
 VM Clusters is selected by default.
- b. Choose your Compartment.
 A list of VM Clusters is displayed for the chosen Compartment.
- c. 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.
- d. Click Database Homes.
- e. Click the name of the Database Home where you want to create the database.
- f. 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

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.
- **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 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.

- **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 <code>Oracle or Table</code>, 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:
 - None: Select to not define a backup configuration for the database.
 - Local: Select to store backups locally in the Oracle Exadata Storage Servers on your Oracle Exadata Cloud at Customer system.
 This option is available only if you enabled backups on local Oracle Exadata storage in the VM cluster that you want to host the database.
 - 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. However, if you select **None** now, then you can select a backup destination after the database is created.

 Enable automatic backups: Select this option to enable daily backups using the policy for automatic backups.

This option is only enabled when you select a **Backup Destination Type** other than **None**. You can change this setting after database creation.

- (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 <code>instance_name</code> 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
- Backup retention period: From the list, you can choose the length of time that you want automatic backups to be 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 Oracle Zero Data Loss 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.

- 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 Database.

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.

Using the Console to Manage SYS User and TDE Wallet Passwords

Learn to manage administrator (SYS user) and TDE wallet passwords.

- 1. Open the navigation menu. Click Oracle Database, then click 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.
- 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**.
- 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.
- In the resulting Manage passwords dialog, click Update Administrator Password or Update TDE Wallet Password.
 Depending on the option way called the system diaplays the fields to edit.

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 Move a Database to Another Database Home

Learn to move a database to another Database Home.

1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on 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 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 Move Database.
- 7. In the resulting dialog, select the target Database Home.

Note:

Oracle recommends using Database Homes, which are running the latest (N) to 3 versions from the latest (N-3) RU versions when updating the software version of the database by moving them to a target DB Home. Only DB Homes provisioned with database versions, which meet this best practice criterion are available as target homes to move your database.

8. Click Move Database.

The database will be stopped in the current home and then restarted in the destination home. While the database is being moved, the Database Home status displays as **Moving Database**. When the operation completes, Database Home is updated with the current home. If the operation is unsuccessful, the status of the database displays as **Failed**, and the Database Home field provides information about the reason for the failure.

Using the 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 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 Terminate.
- 7. In the resulting dialog, enter the name of the database, and then click **Terminate Database** to confirm the action.

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:

- 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.

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 Exadata Database Service on Cloud@Customer

Learn to manage pluggable databases on Exadata Cloud@Customer.

Pluggable Database Operations You can create and manage pluggable databases (PDBs) in Oracle Exadata Cloud@Customer systems using the Console and APIs.

Pluggable Database Operations

You can create and manage pluggable databases (PDBs) in Oracle Exadata 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

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

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.

- Oracle recommends using the Console or API-based tools (including the OCI CLI, SDKs, and Terraform) to create and manage PDBs. However, there would be periodic sync of the PDBs created through DBAASCLI and SQL*Plus.
- PDB management using the OCI Console and API is available only for Oracle Database versions 19c and later.



- PDBs are backed up at the CDB level, and each backup includes all the PDBs in the database. OCI Control Plane does not support the creation of backups for individual PDBs. However, bkup_api Backup operations. For more information, see Configuring and Customizing Backups with bkup_api. Examples:
 - List backups:

/var/opt/oracle/bkup api/bkup api list --dbname psarch

Create initial PDB backup manually:

```
/var/opt/oracle/bkup_api/bkup_api bkup_start --level1 --dbname
psarch --pdb NEWPDBA
```

- Restore operations are performed at the CDB level. OCI Control Plane does not support restoring individual PDBs. However, bkup_api supports PDB restore operations.
 Examples:
 - Recover a PDB with least or no data loss possible:

```
/var/opt/oracle/bkup_api/bkup_api recover_start --latest --
dbname psarch --pdb NEWPDBA
```

Recover a PDB back to a point in time:

/var/opt/oracle/bkup_api/bkup_api recover_start -t '17-AUG2021
21:15:00' --dbname psarch --pdb NEWPDBA

Recover until SCN:

/var/opt/oracle/bkup_api/bkup_api recover_start -scn 138935800 pdb=NEWPDBA -uuid=fec6579e077211ec8b0a00102ee75632 -bname=psarch

Related Topics

Configuring and Customizing Backups with bkup_api
 In addition to the console-based automated backup option, there is a command line backup utility, bkup_api, which can allow for further customization. If configuring backups using bkup_api instead of the console, then do not enable backups for your database in the console.

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.
- Using the API to Create a Pluggable Database
 Use various API features to help create your pluggable databases on Oracle
 Exadata Cloud@Customer.



Using the Console to Create a Pluggable Database

To create a pluggable database with the console, use this procedure.

- 1. Open the navigation menu Under Oracle Database, and 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 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.



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< td=""><td>NAME</td><td>></td><td></td><td></td></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 Create Pluggable Database.

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.

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.

- Open the navigation menu Under Oracle Database, and 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.
- Click the name of the PDB that you want to relocate.
 From the Pluggable Database details page, click More Actions, and then select Relocate.

(or)

Click the Actions menu (three dots) and select Relocate.

- 6. In the resulting Relocate Pluggable Database window, enter the following:
 - VM Cluster: Use the menu to select the destination VM cluster.
 - Destination database: Use the menu to select an existing database where the PDB will be created. This database can be of the same version as the CDB the source PDB is in or of a higher version.
 - New PDB name for the clone: The name must begin with an alphabetic character and can contain up to 30 characters. To keep the PDB name the same, just re-enter the source PDB name.
 - Database TDE wallet password: Enter the TDE wallet password for the parent CDB of the source PDB.



- Unlock my PDB Admin Account:
 - To enter the administrator's password, check this check box.
 - * **PDB Admin Password:** Enter PDB admin password. The password must contain:
 - * a minimum of 9 and a maximum of 30 characters
 - * at least two uppercase characters
 - * at least two lowercase characters
 - * at least two special characters. The valid special characters are underscore (_), a pound or hash sign (#), and dash (-). You can use two of the same characters or any combination of two of the same characters.
 - * at least two numeric characters (0 9)
 - * **Confirm PDB Admin Password:** Enter the same PDB Admin password in the confirmation field.
 - To skip entering the administrator's password, uncheck this check box. If you uncheck this check box, then the PDB is created but you cannot use it. To use the PDB, you must reset the administrator password.



When you create a new PDB, a local user in the PDB is created as the administrator and granted the $\tt PDB_DBA$ 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
```

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 autobackup 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 pluggable database.

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
- Resource Tags



Using the API to Create a Pluggable Database

Use various API features to help create your pluggable databases on Oracle Exadata 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 Exadata 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.
- Using the API to Manage Pluggable Databases
 Use various API features to help manage your pluggable databases on Oracle
 Exadata 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.

- Open the navigation menu Under Oracle Database, and 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 PDB 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.

- Open the navigation menu Under Oracle Database, and 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 PDB** to confirm the stop operation.



Using the Console to Delete a Pluggable Database

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

- Open the navigation menu Under Oracle Database, and 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 PDB** to confirm the delete operation.

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.

- Open the navigation menu Under Oracle Database, and 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.
- 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.

The following types of clones are supported:

- Local clone: A clone of the PDB is created within the same CDB.
- **Remote clone**: A clone of the PDB is created in a different CDB.
- **Refreshable clone**: A clone of the PDB is created in a different CDB, and you will be able to refresh the cloned PDB.
- Using the Console to Create a Local Clone of a Pluggable Database (PDB) 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.

- Open the navigation menu Under Oracle Database, and 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 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.

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.
```

• Take a backup of the PDB immediately after creating it: You must enable autobackup 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.

9. Click Clone pluggable database.

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.

- Open the navigation menu Under Oracle Database, and 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 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:

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.
- Take a backup of the PDB immediately after creating it: You must enable autobackup 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.
- 9. Click Clone pluggable database.

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.

- Open the navigation menu Under Oracle Database, and 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.
 - **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.

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
```

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.
- 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 pluggable database.
- 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.

Related Topics

Resource Tags

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

To create a refresh a cloned PDB, follow this procedure.

- Open the navigation menu Under Oracle Database, and 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.

- Open the navigation menu Under Oracle Database, and 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.
- 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.

 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 autobackup 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.

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.

- Open the navigation menu Under Oracle Database, and 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.
- 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.

Using the API to Manage Pluggable Databases

Use various API features to help manage your pluggable databases on Oracle Exadata 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
- 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

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.
- Enable, Disable, and Re-enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database Learn to enable, disable, and re-enable Identity and Access Management (IAM) Authentication on Oracle Database. Also, to change the external identity provider from (IAM) authentication and authorization to another and vice-versa.
- Manage Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies, Users, Roles, and Database Passwords
 Learn to create and manage IAM policies, add IAM users to the Oracle Database and grant global roles, and create IAM database passwords for IAM users.



- Connect to Oracle Database with Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication
 After the DBA user enables Oracle Cloud Infrastructure IAM on Oracle Database, users log in to the Oracle Database instance using their Oracle Cloud Infrastructure IAM credentials or access the database through an Oracle Cloud Infrastructure IAM database token.
- Database Links in an Oracle DBaaS-to-IAM Integration The use of database links when accessing the Oracle DBaaS database using IAM credentials is supported.
- Configuring Authorization for IAM Users and Oracle Cloud Infrastructure
 Applications

An Oracle DBaaS database administrator can map IAM users and Oracle Cloud Infrastructure (OCI) applications to the Oracle Database global schemas and global roles.

Configure Proxy Authentication

Proxy authentication allows an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) user to proxy to a database schema for tasks such as application maintenance.

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.

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 non-default custom domains are only available with Oracle Database Release 19c, Version 19.21 and higher (but not Oracle Database Release 21c).

Oracle Cloud Infrastructure IAM integration with Oracle Exadata Database Service on Cloud@Customer supports the following:

- IAM Database Password Authentication
- Identity and Access Management (IAM) SSO Token Based Authentication

See Authenticating and Authorizing IAM Users for Oracle DBaaS Databases for complete details about the architecture for using IAM users on Oracle Exadata Database Service on Cloud@Customer.

Related Topics

Authenticating and Authorizing IAM Users for Oracle DBaaS Databases

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 *Managing User Credentials*.

Related Topics

• Managing User Credentials



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.

For token 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.

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 dbtoken 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.
- 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.

Related Topics

- Supported Client Drivers for IAM Connections
- db-token get
- Configuring a Client Connection for SQL*Plus That Uses an IAM Token



Authenticating and Authorizing IAM Users for Oracle DBaaS Databases

Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database

Review the prerequisites for Identity and Access Management (IAM) authentication on an Oracle Database.

- 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 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.
- 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.

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.

- 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 Logins*.

Related Topics

- Identity and Access Management Service API
- Network Requirements for Oracle Exadata Database Service on Cloud@Customer
- Troubleshooting IAM Logins

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 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) Oneway TLS is being configured where the client does not have its own certificate and 2) the root certificate that signed the database server certificate is stored in the system default certificate store. The root certificate would most likely already be there if the server certificate is signed by a common certificate authority. If it's a self-signed certificate, then the root certificate would need to be installed in the system default certificate store to avoid using a client wallet.

For details on how to configure TLS between the database client and database server including the options described above, see *Configuring Transport Layer Security Authentication* in the *Oracle Database Security Guide*.

If you choose to use self-signed certificates and for additional wallet related tasks, *Managing Public Key Infrastructure (PKI) Elements* in the *Oracle Database Security Guide*.

Related Topics

- Configuring Transport Layer Security Authentication
- Managing Public Key Infrastructure (PKI) Elements

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
```

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

Learn to enable, disable, and re-enable Identity and Access Management (IAM) Authentication on Oracle Database. Also, to change the external identity provider from (IAM) authentication and authorization to another and vice-versa.

- Enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database Review the steps to enable IAM user access to Oracle Database.
- Change External Identity Providers on Oracle Exadata Database Service on Cloud@Customer
 Review the steps to change the external identity provider from (IAM) authentication and authorization to another and vice-versa.
- Re-enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication and Authorization Review the steps to re-enable IAM users to connect to Oracle Database using Oracle Cloud Infrastructure (IAM) Authentication and Authorization
- Disable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database Describes the steps to disable IAM external authentication user access for Oracle Database.
- Using Oracle Database Tools with Identity and Access Management (IAM) Authentication Review the notes for using Oracle Database tools with IAM authentication enabled.

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

Review the steps to enable IAM user access to Oracle Database.

Note:

Oracle Database supports the Oracle DBaaS integration for Oracle Cloud Infrastructure (OCI) IAM with identity domains as well as the legacy IAM, which does not include identity domains. Both default and non-default domain users and groups are supported when using IAM with Identity Domains.

1. Perform the prerequisites for IAM authorization and authentication on Oracle Database.

See Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database for more information.



2. Enable Oracle Cloud Infrastructure (IAM) Authentication and Authorization using the ALTER SYSTEM command.

ALTER SYSTEM SET IDENTITY PROVIDER TYPE=OCI IAM SCOPE=BOTH;

3. Verify the value of IDENTITY PROVIDER TYPE system parameter.

Related Topics

- Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database Review the prerequisites for Identity and Access Management (IAM) authentication on an Oracle Database.
- Disable External Authentication Scheme Review the prerequisites for enabling IAM user access to Oracle Database.

Change External Identity Providers on Oracle Exadata Database Service on Cloud@Customer

Review the steps to change the external identity provider from (IAM) authentication and authorization to another and vice-versa.

If Oracle Cloud Infrastructure (IAM) authentication and authorization for users is enabled and you wish to switch to a different external service (CMU for Active Directory, EUS for OID or OUD), you must first disable IAM integration before you enable the other integration. There can only be one external authentication scheme enabled at any given time. If IAM and another directory service is configured at the same time, IAM integration will take precedence.

1. Disable IAM integration using the ALTER SYSTEM command.

ALTER SYSTEM RESET IDENTITY PROVIDER TYPE SCOPE=BOTH;

2. Verify the value of IDENTITY PROVIDER TYPE system parameter.

3. Configure the other directory service integration.

Re-enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication and Authorization

Review the steps to re-enable IAM users to connect to Oracle Database using Oracle Cloud Infrastructure (IAM) Authentication and Authorization

1. Disable the integration with the other identity provider or directory service.



2. Enable IAM integration as described in *Enable Identity and Access Management* (IAM) Authentication on Oracle Database.

Related Topics

 Enable Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database Review the steps to enable IAM user access to Oracle Database.

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

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

To disable IAM user access on your Oracle Database instance:

1. Disable IAM integration using the ALTER SYSTEM command.

ALTER SYSTEM RESET IDENTITY PROVIDER TYPE SCOPE=BOTH;

2. If you also want to remove IAM user access to the database, you may need to remove or modify the IAM group and the IAM policies you set up to allow access to the database.

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

Review the notes for using Oracle Database tools with IAM authentication enabled.

- Oracle APEX is not supported for IAM users with Oracle Database.
- Database Actions is not supported for IAM users with Oracle Database. See *Provide Database Actions Access to Database Users* for information on using regular database users with Oracle Database.
- Oracle Machine Learning Notebooks and other components are not supported for IAM Authorized users with Oracle Database. See Add Existing Database User Account to Oracle Machine Learning Components for information on using regular database users with Oracle Database.

Related Topics

- Provide Database Actions Access to Database Users
- Add Existing Database User Account to Oracle Machine Learning Components

Manage Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies, Users, Roles, and Database Passwords

Learn to create and manage IAM policies, add IAM users to the Oracle Database and grant global roles, and create IAM database passwords for IAM users.

 Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies for IAM Users
 Review the the steps to write policy statements for an IAM group to enable IAM user access to OCI resources, specifically the Oracle Database instances.



- Add Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users on Oracle Database
 - Review the steps to authorize IAM users on an Oracle Database instance.
- Add Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Roles on Oracle Database
 Review the steps to map Oracle Database global roles to IAM groups.
- Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM)
 Database Password for IAM Users
 To add an IAM user and allow the IAM user to login to Oracle Database by supplying a
 username and password, you must create an IAM database password.

Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Groups and Policies for IAM Users

Review the the steps to write policy statements for an IAM group to enable IAM user access to OCI resources, specifically the Oracle Database instances.

A policy is a group of statements that specifies who can access particular resources, and how. Access can be granted for the entire tenancy, databases in a compartment, or individual databases. This means you write a policy statement that gives a specific group a specific type of access to a specific type of resource within a specific compartment.

Note:

Defining a policy is required to use IAM tokens to access the database. A policy is not required when using IAM database passwords to access the Oracle Database.

To enable the Oracle Database to allow OCI IAM users to connect to the database using OCI IAM tokens:

1. Create an IAM group for IAM users to access OCI databases. Add users to this group.

For example, create the group *sales_dbusers*. For more information, see *Managing Groups*.

- 2. Write policy statements to enable access to Oracle Cloud Infrastructure resources.
 - a. In the Oracle Cloud Infrastructure console, click **Identity and Security**, and then click **Policies**.
 - **b.** To a write policy, click **Create Policy**, and then enter a **Name** and a **Description**.
 - c. Use the Policy Builder to create a policy.

For example, to create a policy to allow users in IAM group *DBUsers* to access any Oracle Database in their tenancy:

Allow group DBUsers to use database-connections in tenancy

For example to create a policy that limits members of *DBUsers* group to access Oracle Databases in the compartment *testing_compartment* only:

allow group *DBUsers* to use database-connections in compartment *testing compartment*



For example, to create a policy that limits group access to a single database in a compartment:

allow group DBUsers to use database-connections in compartment testing_compartment where target.database.id = 'ocid1.autonomousdatabase.oc1.iad.aaaabbbbbcccc'

d. Click Create.

Note:

Note the following for creating policies for use with IAM users on database in the ExaDB-C@C service.

- Policies can allow IAM users to access Oracle Database instances across the entire tenancy, in a compartment, or can limit access to a single Oracle Database instance.
- You can use either instance principal or resource principal to retrieve database tokens to establish a connection from your application to an Oracle Database instance. If you are using an instance pricipal or resource principal, you must map a dynamic group. Thus, you cannot exclusively map instance and resource principals; you only can map them through a shared mapping and putting the instance or resource instance in an IAM dynamic group. You can create Dynamic Groups and reference dynamic groups in the policies you create to access Oracle Cloud Infrastructure. See Accessing Cloud Resources by

Configuring Policies and Roles and Managing Dynamic Groups for details.

Related Topics

- Managing Groups
- Accessing Cloud Resources by Configuring Policies and Roles
- Managing Dynamic Groups

Add Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Users on Oracle Database

Review the steps to authorize IAM users on an Oracle Database instance.

To add IAM users to allow access to Oracle Database, map database global users to IAM groups or users with CREATE USER or ALTER USER statements with IDENTIFIED GLOBALLY AS clause.

An IAM user must be mapped to one schema to be authorized to access the database. This could be an exclusive schema or a shared schema.

The authorization of IAM users to an Oracle Database instance works by mapping IAM global users (schemas) to IAM users (exclusive mapping) or IAM groups (shared schema mapping).

1. Log in as the ADMIN user to the database that is enabled to use IAM.

The ADMIN user has the required CREATE USER and ALTER USER system privileges that you need for these steps.



2. Create a mapping between the Oracle Database user (schema) with CREATE USER or ALTER USER statements and include the IDENTIFIED GLOBALLY AS clause, specifying the IAM group name.

Use the following syntax to map a global user to an IAM group:

```
CREATE USER global_user IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=IAM_GROUP_NAME';
```

For example, to map an IAM group named *db_sales_group* to a shared database global user named *sales_group*:

```
CREATE USER sales_group IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=db_sales_group';
```

This creates a shared global user mapping. The mapping, with the global user *sales_group* is effective for all users in the IAM group. Thus, anyone in the *db_sales_group* can log in to the database using their IAM credentials through the shared mapping of the *sales_group* global user.

If you want to create additional global user mappings for other IAM groups or users, follow these steps for each IAM group or user.

Note:

Database users that are not IDENTIFIED GLOBALLY can continue to login as before, even when the Oracle Database is enabled for IAM authentication.

To Exclusively Map a Local IAM User to an Oracle Database Global User

1. Log in as the ADMIN user to the database that is enabled to use IAM.

The ADMIN user has the required CREATE USER and ALTER USER system privileges that you need for these steps.

2. Create a mapping between the Oracle Database user (schema) with CREATE USER or ALTER USER statements and include the IDENTIFIED GLOBALLY AS clause, specifying the IAM local IAM user name.

For example, to create a new database global user named *peter_fitch* and map this user to an existing local IAM user named *peterfitch*:

CREATE USER peter_fitch IDENTIFIED GLOBALLY AS 'IAM PRINCIPAL NAME=peterfitch'

Add Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Roles on Oracle Database

Review the steps to map Oracle Database global roles to IAM groups.

Optionally, create global roles to provide additional database roles and privileges to IAM users when multiple IAM users are mapped to the same shared global user.



Using global roles is optional when a user is mapped exclusively to a global schema or mapped to a shared schema. For example, all privileges and roles can be granted to the shared schema and all IAM users who map to the shared schema would be granted the privileges and roles assigned to the shared schema.

Use a global role to optionally differentiate users who use the same shared schema. For example, a set of users can all have the same shared schema and the shared schema could have the CREATE SESSION privilege. Then global roles can be used to provide differentiated privileges and roles assigned to different groups of users who all use the same shared schema.

Granting additional roles to IAM users in Oracle Database works by mapping Oracle Database global roles to IAM groups.

1. Log in as the ADMIN user to the database that is enabled to use IAM.

The ADMIN user has the required CREATE USER and ALTER USER system privileges that you need for these steps.

2. Set database authorization for Oracle Database roles with CREATE ROLE or ALTER ROLE statements and include the IDENTIFIED GLOBALLY AS clause, specifying the IAM group name.

Use the following syntax to map a global role to an IAM group:

CREATE ROLE global_role IDENTIFIED GLOBALLY AS 'IAM GROUP NAME=IAM GROUP of WHICH the IAM USER IS a MEMBER';

For example, to map an IAM group named *ExporterGroup* to a shared database global role named *export_role*:

CREATE ROLE export_role IDENTIFIED GLOBALLY AS
'IAM GROUP NAME=ExporterGroup';

3. Use the GRANT statements to grant the required privileges or other roles to the global role.

GRANT CREATE SESSION TO export role;

GRANT DWROLE TO export role;

4. If you want an existing database role to be associated with an IAM group, then use the ALTER ROLE statement to alter the existing database role to map the role to an IAM group.

Use the following syntax to alter an existing database role to map it to an IAM group:

```
ALTER ROLE existing_database_role IDENTIFIED GLOBALLY AS 'IAM GROUP NAME=IAM Group Name';
```

Follow these steps for each IAM group to add additional global role mappings for other IAM groups.



Create Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password for IAM Users

To add an IAM user and allow the IAM user to login to Oracle Database by supplying a username and password, you must create an IAM database password.

For more information, see Working with IAM Database Passwords.

Related Topics

Working with IAM Database Passwords

Connect to Oracle Database with Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication

After the DBA user enables Oracle Cloud Infrastructure IAM on Oracle Database, users log in to the Oracle Database instance using their Oracle Cloud Infrastructure IAM credentials or access the database through an Oracle Cloud Infrastructure IAM database token.

After you enable Oracle Cloud Infrastructure IAM user access, you can also log in to the Oracle Database using your local database account username and password (non-global database user account).

You can use a database client to access an Oracle Database instance as an Oracle Cloud Infrastructure IAM user. Enter the IAM user name and IAM database password (not the Oracle Cloud Infrastructure console password) using any currently supported database client. The only constraint is that the database client version be either Oracle Database release 12.1.0.2 or later (or patched) to allow Oracle Database 12c passwords. The database client must be able to use the 12C password verifier. Using the 11G verifier encryption is not supported with IAM.

Alternatively, you can use an Oracle Cloud Infrastructure IAM database token to access an Oracle Database instance with supported clients. IAM database token usage requires the Oracle Database client 19.16 and above (not 21c). Limited (not full) IAM database token capabilities are available with some Oracle Database clients 21.5 and above.

The following examples show password verifier with SQL*Plus to access the database with an Oracle Cloud Infrastructure IAM username and password and the steps required to use SQL*Plus with an Oracle Cloud Infrastructure IAM database token.

Note:

If your Oracle Database instance is in Restricted mode, only the users with the RESTRICTED SESSION privilege can connect to the database.

- About Connecting to an Oracle Database Instance Using Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM)
 IAM users can connect to the Oracle Database instance by using either an IAM database password verifier or an IAM token.
- Configure a Client Connection for SQL*Plus That Uses an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password You can configure SQL*Plus to use an IAM database password.



- Configure a Client Connection for SQL*Plus That Uses an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Token You can configure a client connection for SQL*Plus that uses an IAM token.
- Use Instance Principal to Access Oracle Database with Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication After the DBA user enables Oracle Cloud Infrastructure IAM on Oracle Database, an application can access the database through an Oracle Cloud Infrastructure IAM database token using an instance principal.

Related Topics

- JDBC and UCP Downloads
- Support for Token-Based Authentication
- Identity and Access Management (IAM) Token-Based Authentication
- Oracle Data Access Components .NET Downloads
- NuGet Gallery
- Visual Studio Code Marketplace

About Connecting to an Oracle Database Instance Using Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM)

IAM users can connect to the Oracle Database instance by using either an IAM database password verifier or an IAM token.

Using the IAM database password verifier is similar to the Oracle Database password authentication process. However, instead of the password verifier (encrypted hash of the password) being stored in the Oracle Database, the verifier is instead stored as part of the Oracle Cloud Infrastructure (OCI) IAM user profile.

The second connection method, the use of an IAM token for the database, is more modern. The use of token-based access is a better fit for Cloud resources such as Oracle Database. The token is based on the strength that the IAM endpoint can enforce. This can be multi-factor authentication, which is stronger than the use of passwords alone. Another benefit of using tokens is that the password verifier (which is considered sensitive) is never stored or available in memory. A TCPS (TLS) connection is required when using tokens for database access.

Note:

You cannot configure native network encryption when passing an IAM token. Only Transport Layer Security (TLS) by itself is supported, not native network encryption or native network encryption with TLS.

- Client Connections That Use an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier
 After you have configured the authorization needed for the IAM user, this user can log in using an existing client application, such as SQL*Plus or SQLcl without additional configuration.
- Client Connections That Use a Token Requested by a Client Application or Tool For IAM token access to the Oracle DBaaS, the client application or tool requests a database token from IAM for the IAM user.



Client Connections That Use a Token Requested by an IAM User Name and Database
 Password

You can create a client connection that uses a token requested by an IAM user name and database password.

 Configure a Secure External Password Store Wallet to Retrieve an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Token You can enable an IAM user name and a secure external password store (SEPS) to request the IAM database token.

Client Connections That Use an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier

After you have configured the authorization needed for the IAM user, this user can log in using an existing client application, such as SQL*Plus or SQLcl without additional configuration.

The IAM user enters the IAM user name and IAM database password (not the Oracle Cloud Infrastructure console password) using any currently supported database client. The only constraint is that the database client version be either Oracle Database release 12.1.0.2 or later (or patched) to allow Oracle Database 12c passwords. The database client must be able to use the 12C password verifier.

Using the 11G verifier encryption is not supported with IAM. No special client or tool configuration is needed for the IAM user to connect to the Oracle Database instance.

Client Connections That Use a Token Requested by a Client Application or Tool

For IAM token access to the Oracle DBaaS, the client application or tool requests a database token from IAM for the IAM user.

The client application will pass the database token directly to the database client through the database client API.

If the application or tool has not been updated to request an IAM token, then the IAM user can use Oracle Cloud Infrastructure (OCI) command line interface (CLI) to request and store the database token. You can request a database access token (db-token) using the following credentials:

- Security tokens (with IAM authentication), delegation tokens (in the OCI cloud shell) and API-keys, which are credentials that represent the IAM user to enable the authentication
- Instance principal tokens, which enable instances to be authorized actors (or principals) to perform actions on service resources after authenticating
- Resource principal token, which is a credential that enables the application to authenticate itself to other Oracle Cloud Infrastructure services
- Using an IAM username and IAM database password (can only be requested by database client).

When the IAM users log into the client with a slash / login and the ocl_IAM parameter is configured (sqlnet.ora, tnsnames.ora, or as part of a connect string), then the database client retrieves the database token from a file. If the IAM user submits a username and password, the connection will use the IAM database verifier access described for client connections that use IAM database password verifiers. The instructions in this guide show how to use the OCI CLI as a helper for the database token. If the application or tool has been updated to work with IAM, then follow the instructions for the application or tool. Some



common use cases include the following: SQLPlus on-premises, SQLcl on-premises, SQL*Plus in Cloud Shell, or applications that use SEP wallets.

Related Topics

Client Connections That Use an IAM Database Password Verifier

Client Connections That Use a Token Requested by an IAM User Name and Database Password

You can create a client connection that uses a token requested by an IAM user name and database password.

For more information, see:

- About Client Connections That Use a Token Requested by an IAM User Name and Database Password
- Parameters to Set for Client Connections That Use a Token Requested by an IAM User Name and Database Password
- Configuring the Database Client to Retrieve a Token Using an IAM User Name and Database Password
- Configuring a Secure External Password Store Wallet to Retrieve an IAM Token

Related Topics

- About Client Connections That Use a Token Requested by an IAM User Name and Database Password
- Parameters to Set for Client Connections That Use a Token Requested by an IAM User Name and Database Password
- Configuring the Database Client to Retrieve a Token Using an IAM User Name and Database Password
- Configuring a Secure External Password Store Wallet to Retrieve an IAM Token

Configure a Secure External Password Store Wallet to Retrieve an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Token

You can enable an IAM user name and a secure external password store (SEPS) to request the IAM database token.

- 1. Log in to the Oracle Database client.
- 2. Configure this client to use the secure external password store.
- 3. Set the appropriate parameters to retrieve a token that will be requested by an IAM user name and database password.

Configure a Client Connection for SQL*Plus That Uses an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password

You can configure SQL*Plus to use an IAM database password.

• As the IAM user, log in to the Oracle Database.

```
CONNECT user_name@db_connect_string
Enter password: password
```



In this specification, user_name is the IAM user name. There is a limit of 128 bytes for the combined *domain_name/user_name*.

The following example shows how IAM user *peter_fitch* can log in to an Oracle Database instance.

```
sqlplus /nolog
connect peter_fitch@db_connect_string
Enter password: password
```

Some special characters will require double quotation marks around user_name and password. For example:

```
"peter_fitch@example.com"@db_connect_string
"IAM database password"
```

Configure a Client Connection for SQL*Plus That Uses an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Token

You can configure a client connection for SQL*Plus that uses an IAM token.

- 1. Ensure you have an IAM user account.
- Check with an IAM administrator and Oracle Database administrator to ensure you have a policy allowing you to access the database in the compartment or your tenancy and that you are mapped to a global schema in the database.
- If your application or tool does not support direct IAM integration, then download, install, and configure the OCI CLI.

See OCI Command Line Interface Quickstart.

- 4. Set up an API key as part of the OCI CLI configuration and select default values.
 - a. Set up the API key access for the IAM user.
 - **b.** Retrieve the db-token.

For example:

• Retrieving a db-token with an API-key using the Oracle Cloud Infrastructure (OCI) command-line interface:

oci iam db-token get

• Retrieving a db-token with a security (or session) token:

oci iam db-token get --auth security token

If the security token has expired, a window will appear so the user can log in to OCI again. This generates the security token for the user. OCI CLI will use this refreshed token to get the db-token.



• Retrieving a db-token with a delegation token: When you log in to the cloud shell, the delegation token is automatically generated and placed in the /etc directory. To get this token, execute the following command in the cloud shell:

oci iam db-token get

Retrieving an instance token by using the OCI command-line interface:

oci iam db-token get --auth instance principal

See Required Keys and OCIDs for more information.

5. Ensure that you are using the latest release updates for the Oracle Database client release 19c.

This configuration only works with the Oracle Database client release 19c (Oracle Database release 21c offers limited IAM token features).

- 6. Ensure that TLS is configured for the database connection.
 - a. Confirm that DN matching is enabled by looking for SSL_SERVER_DN_MATCH=ON in the sqlnet.ora file.
 - **b.** Configure the database client to use the IAM token by adding TOKEN AUTH=OCI TOKEN to the sqlnet.ora file.

Because you will be using the default locations for the database token file, you do not need to include the token location.

The TOKEN_AUTH and TOKEN_LOCATION values in the tnsnames.ora connect strings take precedence over the sqlnet.ora settings for that connection.

For example, for the connect string, assuming that the token is in the default location (~/.oci/db-token for Linux):

After the connect string is updated with the <code>TOKEN_AUTH</code> parameter, the IAM user can log in to the Oracle Database instance by running the following command to start SQL*Plus. You can include the connect descriptor itself or use the name of the descriptor from the <code>tnsnames.ora</code> file.

```
connect /@exampledb high
```



(or)

Ensure that the database client is configured to get a db-token by setting TOKEN_AUTH=OCI_IAM in either sqlnet.ora, tnsnames.ora, or in the connect string. The database client gets the db-token and signs it using the private key and then sends the token to the Oracle Database. If an IAM user name and IAM database password are specified instead of slash *I*, then the database client will connect using the password instead of using the db-token.

Use Instance Principal to Access Oracle Database with Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication

After the DBA user enables Oracle Cloud Infrastructure IAM on Oracle Database, an application can access the database through an Oracle Cloud Infrastructure IAM database token using an instance principal.

Database Links in an Oracle DBaaS-to-IAM Integration

The use of database links when accessing the Oracle DBaaS database using IAM credentials is supported.

The method of configuring database links for Oracle DBaaS connections to IAM depends on the Oracle DBaaS platform. Review the topic below that corresponds to your Oracle DBaaS platform and then click on the associated link for more information.

- Oracle Autonomous Database on Shared Exadata Infrastructure: You can use fixed user database links in which a database user is used for the fixed database link. The database user for creating the database link can only use password authentication with the database link. The IAM user can authenticate to the source database using either password or token access. You cannot configure IAM users as fixed database links, nor can you use connected or current user database links. See, Use Database Links with Autonomous Database.
- Oracle Autonomous Database on Dedicated Exadata Infrastructure and all non-Autonomous Database DBaaS platforms: You can use connected user and fixed user database links, but not current user database links. For connected user database links, an IAM user must be provisioned to both the source and target link databases. You can use a database password verifier or an IAM database token to connect and use connected user database links. For a fixed user database link, a user can connect to the target database using a target database user with password authentication. In addition, an IAM user can connect to the first PDB by using an IAM user name and password or an IAM token. See, Use Identity and Access Management (IAM) Authentication with Autonomous Database.



Related Topics

- Use Database Links with Autonomous Database
- Use Identity and Access Management (IAM) Authentication with Autonomous
 Database

Configuring Authorization for IAM Users and Oracle Cloud Infrastructure Applications

An Oracle DBaaS database administrator can map IAM users and Oracle Cloud Infrastructure (OCI) applications to the Oracle Database global schemas and global roles.

About Configuring Authorization for IAM Users and Oracle Cloud Infrastructure
 Applications

You create the mappings for IAM users and Oracle Cloud Infrastructure (OCI) applications to database users (schemas) in the Oracle DBaaS.

- Mapping an IAM Group to a Shared Oracle Database Global User Oracle Database global users that are mapped to IAM groups and IAM dynamic groups give IAM users and OCI applications a schema when they log in along with the privileges and roles granted to that schema.
- Mapping an IAM Group to an Oracle Database Global Role Oracle Database global roles that are mapped to IAM groups and dynamic groups give member users and applications additional privileges and roles above what they have been granted through their login schemas.
- Exclusively Mapping an IAM User to an Oracle Database Global User You can map an IAM user exclusively to an Oracle Database global user.
- Altering or Migrating an IAM User Mapping Definition You can update an IAM user to a database global user mapping by using the ALTER USER statement.
- Mapping Instance and Resource Principals Instance principals and resource principals can be used by applications to retrieve database tokens to establish a connection to an Oracle DBaaS instance.
- Verifying the IAM User Logon Information After you configure and authorize an IAM user for the Oracle DBaaS instance, you can verify the user logon information by executing a set of SQL queries on the Oracle database side.

About Configuring Authorization for IAM Users and Oracle Cloud Infrastructure Applications

You create the mappings for IAM users and Oracle Cloud Infrastructure (OCI) applications to database users (schemas) in the Oracle DBaaS.

There is a difference with authorization between IAM database password authentication and using IAM token based authentication. IAM database password verifier authorization is only based on mappings of database schemas and global roles to IAM users and group. With IAM token based authentication, IAM policies are an additional authorization for IAM users to access their tenancy databases. An IAM user must be authorized through an IAM policy **and** be authorized through a mapping to a database global schema (exclusive or shared).



For both token and password verifier database access, you create the mappings for IAM users and OCI applications to the Oracle DBaaS 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.

When the IAM user accesses the Oracle DBaaS instance with a token, the database will perform an authorization check against IAM policies to ensure the user is allowed to access the database. If the IAM user is allowed to access the database by IAM policy, then the database will guery IAM for the user groups. When using password verifier authentication, the database will query IAM for user groups once the IAM user successfully completes authentication. The database gueries the IAM endpoint to find the groups of which the user is a member. If your deployment is using shared schemas, then one of the IAM groups will map to a shared database schema and the IAM user will be assigned to that database schema. The IAM user will have the roles and privileges that are granted to the database schema. Because multiple IAM users can be assigned to the same shared database schema, only the minimal set of roles and privileges should be granted to the shared schema. In some cases, no privileges and roles should be granted to the shared schema. Users will be assigned the appropriate set of roles and schemas through database global roles. Global roles are mapped to IAM groups. This way, different users can have different roles and privileges even if they are mapped to the same database shared schema. A newly hired user will be assigned to an IAM group mapped to a shared schema and then to one or more additional groups mapped to global roles to gain the additional roles and privileges required to complete their tasks. The combination of shared schemas and global roles allows for centralized authorization management with minimal changes to the database operationally. The database must be initially provisioned with the set of shared schemas and global roles mapped to the appropriate IAM groups, but then user authorization management can happen within IAM.

Ensure that the IAM user is only mapped to one schema, either through exclusive mapping to a database schema or as a member of one IAM group that is mapped to a shared database schema. If more than one schema is mapped for an IAM user, then the database will take exclusive mapping as precedence over any group mapping to a shared schema. If more than one group is mapped for a user, then the database will select the oldest mapping.

When using global roles to grant privileges and roles to the user, remember that the maximum number of enabled roles in a session is 150.

If you drop and recreate IAM users and groups using the same names, then the mappings from the database to IAM using the same names will continue to work. However, recreating an IAM user will require the IAM user to do one or more of the following: create the IAM database password, re-upload the API public key, update the OCI configuration file, and then re-examine the IAM policy for database authentication and authorization with IAM. If the IAM policy specifies a group that can use or manage the database-connections and autonomous-database-family resource types, then the user will need to be added to that group to allow IAM authentication and authorization.

Accessing the database with tokens requires the user to be authorized by IAM policy and by database mapping. Accessing the database with the IAM database password verifier requires authorization through database mapping. If no database schema mapping exists for the IAM user, the IAM user is prevented from accessing the database even if they have a valid token or password.

IAM users get their authorizations to perform various tasks based on the roles that they have been granted. The following scenarios are possible:

• IAM group mapped to a shared Oracle Database global user: With the shared database global user account, an IAM user is assigned to a shared database schema (user) through the mapping of an IAM group to the shared schema. The IAM users that



are members of the group can connect to the database through this shared schema. Use of shared schemas allows for centralized management of user authorization in IAM.

- IAM group mapped to an Oracle Database global role: The privileges that have been granted to the shared Oracle Database global role become available to the users who have added to the IAM group.
- Local IAM user exclusively mapped to an Oracle Database global user: With an exclusive global user mapping, a dedicated database user is exclusively mapped to a local IAM user. Not as common as the shared database schema, this user is created for when the user requires their own schema objects. Oracle recommends that you grant database privileges to these users through global roles, which facilitates authorization management. These users can also have direct privilege and role grants to their exclusive schema.
 In IAM with Identity Domains, users and groups are supported in the default domain as well as custom non-default domains. When you specify users and groups in the default domain, then no domain prefix is required. When you specify users and groups in a non-default domain, then the domain must be prefixed.

Mapping an IAM Group to a Shared Oracle Database Global User

Oracle Database global users that are mapped to IAM groups and IAM dynamic groups give IAM users and OCI applications a schema when they log in along with the privileges and roles granted to that schema.

- 1. Log in to the Oracle DBaaS instance as a user who has the CREATE USER or ALTER USER system privilege.
- 2. Run the CREATE USER or ALTER USER statement with the IDENTIFIED GLOBALLY AS clause specifying the IAM group name (which can be a dynamic group).

For example, to create a new database global user account named shared_sales_schema and map it to an existing IAM group named WidgetSalesGroup:

```
CREATE USER shared_sales_schema IDENTIFIED GLOBALLY AS
'IAM GROUP NAME=WidgetSalesGroup';
```

The following example shows how to accomplish this for a non-default domain:

CREATE USER shared_sales_schema IDENTIFIED GLOBALLY AS 'IAM GROUP NAME=sales domain/WidgetSalesGroup';

Mapping an IAM Group to an Oracle Database Global Role

Oracle Database global roles that are mapped to IAM groups and dynamic groups give member users and applications additional privileges and roles above what they have been granted through their login schemas.

Global roles cannot be granted to a database schema (user), they can only be mapped to a group and be assigned to an IAM user when accessing the database.

1. Log in to the Oracle DBaaS instance as a user who has been granted the CREATE ROLE or ALTER ROLE system privilege



2. Run the CREATE ROLE or ALTER ROLE statement with the IDENTIFIED GLOBALLY AS clause specifying the name of the IAM group (which can be a dynamic group).

For example, to create a new database global role named widget_mgr_role and map it to an existing IAM group named WidgetManagerGroup, using the default domain:

```
CREATE ROLE widget_mgr_role IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=WidgetManagerGroup';
```

The following example shows how to create the role by specifying a non-default domain, sales domain:

```
CREATE ROLE widget_sales_role IDENTIFIED GLOBALLY AS
'IAM_GROUP_NAME=sales_domain/WidgetManagerGroup';
```

All members of the WidgetManagerGroup in the sales_domain domain will be authorized with the database global role widget sales role when they log in to the database.

Exclusively Mapping an IAM User to an Oracle Database Global User

You can map an IAM user exclusively to an Oracle Database global user.

- 1. Log in to the Oracle DBaaS instance as a user who has been granted the CREATE USER or ALTER USER system privilege.
- 2. Run the CREATE USER or ALTER USER statement with the IDENTIFIED GLOBALLY AS clause specifying the IAM database user name.

By default, the IAM database user name is the same as the IAM user name, including the domain name. You can also create a unique IAM database user name for ease of authentication to the database. In your OCI IAM user profile, you can create a unique IAM database user name for ease of authentication to the database. This can be set when you create and manage your IAM database password in your IAM profile. Adding or changing the IAM database user name will invalidate the IAM user to schema mapping, so the database schema will need to be remapped to the new IAM database user name.

For example, to create a new database global user named <code>peter_fitch</code> and map this user to an existing IAM user named with an IAM database user name of <code>peterfitch</code>, using the default domain:

```
CREATE USER peter_fitch IDENTIFIED GLOBALLY AS 'IAM PRINCIPAL NAME=peterfitch';
```

The following example shows how to create the user by specifying a non-default domain, sales domain:

```
CREATE USER peter_fitch2 IDENTIFIED GLOBALLY AS 'IAM PRINCIPAL NAME=sales domain/peterfitch';
```



Altering or Migrating an IAM User Mapping Definition

You can update an IAM user to a database global user mapping by using the ALTER USER statement.

You can update database schemas that were mapped to an IAM user, and whose accounts were created using any of the CREATE USER statement clauses: IDENTIFIED BY password, IDENTIFIED EXTERNALLY, or IDENTIFIED GLOBALLY. This is useful when migrating existing schemas to using IAM. If you delete and recreate an IAM user or an IAM group using the exact same name as the previous IAM user or group, then the existing mapping from the database that uses that IAM user or IAM group name will continue to work.

- 1. Log in to the Oracle DBaas instance as a user who has been granted the ALTER USER system privilege.
- 2. Run the ALTER USER statement with the IDENTIFIED GLOBALLY AS clause.

For example, suppose you want to change the existing schema shared_sales_schema to a different IAM group:

```
ALTER USER shared_sales_schema IDENTIFIED GLOBALLY AS 'IAM GROUP NAME=BiggerWidgetSalesGroup';
```

The following example shows how to modify the schema by specifying a nondefault domain, sales domain:

ALTER USER shared_sales_schema IDENTIFIED GLOBALLY AS 'IAM GROUP NAME=sales domain/BiggerWidgetSalesGroup';

Mapping Instance and Resource Principals

Instance principals and resource principals can be used by applications to retrieve database tokens to establish a connection to an Oracle DBaaS instance.

Only dynamic groups can be mapped when you use instance and resource principals. You cannot exclusively map instance and resource principals; you only can map them through a shared mapping and putting the instance or resource instance in an IAM dynamic group.

Related Topics

- Managing Dynamic Groups
- Calling Services from an Instance
- Accessing Other Oracle Cloud Infrastructure Resources from Running Functions

Verifying the IAM User Logon Information

After you configure and authorize an IAM user for the Oracle DBaaS instance, you can verify the user logon information by executing a set of SQL queries on the Oracle database side.

1. Log in to the Oracle DBaaS instance as an IAM user that you have just configured and authorized.



For example, to log in to the database instance inst1 as the database global user peterfitch, who is using the default domain in IAM:

sqlplus /nolog CONNECT "peterfitch"@inst1 Enter password: *password*

This example shows how to log in if user peterfitch is in a non-default domain, sales domain:

sqlplus /nolog CONNECT "sales_domain/peterfitch"@inst1 Enter password: password

2. Verify the mapped global user.

The mapped global user is the database user account that has the IAM user authorization. User <code>PETER_FITCH_SCHEMA</code> is considered a global user with exclusive mapping for the IAM user <code>peterfitch</code>, while user <code>WIDGET_SALES</code> is considered a global user with shared mapping for IAM group <code>widget_sales_group</code> of which <code>peterfitch</code> is a member.

SHOW USER;

Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

USER is "PETER_FITCH_SCHEMA"

Or

USER is "WIDGET SALES"

3. Find the roles that have been granted to the centrally managed user.

SELECT ROLE FROM SESSION_ROLES ORDER BY ROLE;

Output similar to the following appears:

```
ROLE
-----WIDGET_SALES_ROLE
...
```

- 4. Run the following queries to check the SYS_CONTEXT namespace values for the current schema being used in this database session, current user name, session user name, authentication method, authenticated identity, enterprise identity, identification type, and server type.
 - Verify the current schema that is being used in this database session. A database schema is an object container that identifies the objects it contains. The current schema is the default container for objects name resolution in this database session.

SELECT SYS CONTEXT ('USERENV', 'CURRENT SCHEMA') FROM DUAL;

Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

Or



 Verify the current user. In this case, the current user is the same as the current schema.

SELECT SYS CONTEXT ('USERENV', 'CURRENT USER') FROM DUAL;

Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

```
SYS_CONTEXT('USERENV','CURRENT_USER')
```

PETER_FITCH_SCHEMA

Or

Verify the session user.

SELECT SYS CONTEXT ('USERENV', 'SESSION USER') FROM DUAL;

Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

_ _ _

Or

• Verify the authentication method.

SELECT SYS CONTEXT ('USERENV', 'AUTHENTICATION METHOD') FROM DUAL;

Output similar to the following appears:

If the user is authenticating with a token, then the output is TOKEN GLOBAL.

• Verify the authenticated identity for the enterprise user. The IAM authenticated user identity is captured and audited when this user logs on to the database.

SELECT SYS_CONTEXT('USERENV', 'AUTHENTICATED_IDENTITY') FROM DUAL;

Output similar to the following appears:

SYS_CONTEXT('USERENV', 'AUTHENTICATED_IDENTITY')

```
sales domain/peterfitch
```



• If a user nickname has been set for the enterprise user, then verify this nickname.

SELECT SYS CONTEXT ('USERENV', 'USER NICKNAME') FROM DUAL;

Output similar to the following appears:

```
SYS_CONTEXT('USERENV','USER_NICKNAME')
_____
pfitch
```

• Verify the centrally managed user's enterprise identity.

SELECT SYS CONTEXT ('USERENV', 'ENTERPRISE IDENTITY') FROM DUAL;

Enterprise Identity will show the OCI Identity (OCID) of the IAM user or OCI application. Output similar to the following appears:

Verify the identification type.

SELECT SYS CONTEXT ('USERENV', 'IDENTIFICATION TYPE') FROM DUAL

Output similar to the following appears, depending on if it is an exclusive mapping or a shared mapping:

Or

```
SYS_CONTEXT('USERENV','IDENTIFICATION_TYPE')
______GLOBAL SHARED
```

• Verify the server type.

SELECT SYS CONTEXT ('USERENV', 'LDAP SERVER TYPE') FROM DUAL;

Output similar to the following appears. In this case, the LDAP server type is IAM.

Configure Proxy Authentication

Proxy authentication allows an Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) user to proxy to a database schema for tasks such as application maintenance.

About Configuring Proxy Authentication
 Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) users can
 connect to Oracle Database by using proxy authentication.



- Configure Proxy Authentication for the Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) User
 To configure proxy authentication for an IAM user, the IAM user must already have a mapping to a global schema (exclusive or shared mapping). A separate database schema for the IAM user to proxy to must also be available.
- Validate the Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) User Proxy Authentication
 You can validate the IAM user proxy configuration for both password and token authentication methods.

About Configuring Proxy Authentication

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) users can connect to Oracle Database by using proxy authentication.

Proxy authentication is typically used to authenticate the real user and then authorize them to use a database schema with the schema privileges and roles in order to manage an application. Alternatives such as sharing the application schema password are considered insecure and unable to audit which actual user performed an action.

A use case can be in an environment in which a named IAM user who is an application database administrator can authenticate by using their credentials and then proxy to a database schema user (for example, *hrapp*). This authentication enables the IAM administrator to use the *hrapp* privileges and roles as user *hrapp* in order to perform application maintenance, yet still use their IAM credentials for authentication. An application database administrator can sign in to the database and then proxy to an application schema to manage this schema.

You can configure proxy authentication for both password authentication and token authentication methods.

Configure Proxy Authentication for the Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) User

To configure proxy authentication for an IAM user, the IAM user must already have a mapping to a global schema (exclusive or shared mapping). A separate database schema for the IAM user to proxy to must also be available.

After you ensure that you have this type of user, alter the database user to allow the IAM user to proxy to it.

- 1. Log in to the Oracle Database instance as a user who has the ALTER USER system privileges.
- 2. Grant permission for the IAM user to proxy to the local database user account.

An IAM user cannot be referenced in the command so the proxy must be created between the database global user (mapped to the IAM user) and the target database user. In the following example, hrapp is the database schema to proxy to, and *peterfitch_schema* is the database global user exclusively mapped to user *peterfitch*.

ALTER USER hrapp GRANT CONNECT THROUGH peterfitch schema;

At this stage, the IAM user can log in to the database instance using the proxy.



For example, to connect using a password verifier:

```
CONNECT peterfitch[hrapp]@connect_string
Enter password: password
```

3. To connect using a token:

```
CONNECT [hrapp]/@connect string
```

Validate the Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) User Proxy Authentication

You can validate the IAM user proxy configuration for both password and token authentication methods.

- 1. Log in to the Oracle Database instance as a user who has the CREATE USER and ALTER USER system privileges.
- 2. Connect at the IAM user and execute the SHOW USER and SELECT SYS_CONTEXT commands.

For example, suppose you want to check the proxy authentication of the IAM user peterfitch when they proxy to database user hrapp. You will need to connect to the database using the different types of authentication methods shown here, but the output of the commands that you execute will be the same for all types.

• For password authentication:

```
CONNECT peterfitch[hrapp]/password\!@connect_string

SHOW USER;

--The output should be USER is "HRAPP"

SELECT SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD') FROM DUAL;

--The output should be "PASSWORD_GLOBAL"

SELECT SYS_CONTEXT('USERENV','PROXY_USER') FROM DUAL;

--The output should be "PETERFITCH_SCHEMA"

SELECT SYS_CONTEXT('USERENV','CURRENT_USER') FROM DUAL;

--The output should be "HRAPP"
```

• For token authentication:

```
CONNECT [hrapp]/@connect_string

SHOW USER;

--The output should be USER is "HRAPP"

SELECT SYS_CONTEXT('USERENV','AUTHENTICATION_METHOD') FROM DUAL;

--The output should be "TOKEN_GLOBAL"

SELECT SYS_CONTEXT('USERENV','PROXY_USER') FROM DUAL;

--The output should be "PETERFITCH_SCHEMA"

SELECT SYS_CONTEXT('USERENV','CURRENT_USER') FROM DUAL;

--The output should be "HRAPP"
```



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
- Oracle Database Backup Methods in Exadata Cloud
 Oracle Exadata 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 bkup_api
 In addition to the console-based automated backup option, there is a command line backup utility, bkup_api, which can allow for further customization. If configuring backups using bkup_api instead of the console, then do not enable backups for your database in the console.
- Creating an On-Demand Backup by Using the bkup_api Utility You can use the bkup_api utility 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.

Backup Destinations

 About Managing Backup Destinations for Exadata Database Service on Cloud@Customer
 For backups, you can either use the Exadata Database Service on Cloud@Customer backup facility or you can configure a backup location on

Cloud@Customer backup facility, or you can configure a backup location on a location you manage.

Prerequisites for Backup Destinations for 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 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 Exadata Cloud@Customer Backup Destinations Review the list of API calls to manage your Exadata Database Service on Cloud@Customer backup destinations.

About Managing Backup Destinations for Exadata Database Service on Cloud@Customer

For backups, you can either use the Exadata Database Service on Cloud@Customer backup facility, or you can configure a backup location on a location you manage.

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 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 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 Exadata Database Service on Cloud@Customer Systems

Learn to manage Oracle Database homes on 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 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.

- 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 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:
 - 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.
 - Autonomous Database on Cloud@Customer:

Note:

Oracle Autonomous 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.



Using the Console for Backup Destinations for 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 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*.

ORACLE

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.
 - **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 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

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 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 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 Exadata Cloud@Customer Backup Destinations

Review the list of API calls to manage your 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 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 Cloud@Customer offers two approaches to configure and take backups as a recommended solution: Oracle Managed Backup and User Configured Backup.



Oracle Managed Backup

Database backups are managed entirely by the Oracle Exadata 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 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
- 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

About Managing Backup Destinations for Exadata Database Service on Cloud@Customer

For backups, you can either use the 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 Exadata Database Service on Cloud@Customer
 To configure backup destinations on a Zero Data Less Dessuant Appliant

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 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 Exadata Cloud@Customer Backup Destinations Review the list of API calls to manage your Exadata Database Service on Cloud@Customer backup destinations.



About Managing Backup Destinations for Exadata Database Service on Cloud@Customer

For backups, you can either use the Exadata Database Service on Cloud@Customer backup facility, or you can configure a backup location on a location you manage.

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 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 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 Exadata Database Service on Cloud@Customer Systems
 Learn to manage Oracle Database homes on Exadata Database Service on

Learn to manage Oracle Database homes on 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 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.

• 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 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:
 - 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.
 - Autonomous Database on Cloud@Customer:

Note:

Oracle Autonomous 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.

Using the Console for Backup Destinations for 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 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.

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

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 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

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 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 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 Exadata Cloud@Customer Backup Destinations

Review the list of API calls to manage your 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 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.



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.
 - For backups to Local Exadata Storage:

The retention period option for Local Exadata Storage is 7 or 14 days. Regardless of the retention window selected for backups to Local Exadata Storage, incremental level 1 backups are always performed after the initial level 0 image copy is taken. Also, the incremental level 1 backups are merged into the level 0 image copy backup when they become older than the retention period.

For example: A 14 day for Local retention window includes one "merged" level 0, 14 incremental level 1's plus archivelogs for the 14 days.

- 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 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.
- 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).

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 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 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.
- 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 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 Edit Backup Settings.
- 8. Your current backup configuration determines the changes that you can make in the **Backup Settings** dialog, as follows:
 - If automatic backups are not configured (Backup Destination Type is set to None), then you can use the following settings to define the backup configuration for the database:
 - Backup Destination Type: From the list, choose an option.



- * None Select if you do not define a backup configuration for the database.
- * **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 .
- * 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 (other than **None**), 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.

This option is only enabled when you select a **Backup Destination Type** other than **None**. You can change this setting later.

 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 Changes.

Related Topics

- Backup Destinations
- 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 Exadata Database Service on Cloud@Customer Console.

Use the following procedure to restore a database:

- 1. Open the navigation menu. Under Oracle 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.

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 <code>dbaascli</code> 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
- 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

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.

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.
- 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 5-19 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.	Yes
	Default: low.	
	NONE disables RMAN compression.	
bkup_set_section_size	Enables the use of the RMAN multisection backup feature. Valid values are yes and no.	Yes
bkup_section_size	RMAN section size that is used for automatic backups. The default value is 64G.	Yes
	Applicable only when bkup_set_section_size is set to 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.	Yes
	Default: 30 for Exadata Cloud@Customer.	



Parameter	Description	Compatible with Oracle Managed Backup
bkup_10_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.

Table 5-19 (Cont.) General Configuration Parameters (valid for all backupdestinations except Local Storage (FRA))

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
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

Table 5-20 Configuration Parameters for Object Storage Service (OSS) Destination



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 5-21Configuration 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
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 5-22Configuration 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



Parameter	Description	Compatible with Oracle Managed Backup
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_windo w	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.	No
	Default: 30.	

Table 5-22(Cont.) Configuration Parameters for Network File System (NFS)Destination

Note:

Currently, Network File System (NFS) destination is supported only on Exadata Cloud@Customer.

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

Table 5-23	Configuration	Darameters f	orlocal	Storage		Destination
Table 5-25	Configuration	Parameters n	UI LUCAI	Slurage	FRA	Destination

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)

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 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 bkup_api

In addition to the console-based automated backup option, there is a command line backup utility, bkup_api, which can allow for further customization. If configuring backups using



 $\tt bkup_api$ instead of the console, then do not enable backups for your database in the console.

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.

- 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.
- Customizing Which System Files Are Backed Up By default, backups via the console or bkup_api backup certain system files in addition to the database files themselves.
- Customizing Which Database Configuration Files Are Backed Up By default, backups through the console or bkup_api backup certain database configuration files in addition to the database files themselves.

Related Topics

• User Configured Backup It is the user's responsibility to configure and execute backup operations using dbaascli according to their preferences.

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.

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

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

Where:

filename is an optional parameter used to specify a name for the file that is generated

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

4. Edit the parameter values in the generated file to change 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 5-24Backup Configuration Parameters - Cron Parameters Specific to
bkup_api

Parameter	Description	Compatible with Console Automatic Backups*
bkup_cron_entry	Enables the automatic backup configuration.	No
	Valid values are yes and no.	
bkup_cfg_files	For ZDLRA, backups of the system/ database configuration will not occur.	No



Table 5-25Backup Configuration Parameters - General RMANConfiguration Parameters (valid for all backup destinations except LocalStorage (FRA))

Parameter	Description	Compatible with Console Automatic Backups*
bkup_rman_compression	Level of compression applied to automatic backups.	Yes
	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.	
bkup_set_section_size	Enables the use of the RMAN multisection backup feature.	Yes
	Valid values are yes and no.	
bkup_section_size	RMAN section size that is used for automatic backups.	Yes
	Default value is 64G.	
	Applicable only when bkup_set_section_size is set to yes.	
bkup_channels_node	Number of RMAN channels per node used for automatic backups.	Yes
	Valid values are between 1 and 32.	
	Default value is 4.	
bkup_daily_time	Start time of the automatic daily backup expressed in 24-hour time as hh:mm.	Yes
bkup_archlog_frequency	Interval in minutes between automatic backups of archived database log files.	Yes
	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.	



Parameter	Description	Compatible with Console Automatic Backups*
bkup_disk	Enables backups to local Exadata storage. Valid values are yes and no.	No
bkup_disk_recovery_windo w	Retention period for backups on local Exadata storage, expressed as a number of days up to 14.	No
	Applicable only when bkup_disk is set to yes.	
	Default value is 7.	

Table 5-26 Backup Configuration Parameters - Local Storage (FRA) Parameters

Table 5-27Backup Configuration Parameters - Network File System (NFS)Parameters

Parameter	Description	Compatible with Console Automatic Backups*
bkup_nfs	Enables backups to NFS mounted directory.	No
	Valid values are yes and no.	
bkup_nfs_loc	The NFS mounted location, the directory provided must be a mountpoint and available on all nodes.	No
bkup_nfs_recovery_window	Retention period for backups on NFS storage, expressed as a number of days up to 90.	No
	Applicable only when bkup_nfs is set to yes.	
	Default value is 30.	
bkup_oss_10_day	This parameter control the Level 0 day of the week for both OSS and NFS.	Yes
	Day of the week when a level 0 backup is taken and stored on cloud storage or for NFS.	
	Valid values are mon, tue, wed, thu, fri, sat, sun.	
	Applicable only when bkup_nfs is set to yes.	
	Default value is sun.	

Parameter	Description	Compatible with Console Automatic Backups*
bkup_oss	Enables backups to cloud storage.	No
	Valid values are yes and no.	
bkup_oss_recovery_wind ow	Retention period for backups to cloud storage, expressed as a number of days up to 90.	No
	Applicable only when bkup_oss is set to yes.	
	Default value is 30.	
bkup_oss_url	Location of the storage container that is used for backup to cloud storage.	No
	Applicable only when bkup_oss is set to yes.	
bkup_oss_user	User name of the Oracle Cloud user having write privileges on the cloud storage container specified in bkup_oss_url.	No
	Applicable only when bkup_oss is set to yes.	
bkup_oss_passwd	Password of the Oracle Cloud user having write privileges on the cloud storage container specified in bkup oss url.	No
	Applicable only when bkup_oss is set to yes.	
bkup_oss_10_day	This parameters control the Level 0 day of the week for both OSS and NFS.	Yes
	Day of the week when a level 0 backup is taken and stored on cloud storage.	
	Valid values are mon, tue, wed, thu, fri, sat, sun.	
	Applicable only when bkup_oss is set to yes.	
	Default value is sun.	

Table 5-28Backup Configuration Parameters - Object Storage Service(OSS) Parameters

Note:

Zero Data Loss Recovery Appliance (ZDLRA) parameters are only valid for Exadara Cloud@Customer installations.

Parameter	Description	Compatible with Console Automatic Backups*
bkup_zdlra	Enables backups to a Recovery Appliance.	No
	Valid values are yes and no.	
bkup_zdlra_url	Location of the Recovery Appliance that is being used for backups.	No
	Applicable only when bkup_zdlra is set to yes.	
bkup_zdlra_user	The virtual private catalog (VPC) user name for the Recovery Appliance specified in bkup_zdlra_url.	No
	Applicable only when bkup_zdlra is set to yes.	
bkup_zdlra_passwd	Password of the Recovery Appliance user specified in bkup_zdlra_url.	No
	Applicable only when bkup_zdlra is set to yes.	

Table 5-29Backup Configuration Parameters - Zero Data Loss RecoveryAppliance (ZDLRA) Parameters

Table 5-30Backup Configuration Parameters - RMAN Catalog SupportParameters

Parameter	Description	Compatible with Console Automatic Backups*
bkup_use_rcat	Enables the use of an existing RMAN recovery catalog. Valid values are yes and no.	Yes (only for NFS and OSS backups)
bkup_rcat_user	Recovery catalog user name. Applicable only when bkup_use_rcat is set to yes.	Yes (only for NFS and OSS backups)
bkup_rcat_passwd	Password for recovery catalog user specified in bkup_rcat_user.	Yes (only for NFS and OSS backups)
	Applicable only when bkup_use_rcat is set to yes.	
bkup_rcat_conn	Connection string for the RMAN recovery catalog.	Yes (only for NFS and OSS backups)
	Applicable only when bkup_use_rcat is set to yes.	

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 bkup_api set config command to update the backup settings using the file containing your updated backup settings:

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

Where:

filename is used to specify the name of the file that contains the updated backup settings

dbname is the database name for the database that you are acting on

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

/var/opt/oracle/bkup api/bkup api configure status

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 Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Customizing Which System Files Are Backed Up

By default, backups via the console or bkup_api backup certain system files in addition to the database files themselves.

If you need different system files backed up, then use these steps to change which system files get backed up. It is safe to alter parameters pertaining to system file backups in conjunction with the use of console-based automated backups.

If your backup configuration includes bkup_cfg_files=yes, then each backup includes system configuration files and directories specified in the oscfg.spec file.

To change which system files and directories are backed up:

1. Connect to a virtual machine as opc user.

For detailed instructions, see Connecting to a Virtual Machine with SSH.

2. Switch to oracle user.



3. Edit the contents of the oscfg.spec file.

This file is located under /var/opt/oracle/dbaas_acfs/bkup/dbname, where dbname is the name of the database that is associated with the backup configuration.

Following is an example of the default contents of the oscfg.spec file:

```
## OS Configuration Files
#
# Doc Spec
oscfg.spec
# Directories
/etc/rc.d
/home/oracle/bkup
# Single files
/home/oracle/.bashrc
/etc/crontab
/etc/sysctl.conf
/etc/passwd
/etc/group
/etc/oraInst.loc
/etc/oratab
/etc/fstab
```

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Customizing Which Database Configuration Files Are Backed Up

By default, backups through the console or bkup_api backup certain database configuration files in addition to the database files themselves.

If you need different database configuration files backed up, then use these steps to change which database configuration files get backed up. It is safe to alter parameters pertaining to database configuration file backups in conjunction with the use of console-based automated backups.

If your backup configuration includes bkup_cfg_files=yes, then each backup includes
database configuration files and directories specified in the dbcfg.spec file.

To change which database configuration files are backed up:

1. Connect to a virtual machine as opc user.

For detailed instructions, see Connecting to a Virtual Machine with SSH.

- 2. Switch to oracle user.
- 3. Edit the contents of the dbcfg.spec file.

This file is located under /var/opt/oracle/dbaas_acfs/bkup/dbname, where dbname is the name of the database that is associated with the backup configuration.



Following is an example of the contents of the dbcfg.spec file:

```
### Oracle Home configuration files.
#
# Doc Spec
dbcfg.spec
# DB id
dbid
# Directories
/u02/app/oracle/product/dbversion/dbhome n/admin/dbname/xdb wallet
/u02/app/oracle/admin/dbname/xdb wallet
/u02/app/oracle/admin/dbname/db wallet
# Note: tde wallet must be backed up in a different location than
DATA bkup.
/u02/app/oracle/admin/dbname/tde wallet
/u02/app/oracle/admin/dbname/cat wallet
#/u01/app/oraInventory
# Single files
/var/opt/oracle/dbaas acfs/dbname/opc/opcdbname.ora
/u02/app/oracle/product/dbversion/dbhome n/dbs/opcdbname.ora
/u02/app/oracle/product/dbversion/dbhome n/dbs/orapwinstancename
/u02/app/oracle/product/dbversion/dbhome n/network/admin/
listener.ora
/u02/app/oracle/product/dbversion/dbhome n/network/admin/sqlnet.ora
/u02/app/oracle/product/dbversion/dbhome n/network/admin/
tnsnames.ora
/u02/app/oracle/product/dbversion/dbhome n/rdbms/lib/env rdbms.mk
/u02/app/oracle/product/dbversion/dbhome n/rdbms/lib/ins rdbms.mk
# Creg
/var/opt/oracle/creg/instancename.ini
```

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Creating an On-Demand Backup by Using the bkup_api Utility

You can use the bkup_api utility 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 bkup api command:
 - To create a backup that follows the current retention policy, use the following bkup_api command:

/var/opt/oracle/bkup_api/bkup_api bkup_start --dbname=dbname

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 of the complete database that persists until you delete it, use the following bkup_api command:

/var/opt/oracle/bkup_api/bkup_api bkup_start --keep --dbname=dbname

By default, the long-term backup is given a timestamp-based tag. To specify a custom backup tag, add the --tag option to the bkup api command.

For example, to create a long-term backup with the tag monthly, use the following command:

/var/opt/oracle/bkup_api/bkup_api bkup_start --keep --tag=monthly -dbname=dbname

• To create an on-demand RMAN level 0 backup, use the following bkup_api command:

/var/opt/oracle/bkup_api/bkup_api bkup_start --level0 -dbname=dbname



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.

• To create an on-demand backup that includes an image copy of the database data files, use the following bkup api command:

/var/opt/oracle/bkup_api/bkup_api bkup_start --datafiles -dbname=dbname

You can use this option to manually perform a full image backup to cloud storage if the scheduled weekly full 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 and local Exadata storage.

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 bkup_api command on the same compute node where the backup is running:

/var/opt/oracle/bkup_api/bkup_api bkup_status --dbname=dbname

- 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 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 Database Service on Cloud@Customer console, API or bkup_api work for a variety of backup and recovery use cases. If you require use cases not supported by the cloud-managed backups, then you can manage database backup and recovery manually, using the Oracle Recovery Manager (RMAN) utility. For information about using RMAN, see the Oracle Database Backup and Recovery User's Guide for Release 19.

Managing backup and recovery, using RMAN, on Exadata Database Service on Cloud@Customer requires taking full ownership of both database and archive log backups, and the cloud-managed backups should no longer be used. Before manual backups are started, the cloud-managed backup functionality should be disabled. This is needed so the cloud backup jobs do not purge archive logs before they are manually backed up and do not conflict with the manual backups.

You can use the bkup_api utility to disable cloud-managed backups, including disabling the automatic archive log purge job, by following this procedure:



Note:

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

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

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

2. Start a root-user command shell:

sudo -s

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

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

Where:

- *filename* is an optional parameter used to specify a name for the file that is generated
- *dbname* is the database name for the database that you want to act on
- 4. Edit the parameter values in the generated file to change the following parameters.

This will remove the backup crontab entries and disable all automatic backups. If the values are set to ${\tt yes},$ then set to no.

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

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

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

Where:

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

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



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

/var/opt/oracle/bkup api/bkup api configure status --dbname=dbname

Where:

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

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

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

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

Where:

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

Note:

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

Note:

Changes made using the bkup_api command are not reflected in the Oracle Exadata Database Service on Cloud@Customer console.

8. Exit the root-user command shell:

exit

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an 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 <code>dbaascli</code>. This is available from <code>dbaastools 21.4.1</code> release.

For example, to enable RTRT for the database *myTestDB*:

dbaascli database backup --dbName myTestDB --configure -enableRTRT

For more information, see dbaascli database backup.

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: bkup_api and Oracle Recovery Manager (RMAN).



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.

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 Database Service on Cloud@Customer 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 bkup_api. If you choose to orchestrate backups manually leveraging RMAN, you should not use either console automated backups, nor should you use bkup_api. You must first completely disable console based automated backups. For more information, see *Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management*.

The bkup_api method offers a middle ground between RMAN and console automated backups in terms of flexibility and simplicity. Use bkup_api if needed functionality is not supported with console automated backups, but when you wish to avoid complexity of using RMAN directly. In certain cases, bkup_api can be used to modify the console automated backup configuration, but this is not generally the case. Generally, bkup_api 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 bkup_api, then you can manually restore that database backup by using the Oracle Recovery Manager (RMAN) utility. For information about using RMAN, see the *Oracle Database Backup and Recovery User's Guide 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 backup_api 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.

Related Topics

Oracle Database Backup and Recovery User's Guide for Release 19

Patch and Update an Exadata Database Service on Cloud@Customer System

Learn to update and patch the Exadata Database Service on Cloud@Customer System

- Perform User Managed Maintenance Updates
- Patching and Updating an 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 Exadata Database Service on Cloud@Customer System Manually

This topic describes the procedures for patching and updating various components in 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".

Perform User Managed Maintenance Updates

Maintaining a secure Exadata 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 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 Exadata Cloud@Customer VM cluster
 nodes in an automated manner from the OCI console and APIs.
- Administering Oracle Clusterware



Patching and Updating an 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 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 Exadata Cloud@Customer VM cluster nodes in an automated manner from the OCI console and APIs.
- Upgrading Oracle Grid Infrastructure on an Exadata Cloud@Customer VM Cluster Learn to upgrade Oracle Grid Infrastructure on an Exadata Cloud@Customer VM cluster using the Oracle Cloud Infrastructure Console or API.
- Upgrading Oracle Databases Learn to upgrade Oracle Database 19c (Long Term Release) using the Console and the API.

Related Topics

• Patching and Updating an Exadata Database Service on Cloud@Customer System Manually

This topic describes the procedures for patching and updating various components in 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 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.
- 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 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.

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 a Grid Infrastructure Patch Operation on a VM
 Cluster

Learn to apply Grid Infrastructure patches on a VM cluster.

- Using the Console to Perform a Patch 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 a Grid Infrastructure Patch Operation on a VM Cluster

Learn to apply Grid Infrastructure patches on a VM cluster.

1. Open the navigation menu. Under Oracle 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. Under Oracle Grid Infrastructure Version, click View Updates.
- 5. Review the scope:
 - VM Cluster: Automatically set to the context from which you have launched this page.
 - **Database Home:** Automatically set to the context from which you have launched this page. If you have not set the context, then select the Database Home first.
- 6. Review the list of available patches for the VM cluster.
- **7.** 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 Patch: Applies the selected patch.
- 8. 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 a Patch Operation on a Database Home

Learn to apply patches on a Database Home.

1. Open the navigation menu. Under Oracle 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. Under Resources, click Database Homes.
- 5. In the list of Database Homes, click the Database Home on which you want to perform a patch operation.
- 6. Under Database Software Version, click View Patches.
- 7. Review the scope:
 - **Database Home:** Automatically set to the context from which you have launched this page.
- 8. Review the list of available patches for the Database Home.

The **Oracle Provided Database Software Images** tab displays generally-available Oracle Database software images that you can use to patch your database. Oracle images that can be used for patching have the update Type of "Patch".

The **Custom Database Software Images** tab allows you to select a 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.

- 9. 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. The Precheck does not cause any availability impact to the cluster, everything remains operational.
 - **Apply:** Applies the selected patch.
- **10.** 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, 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 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. Under Oracle Grid Infrastructure Version, click View Patches.
- 5. 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 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. Under Resources, click Database Homes.

A list of Database Homes is displayed.

- 5. In the list of Database Homes, click the Database Home you are interested in.
- 6. Under Database Software Version, click View Patches.



7. 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 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.

- In the list of VM clusters, click the VM cluster where the database you want to move is located.
- 4. Under Resources, click Database Homes.
- 5. In the list of Database Homes, click the Database Home you are interested in.

A list of databases is displayed.

- 6. In the list of databases, click the database you are interested in.
- 7. Click Move Database.
- 8. Select the target Database Home.
- 9. Click Move Database.

The database will be stopped in the current home and then restarted in the destination home.

10. 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.

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 Exadata 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.



- 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



- * 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 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. In the Version section, to the right of the Updates Available, click View Updates to display the Updates page.
- 5. Review the list of available software updates and locate the 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.



Note:

If the precheck fails, the system displays a message in the **Apply Exadata OS Image Update** dialog that the last precheck has failed. Oracle recommends that you run the precheck again. Click the Actions icon (three dots)at the end of the row listing the OS patch to view the dialog.

- Apply Exadata OS Image Update. This link displays the Apply Exadata Image Update dialog that you use to apply the patch. The dialog shows the name of the database system you are patching, the current version of the database, and the new version of the database after the patch is applied. To start the process, click Apply Exadata OS Image Update.
- **Copy OCID.** This copies the Oracle Cloud ID. This can be used when troubleshooting a patch or to give to Support when contacting them.

Note:

While the patch is running:

- Run Precheck and Apply OS Image Update are not available. When the patch has completed, these actions are available again.
- If the Exadata infrastructure containing this VM cluster is scheduled for maintenance that conflicts with the patching operation, the patch fails and the system displays a message explaining why. After the infrastructure maintenance is complete, run the patch operation again.
- 7. Confirm when prompted.

The patch list displays the status of the operation in the Version section of the database details page. Click **View Updates** to view more details about an individual patch status and to display any updates that are available to run. If no new updates are available, the system displays a message that says **No Updates Available**.

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 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.

Confirm Rollback Operation dialog is displayed.

Click Roll Back to confirm.

- 4. You can also Apply Exadata Image Update from the Updates page.
 - a. In the Version section, to the right of the Updates Available, click View Updates to display the Updates page.
 - b. Click the Actions icon (three dots) and then select the Apply Exadata OS Image Update option.

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 Exadata Cloud@Customer VM Cluster

Learn to upgrade Oracle Grid Infrastructure on an Exadata 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.
- 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 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. Under Version, click the View Updates link beside the Updates Available field.
- 5. Click **View Updates** to view the list of available patches and upgrades.
- 6. Click the Actions icon (three dots) at the end of the row listing the Oracle Grid Infrastructure (GI) upgrade, then click **Run Precheck**.
- 7. In the **Confirm** dialog, confirm you want to upgrade to begin the precheck operation.

Using the Console to Upgrade Oracle Grid Infrastructure of a VM Cluster

- 1. Open the navigation menu. Under Oracle 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. Under Version, click the View Updates link beside the Updates Available field.
- 5. Click **View Updates** to view the list of available patches and upgrades.
- 6. Click the Actions icon (three dots) at the end of the row listing the Oracle Grid Infrastructure (GI) upgrade, then click **Upgrade Grid Infrastructure**.
- 7. In the Upgrade Grid Infrastructure dialog, confirm you want to upgrade the GI by clicking **Upgrade Grid Infrastructure**.

If you haven't run a precheck, you have the option of clicking **Run Precheck** in this dialog to precheck your cloud VM cluster prior to the upgrade.

Using the API to 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
- 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 Oracle Database 19c (Long Term Release) 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 prerequsites to upgrade an Exadata Cloud@Customer Oracle Database instance.

- The Exadata Cloud@Customer system software must use Oracle Linux 7 (OL7). 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 must be version 19c. 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 two most recent versions of Oracle Database 19c 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 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 Oracle Database Software Images
 Learn about Database Software Image resource type and how you can use it to
 create Oracle Databases and Oracle Database Homes and to patch databases.
- 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.
- If your databases use Data Guard, you will need to disable or remove the Data Guard association prior to upgrading.
- An upgrade operation cannot take place while an automatic backup operation is underway. Before upgrading, Oracle recommends disabling automatic backups and performing a manual backup. See *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 bkup_api

In addition to the console-based automated backup option, there is a command line backup utility, <code>bkup_api</code>, which can allow for further customization. If configuring backups using <code>bkup_api</code> instead of the console, then do not enable backups for your database in the console.

• Creating an On-Demand Backup by Using the bkup_api Utility You can use the bkup_api utility 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:

- 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 bkup_api
 In addition to the console-based automated backup option, there is a command line backup utility, bkup_api, which can allow for further customization. If configuring backups using bkup_api instead of the console, then do not enable backups for your database in the console.
- 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 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 Upgrade.
- 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 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 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, under Database Version, click the **View** link that is displayed for databases that have been upgraded.

This link does not appear for databases that have not been updated.

The **Updates History** page is displayed. The table displayed on this page shows precheck and upgrade operations performed on the database.

Using the API to Upgrade 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 Exadata Database Service on Cloud@Customer System Manually

This topic describes the procedures for patching and updating various components in 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 Exadata Database Service on Cloud@Customer Guest VMs.

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 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

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 ${\tt patchmgr}$ 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.
- 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]----+
0.. + .
0.
      o *
                 | E
     . 0 0
                 | . .
        =
| o. S =
                 = .
   +
+ 00
L
  . . + .
      . . .
```

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
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
```

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 <code>-nomodify_at_prereq</code> 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
******
*****
      patchmgr release: 5.180228 (always check MOS 1553103.1 for the
NOTE
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

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
```

ORACLE

**** 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).

- 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-ofedrelease-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 &
```

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 &
```



***** patchmgr release: 5.180228 (always check MOS 1553103.1 for the NOTE 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) :SUCCESS: DONE: Initiate reboot on 2018-02-28 21:48:57 +0000 node(s) 2018-02-28 21:48:57 +0000 :Working: DO: Waiting to ensure node2 is down before reboot. 2018-02-28 21:56:18 +0000 :Working: DO: Initiate prepare steps on node(s). 2018-02-28 21:56:19 +0000 :Working: DO: Check free space and verify SSH equivalence for the root user to node2 2018-02-28 21:57:37 +0000 :SUCCESS: DONE: Check free space and verify SSH equivalence for the root user to node2 2018-02-28 21:57:42 +0000 :SEEMS ALREADY UP TO DATE: node2 2018-02-28 21:57:43 +0000 :SUCCESS: DONE: Initiate update on node(s)

7. After the update operation completes, verify the version of the 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 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

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 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.

Download an Interim Software Update

The patch download path is valid for four days. Download the patch within the specified timeframe.

1. On the Update details page, click **Download**.

The system starts downloading the patch.



- 2. You can also download a patch from the Interim Software Updates page.
 - Click the Actions button (three dots) for the patch you're interested in, and select **Download**.

Note:

You can only download the patches that are in Available state.

Interim Software Updates Lifecycle States:

- Available: Patch has been created successfully and the time-to-live (TTL) has not expired.
- **Creating**: The patch creation process is in progress.
- **Expired**: The lifetime of the patch download link has expired, which means you cannot download it.
- **Failed**: The patch create failed due to some error.
- Terminating: The patch deletion process is in progress.
- **Terminated**: The patch has been deleted.

Delete an Interim Software Update

Be discrete in deleting interim software updates. However, you can delete the interim software updates that have expired to free up space in the Object Store.

- 1. On the Update details page, click Delete.
- 2. In the resulting dialog, enter the name of the patch to confirm and then click **Delete**.
- 3. You can also delete a patch from the Interim Software Updates page.
 - Click the Actions button (three dots) for the patch you're interested in, and select **Delete**.

Move an Interim Software Update Resource to Another Compartment

- 1. On the Update details page, click Move Resource.
- 2. In the resulting dialog, choose a new compartment, and click Move Resource.
- 3. You can also move a patch resource from the Interim Software Updates page.
 - Click the Actions button (three dots) for the patch you're interested in, and select **Move Resource**.

Using the API to Manage Interim Software Updates

ExaDB-C@C and ExaDB-D use the same API to manage interim software updates.



For information about using the API and signing requests, see *REST APIs* and *Security Credentials*. For information about SDKs, see *Software Development Kits and Command Line Interface*.

Use these API operations to manage interim software updates:

- CreateOneoffPatch
- DeleteOneoffPatch
- DownloadOneoffPatch
- UpdateOneoffPatch
- ListOneoffPatches
- GetOneoffPatch
- ChangeOneoffPatchCompartment

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface
- OneoffPatch Reference

Use Oracle Data Guard with Exadata Database Service on Cloud@Customer

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

- About Using Oracle Data Guard with Exadata Database Service on Cloud@Customer This topic explains how to use the Console or the API to manage Data Guard associations in your VM cluster.
- Prerequisites for Using Oracle Data Guard with Exadata Database Service on Cloud@Customer Review the list of prerequisites for using Data Guard with Exadata Database Service on Cloud@Customer.
- Working with Data Guard Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.
- Using the Console to Manage Oracle Data Guard Associations Learn how to enable a Data Guard association between databases, change the role of a database in a Data Guard association using either a switchover or a failover operation, and reinstate a failed database.
- Using the API to Manage Data Guard Associations on an Exadata Database Service on Cloud@Customer System
 Learn how to use the API to manage Data Guard associations on an Exadata Database Service on Cloud@Customer system.



About Using Oracle Data Guard with Exadata Database Service on Cloud@Customer

This topic explains how to use the Console or the API to manage Data Guard associations 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.
- You are limited to one standby database for each primary database.
- The standby database is deployed as an open, read-only database (Active Data Guard).

To configure a Data Guard system between on-premises and Exadata database compute nodes, or to configure your database with multiple standbys, you must access the database host directly and set up Data Guard manually.

For complete information on Oracle Data Guard, see the Data Guard Concepts and Administration documentation on the Oracle Document Portal.

Related Topics

- Data Guard Concepts and Administration
- Oracle Document Portal

Prerequisites for Using Oracle Data Guard with Exadata Database Service on Cloud@Customer

Review the list of prerequisites for using Data Guard with 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 For Data Guard operations to work, the SYS password and the TDE wallet password of the primary and standby databases must all be the same.
- Adding a Node to a VM Cluster
- Removing a Node from a VM Cluster

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

For Data Guard operations to work, the SYS password and the TDE wallet password of the primary and standby databases must all be the same.

If you change any one of these passwords, you must update the rest of the passwords to match.

If you make any change to the TDE wallet (such as adding a master key for a new PDB or changing the wallet password), you must copy the wallet from the primary to the standby so that Data Guard can continue to operate. For Oracle Database versions earlier than 12.2, if you change the SYS password on one of the peers, you need to manually sync the password file between the DB systems.

Adding a Node to a VM Cluster

When adding a node to a VM cluster, an instance of the Data Guard database is automatically created on the new node. However, metadata updation on the remote database, that is, the primary database if addition is done on the standby database and vice versa, must be done manually.

This can be done by copying over the addinstance JSON file, /var/opt/oracle/ dbaas_acfs/<dbname>/addInstance.json created at the end of instance addition and running the /var/opt/oracle/ocde/rops update_instance <dbname> <path to addInstance JSON> command on any node of the remote cluster.

Related Topics

• 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. However, deletion of the corresponding metadata on the remote database, that is, the primary database if removal is done on the standby database and vice versa, must be done manually.

This can be done by running the /var/opt/oracle/ocde/rops remove_instance <dbname> <Instance Name> command on any node of the remote cluster.

Related Topics

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



Working with Data Guard

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

The Data Guard implementation requires two databases, one in a primary role and one in a standby role. The two databases compose a Data Guard association. Most of your applications access the primary database. The standby database is a transactionally consistent copy of the primary database.

Data Guard maintains the standby database by transmitting and applying redo data from the primary database. If the primary database becomes unavailable, you can use Data Guard to switch or fail over the standby database to the primary role.

- Switchover A switchover reverses the primary and standby database roles.
- Failover

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 association.

Switchover

A switchover reverses the primary and standby database roles.

Each database continues to participate in the Data Guard association 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 association 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 association.

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



Note:

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

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

Using the Console to Manage Oracle Data Guard Associations

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

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

- Using the Console to Enable Data Guard on an Exadata Database Service on Cloud@Customer System Learn to enable Data Guard association between databases.
- Using the Console to View Data Guard Associations of Databases in an Exadata VM Cluster
 To view the role of each database in a Data Guard association in an Exadata VM Cluster, follow this procedure.
- Using the Console To View and Edit Data Guard Associations
 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 initiate a switchover operation by using the Data Guard association of the primary database.
- Using the Console To Perform a Database Failover
 You initiate a failover operation by using the Data Guard association of the standby database.
- 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 Association on an Exadata Database
 Service on Cloud@Customer System

On a VM cluster, you remove a Data Guard association by terminating the standby database.



Using the Console to Enable Data Guard on an Exadata Database Service on Cloud@Customer System

Learn to enable Data Guard association 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.

1. Open the navigation menu. Under Oracle 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 database for which you want to assume the primary role, and then click the name of that database.
- 4. Under Resources, click Data Guard Associations.
- 5. Click Enable Data Guard.
- 6. On the Enable Data Guard page, configure your Data Guard association.
 - Data Guard association 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, Far Sync, Global Data Services, and Application Continuity.
 - * **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 association.

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 this Data Guard association.

Async - Asynchronous transport mode used with Maximum Performance protection mode.

Sync - Synchronous transport mode used with Maximum Performance and Maximum Availability protection mode.

- 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: 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.

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.
- 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>.

- 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 database.
 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.
- (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
- 7. Click Enable Data Guard.



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 Associations of Databases in an Exadata VM Cluster

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

- 1. Open the navigation menu. Under Oracle 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.

Using the Console To View and Edit Data Guard Associations

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 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. Under Resources, click Data Guard Associations.

A list of Data Guard Associations is displayed with the **Data Guard Type** you have chosen for each Data Guard Association.

6. To edit Data Guard Association details, click the Actions icon (three dots), and then click Edit Data Guard Association.

Edit Data Guard Association screen is displayed.

- 7. Do the following on the Edit Data Guard Association screen.
 - Select an applicable **Data Guard type**.



• Select the **Protection mode**.



- Set the Database password.
- Click Edit Data Guard to save the changes.

Using the Console To Perform a Database Switchover

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

1. Open the navigation menu. Under Oracle 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.
- 5. Under Resources, click Data Guard Associations.
- 6. For the Data Guard association 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 **OK**.

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

Using the Console To Perform a Database Failover

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

1. Open the navigation menu. Under Oracle 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. Under Resources, click Data Guard Associations.



- 6. For the Data Guard association 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 **OK**.

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 by using its Data Guard association.

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.

Set your <code>ORACLE_UNQNAME</code> 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 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. Under Resources, click Data Guard Associations.

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 **OK**.

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

Using the Console To Terminate a Data Guard Association on an Exadata Database Service on Cloud@Customer System

On a VM cluster, you remove a Data Guard association by terminating the standby database.

1. Open the navigation menu. Under Oracle 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 **OK**.

Using the API to Manage Data Guard Associations on an Exadata Database Service on Cloud@Customer System

Learn how to use the API to manage Data Guard associations on an 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.

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.

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 Exadata Database Service on Cloud@Customer
 Customer-managed keys for 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 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 vaults from one compartment to another.
- Setting Up Your Exadata Cloud@Customer to Work With Oracle Key Vault
- Managing Your Key Store
- Administer Transparent Data Encryption (TDE) Keys
 Use this procedure to change the encryption management configuration or rotate the
 TDE key.
- 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)

Customer-Managed Keys in Exadata Database Service on Cloud@Customer

Customer-managed keys for 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 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 Exadata 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

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 Autonomous Databases 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.

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 vaults from one compartment to another.

After you move a vault to a new compartment, inherent policies apply immediately and affect access to the vault. Moving a vault doesn't affect access to any keys or secrets that the vault contains. You can move a key or secret from one compartment to another independently of moving the vault it's associated with. For more information, see "Managing Compartments".

Related Topics

Managing Compartments

Setting Up Your Exadata Cloud@Customer to Work With Oracle Key Vault

Prerequisites

- 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 to access the OKV server.
- 2. Ensure that the REST interface is enabled from the OKV user interface.
- Create "OKV REST Administrator" user. You can use any qualified username of your choice, for example, "okv_rest_user". For ADB-C@C and ExaDB-C@C, use the same or different REST users. Those databases can be key-managed in the same or different on-prem OKV clusters. ExaDB-C@C needs REST user with create endpoint privilege. ADB-C@C needs REST user with create endpoint and create endpoint group privileges.
- 4. Gather OKV administrator credentials and IP address, which is required to connect to OKV.
- 5. Open the ports 443, 5695, and 5696 for egress on the client network to access the OKV server.

For more information, see *Network Port Requirements*, *Managing Oracle Key Vault Users*, and *Managing Administrative Roles and User Privileges*

• Step 1: Create a Vault in OCI Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password



- Step 2: Create a Dynamic Group and a Policy Statement for Key Store to Access Secret in OCI Vault
- Step 3: Create a Dynamic Group and a Policy Statement for Exadata Infrastructure to Key Store
- Step 4: Create a Policy Statement for Database Service to Use Secret from OCI Vault
 Service
- Step 5: 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 Dynamic Group and a Policy Statement for Key Store to Access Secret in OCI Vault

To grant your Key Store resources permission to access Secret in OCI Vault, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the Secret you created in the OCI Vaults and Secrets.

When defining the dynamic group, you identify your Key Store resources by specifying the OCID of the compartment containing your Key Store.

1. Copy the OCID of the compartment containing your Key Store resource.

You can find this OCID on the Compartment Details page of the compartment.



2. Create a dynamic group by following the instructions in "To create a dynamic group" in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

```
ALL {resource.compartment.id ='<compartment-ocid>'}
```

where **<compartment-ocid>** is the OCID of the compartment containing your Key Store resource.

3. After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your vaults and secrets. Then, add a policy statement of this format:

```
allow dynamic-group <dynamic-group> to use secret-family in compartment <vaults-and-secrets-compartment>
```

where *<dynamic-group>* is the name of the dynamic group you created and *<vaults-and-secrets-compartment>* is the name of the compartment in which you created your vaults and secrets.

Related Topics

• To create a dynamic group

Step 3: Create a Dynamic Group and a Policy Statement for Exadata Infrastructure to Key Store

To grant your Exadata infrastructure resources permission to access Key Store, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the Key Store you created.

When defining the dynamic group, you identify your Exadata infrastructure resources by specifying the OCID of the compartment containing your Exadata infrastructure.

1. Copy the OCID of the compartment containing your Exadata infrastructure resource.

You can find this OCID on the Compartment Details page of the compartment.

2. Create a dynamic group by following the instructions in "To create a dynamic group" in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

ALL {resource.compartment.id = '<compartment-ocid>'}

where **<compartment-ocid>** is the OCID of the compartment containing your Exadata infrastructure resource.

3. After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your Key Store. Then, add a policy statement of this format:

Allow dynamic-group **<dynamic-group>** to use keystores in compartment **<key-store-compartment>**



where *<dynamic-group>* is the name of the dynamic group you created and *<key-store-compartment>* is the name of the compartment in which you created your Key Store.

Step 4: Create a Policy Statement for Database Service to Use Secret from OCI Vault Service

To grant the Autonomous Database service permission to use the secret in OCI Vault to log in to the OKV REST interface, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your OCI Vaults and Secrets. Then, add a policy statement of this format:

allow service database to read secret-family in compartment **<vaults-and**secrets-compartment>

where <vaults-and-secrets-compartment> is the name of the compartment in which you created your OCI Vaults and Secrets.

Once the OCI Vault is set up and the IAM configuration is in place, you are now ready to deploy your Oracle Key Vault 'Key Store' in OCI and associate it with your Exadata Cloud@Customer VM Cluster.

Step 5: Create Key Store

Follow these steps to create a Key Store to connect to an on-premises encryption key appliance such as Oracle Key Vault (OKV).

- 1. Open the navigation menu. Under Oracle Database, click Exadata 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 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 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 Key Store.
- 7. Ensure that you use the same "*okv_rest_user*" user credentials, while provisioning Autonomous Database.

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 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.
- 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 Database, click Exadata 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.



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 Database, click Exadata 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 Cloud@Customer system from one compartment to another compartment.

- 1. Open the navigation menu. Under Oracle Database, click Exadata 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 Move Resource.
- 6. On the Move Resource to a Different Compartment page, select the new compartment.
- 7. Click Move Resource.

Delete a Key Store

You can delete a Key Store only if it is not associated with any CDBs.

- 1. Open the navigation menu. Under Oracle Database, click Exadata 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 **Delete**.
- 6. On the Delete Key Store dialog, click Delete.

View Key Store Associated Container Database Details

Follow these steps to view details of the container database associated with a Key Store.

- 1. Open the navigation menu. Under Oracle 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 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
Choose between customer-managed and Oracle-managed encryption	CreateAutonomousContainerDatabase
Get the Key Store (OKV or Oracle-managed) and OKV wallet name	GetAutonomousContainerDatabase
Rotate OKV and Oracle-managed key	RotateAutonomousContainerDatabaseKey
Get the Key store (OKV or Oracle-managed) and OKV wallet name	GetAutonomousDatabase
Rotate OKV and Oracle-managed key	RotateAutonomousDatabaseKey
Get the Key Store (OKV or Oracle-managed) and OKV wallet name	GetAutonomousDatabaseBackup

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 encryption management configuration or rotate the TDE key.

After you provision a database in an ExaDB-C@C system, you can change the encryption key management to OKV and rotate the TDE key for that database.



Note:

- Oracle supports administering encryption keys on databases after and including Oracle Database 11g release 2 (11.2.0.4).
- You can change encryption key management from Oracle-managed keys to customer-managed keys but you cannot change from customer-managed keys to Oracle-managed keys.
- When you change to customer-managed keys on OKV, the database will experience a shutdown abort operation followed by a restart. Plan to perform the migration to customer-managed keys on OKV in a planned maintenance window.
- To ensure that your Exadata database uses the most current version of the TDE key, rotate the key from the database details page on the Oracle Cloud Infrastructure Console. Do not use the Vault service.
- You can rotate TDE keys only on databases that are configured with customermanaged keys.
- 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 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. Click the More Actions drop-down.
- 6. Click Administer Encryption Key.

To change key management type from Oracle-managed keys to customer-managed keys:

- a. Click Change Key Management Type.
- b. Select Use customer-managed keys.

You must have a valid encryption key in 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 vault from the Vault in the Compartment drop-down.

You can change the compartment by clicking the Change Compartment link.



d. Select an encryption key from the Master encryption key in the **Compartment** drop-down.

You can change the compartment containing the encryption key you want to use by clicking the **Change Compartment** link.

- e. If you want to use an encryption key that you import into your vault, then select Change Compartment and enter the OCID of the key you want to use in the Key version OCID field.
- f. Click Apply.

To rotate an encryption key on a database using customer-managed keys:

- a. Click Rotate Encryption Key to display a confirmation dialog.
- b. Click Rotate Key.

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 to customer-managed keys, deleting the key from the OKV will cause the database to become unavailable.

On the database details page for this database, the Encryption section displays the encryption key name and the encryption key OCID.

Related Topics

Key and Secret Management

How to Manually Clone a Pluggable Database (PDB) from a Remote Container Database (CDB) When Data is Encrypted with Master Encryption Key (MEK) in Oracle Key Vault (OKV)

The 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
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 -0 -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 19.0M
                                         0 --:--:--
                    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]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
server=<source_okv_server_ip1>
```

```
user=<source okv rest user>
client wallet=/var/opt/oracle/dbaas acfs/<source db name>/
okv rest cli/client wallet
log property=/var/opt/oracle/dbaas acfs/<source db name>/okv rest cli/
conf/okvrestcli logging.properties
okv client config=/u02/app/oracle/admin/<source db name>/okv home/
conf/okvclient.ora
[oracle@testserver okv rest cli]$ vi /var/opt/oracle/dbaas acfs/
<source db name>/okv rest cli/conf/okvrestcli logging.properties
[oracle@testserver okv rest cli]$ cat /var/opt/oracle/dbaas acfs/
<source db name>/okv rest cli/conf/okvrestcli logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas acfs/
<source db name>/okv rest cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvFor
matter
```

java.util.logging.ConsoleHandler.level=FINER java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.Okv Formatter

```
[oracle@testserver okv_rest_cli]$
```

d. Create the client_wallet directory.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv rest cli]$
```

e. Create OKV REST wallet using the OKV REST command-line interface.

Enter the source OKV REST password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv rest cli]$ export JAVA HOME=/usr/java/latest;
export OKV RESTCLI CONFIG=/var/opt/oracle/dbaas acfs/<source db name>/
okv rest cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas acfs/
<source db name>/okv rest cli/bin/okv admin client-wallet add --
client-wallet /var/opt/oracle/dbaas acfs/<source db name>/
okv rest cli/client wallet --wallet-user <source okv rest user>
Password:
{
 "result" : "Success"
}
[oracle@testserver okv rest cli]$ ls -ltr /var/opt/oracle/dbaas acfs/
<source db name>/okv rest cli/client wallet
total 8
-rw----- 1 oracle oinstall 0 Jun 16 01:29 ewallet.p12.lck
-rw----- 1 oracle oinstall 0 Jun 16 01:29 cwallet.sso.lck
-rw----- 1 oracle oinstall 976 Jun 16 01:29 ewallet.p12
-rw----- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv rest cli]$
```



- 3. Create a new OKV wallet to store only the PDB MEK obtained in step #1.
 - a. Get the OKV wallet name from the source PDB in the format EXA_DB_NAME_DBID_PDB_NAME_WL.

```
For example, the wallet name would be
   EXA NEWDB1 37508325141 PDB NAME WL.
   [root@testserver 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$database;
   NAME
          DB_UNIQUE_NAME
                                       DBID
                                     3750832514
   NEWDB1
            newdb1 uniq
   SQL> select value from v$parameter where name='instance name';
   VALUE
        _____
   _____
   newdb11
   SOL> exit
   Disconnected from Oracle Database 19c EE Extreme Perf Release
   19.0.0.0 - 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> --
```

```
{
    "result" : "Success",
    "value" : {
        "status" : "PENDING",
        "locatorID" : "BA5FBFE1-DB41-4425-8EE4-D58541A1E41A"
```

description "Wallet to clone <source pdb name> pdb from

<source db name>" --unique FALSE

```
}
}
[root@testserver oracle]#
```

c. Check the status until it is ACTIVE.

```
[root@testserver oracle]# export JAVA HOME=/usr/java/latest; export
OKV RESTCLI CONFIG=/var/opt/oracle/dbaas acfs/<source db name>/
okv rest cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas acfs/
<source db name>/okv rest cli/bin/okv manage-access wallet check-
status --wallet <SOURCE PDB OKV WALLET>
{
  "result" : "Success",
  "value" : {
   "status" : "PENDING"
  }
}
[root@testserver oracle]# export JAVA HOME=/usr/java/latest; export
OKV RESTCLI CONFIG=/var/opt/oracle/dbaas acfs/<source db name>/
okv rest cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas acfs/
<source db name>/okv rest cli/bin/okv manage-access wallet check-
status --wallet <SOURCE PDB OKV WALLET>
{
  "result" : "Success",
  "value" : {
   "status" : "ACTIVE",
    "wallet" : "<SOURCE PDB OKV WALLET>"
  }
}
[root@testserver oracle]#
```

- 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;
```



```
DB UNIQUE NAME
NAME
                          DBID
NEWDB1
       newdbl uniq
                        3750832514
SQL> select value from v$parameter where name='instance name';
VALUE
_____
_____
newdb11
SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdb1]$
```

b. Add **Read 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]#
```

- Store MEK from the source PDB obtained in step #1 into the OKV wallet created in step #3.
 - Add MEK (uuid obtained in step #1.b) using the OKV REST command-line interface.

Enter the source OKV Endpoint password when prompted.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest;
export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini; /var/opt/
oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/bin/okv managed-
object wallet add-member --uuid E5344379-8B16-4FE9-BF35-
F8ECB057571A --wallet <SOURCE_PDB_OKV_WALLET>
Password:
{
```

```
"result" : "Success"
}
[root@testserver oracle]#
```

- 6. Install OKV REST wallet in the target database.
 - a. Create the okv_rest_cli directory if it does not exist.

```
[root@testserver 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 ${\tt ALL}$ 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 -0 -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]$ cat /var/opt/oracle/dbaas_acfs/
<target_db_name>/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.0
kvFormatter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.lo
g.OkvFormatter
[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]$
```

- 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; NAME DB UNIQUE NAME DBID - _ _ NEWDB1 newdbl uniq 3750832514 SQL> select value from v\$parameter where name='instance name'; VALUE _____ _____ newdb11 SOL> 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 commandline interface.

[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/ okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/ <target_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --wallet <SOURCE_PDB_OKV_WALLET> --endpoint <TARGET_OKV_EP1> --access RO_MW {

```
"result" : "Success"
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access
--wallet <SOURCE_PDB_OKV_WALLET> --endpoint <TARGET_OKV_EP2> --access
RO_MW
{
    "result" : "Success"
}
[root@testserver oracle]#
```

- 8. Clone the PDB.
 - Run dbaascli to clone the PDB.



Enter the source DB SYS user password when prompted.

```
[root@testserver oracle]# dbaascli pdb remoteClone --pdbName
<source pdb name> --dbName <target db name> --
sourceDBConnectionString <source db connection string> --
targetPDBName <target pdb name>
DBAAS CLI version 23.2.1.0.0
Executing command pdb remoteClone --pdbName <source_pdb_name> --
dbName <target_pdb_name> --sourceDBConnectionString
scaqar06dvclu01-scan1.us.oracle.com:1521/
<source db unique name>.us.oracle.com --targetPDBName
<target pdb name>
Job id: 197f30e9-209e-4ec5-9700-a13f7915f8b9
Session log: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
dbaastools 2023-06-12 10-32-17-PM 188384.log
Enter REMOTE DB SYS PASSWORD:
Enter REMOTE DB SYS PASSWORD (reconfirmation):
Loading PILOT...
Session ID of the current execution is: 6848
Log file location: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
pilot 2023-06-12 10-32-35-PM 204184
_____
Running Plugin initialization job
Enter REMOTE DB SYS PASSWORD
*****
Completed Plugin_initialization job
_____
Running Validate input_params job
Completed Validate input_params job
_____
Running Perform dbca prechecks job
Completed Perform dbca prechecks job
_____
Running PDB creation job
Completed PDB creation job
_____
Running Load pdb details job
Completed Load pdb details job
_____
Running Configure pdb service job
Completed Configure pdb service job
_____
Running Configure tnsnames ora job
Completed Configure tnsnames ora job
Running Set_pdb_admin_user_profile job
Completed Set_pdb_admin_user_profile job
_____
Running Lock pdb admin user job
Completed Lock pdb admin user job
_____
Running Register ocids job
Skipping. Job is detected as not applicable.
```

```
Running Prepare_blob_for_standby_in_primary job
Skipping. Job is detected as not applicable.
______
Running Generate_dbsystem_details job
Completed Generate_dbsystem_details job
dbaascli execution completed
[root@testserver oracle]#
```

 Delete the source PDB OKV wallet created in step #3 using the OKV REST commandline interface.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet delete --
wallet <SOURCE_PDB_OKV_WALLET>
{
    "result" : "Success"
}
[root@testserver oracle]#
```

- **10.** Delete the OKV REST wallet created in step #2.
 - Delete the wallet files in the dbaas acfs directory.

```
[root@testserver oracle] # rm -f /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet/*
[root@testserver oracle]#
```

- 11. Delete OKV REST wallet created in step #6.
 - Delete the wallet files in the dbaas acfs directory.

```
[root@testserver oracle]# rm -f /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet/*
[root@testserver oracle]#
```

Source CDB and Target CDB are Encrypted with MEK in a Different OKV Server

- 1. Get the OKV object ID of the source PDB.
 - a. Get the latest encryption key of the source PDB using SQL*Plus.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<source_db_name>.env
[oracle@testserver oracle]$ sqlplus / as sysdba
SQL*Plus: Release 19.0.0.0.0 - Production on Mon Jun 12 23:13:12 2023
Version 19.19.0.0.0
Copyright (c) 1982, 2022, Oracle. All rights reserved.
Connected to:
```



```
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 -
Production
Version 19.19.0.0.0
SQL> set heading off;
SQL> alter session set container=<SOURCE PDB>;
Session altered.
SQL> select key id, keystore type, activation time from
v$encryption keys order by activation time;
0648E5D8D5559B4F0EBFB8AA5EE730401A
SOFTWARE KEYSTORE
25-MAR-23 12.01.41.075932 AM +00:00
06AFF5B6E27A954F6EBFFC77296B27C9EC
SOFTWARE KEYSTORE
25-MAR-23 11.42.51.336955 AM +00:00
SOL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0 - Production
Version 19.19.0.0.0
```

 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.

[oracle@testserver oracle]\$

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 -0 -k https://
```



```
<source_okv_server_ip1>:5695/okvrestclipackage.zip
% Total % Received % Xferd Average Speed Time Time
Time Current
Dload Upload Total Spent
Left Speed
100 3784k 100 3784k 0 0 19.0M 0 --:--:--
--:-- 19.1M
[oracle@testserver okv_rest_cli]$ unzip -q okvrestclipackage.zip
[oracle@testserver okv rest cli]$
```

c. Modify the okvrestcli.ini and okvrestcli_logging.properties files as follows.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
server=<source_okv_server_ip1>
user=<source_okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/client_wallet
log_property=/var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
conf/okvrestcli_logging.properties
okv_client_config=/u02/app/oracle/admin/<source_db_name>/okv_home/
conf/okvclient.ora
```

```
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvFor
matter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.Okv
Formatter
[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.
 - 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]#
```

- 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$database;
                                  DBID
NAME DB_UNIQUE_NAME
_____
                    ----- ------
NEWDB1 newdb1 uniq
                                3750832514
SQL> select value from v$parameter where name='instance name';
VALUE
_____
newdb11
SOL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0 - 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]#
```

- 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.

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]$ $OKV_HOME/bin/okvutil download -1 /home/
oracle/<source_pdb_wallet_dir> -t wallet -g <SOURCE_PDB_OKV_WALLET>
Enter new wallet password (<enter> for auto-login):
Confirm new wallet password:
Enter Oracle Key Vault endpoint password:
Download succeeded
[oracle@testserver oracle]$
```

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 --
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 19.0M
                    0
                                         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_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.OkvFor matter java.util.logging.ConsoleHandler.level=FINER java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.Okv Formatter [oracle@testserver_okv_rest_cli]\$

d. Create the client_wallet directory.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv rest cli]$
```

e. Create OKV REST wallet using the OKV REST command-line interface.

Enter the target OKV REST password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv rest cli]$ export JAVA HOME=/usr/java/latest;
export OKV RESTCLI CONFIG=/var/opt/oracle/dbaas acfs/<target db name>/
okv rest cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas acfs/
<target db name>/okv rest cli/bin/okv admin client-wallet add --
client-wallet /var/opt/oracle/dbaas acfs/<target db name>/
okv rest cli/client wallet --wallet-user <target okv rest user>
/var/opt/oracle/dbaas acfs/newdb1/okv rest cli/logs/okvrest.log.lck
Password:
{
 "result" : "Success"
}
[oracle@testserver okv rest cli]$ ls -ltr /var/opt/oracle/dbaas acfs/
newdb1/okv rest cli/client wallet
total 8
-rw----- 1 oracle oinstall 0 Jun 16 01:29 ewallet.p12.lck
-rw----- 1 oracle oinstall 0 Jun 16 01:29 cwallet.sso.lck
-rw----- 1 oracle oinstall 976 Jun 16 01:29 ewallet.p12
-rw----- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv rest cli]$
```

12. Upload the source PDB wallet created in step #6 and copied to the target Cluster VM in step #9.



a. Unzip the source PDB wallet.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /home/oracle/
[oracle@testserver ~]$ unzip <source_pdb_wallet_dir>.zip
Archive: nfsa_1672104454_NFSPDB_wallet.zip
    creating: <source_pdb_wallet_dir>/
    extracting: <source_pdb_wallet_dir>/ewallet.p12
[oracle@testserver ~]$
```

b. Get OKV wallet name from the target database in the format EXA_DB_NAME_DBID_WL.

For example, the wallet name would be EXA NEWDB1 37508325141 WL.

```
[root@testserver newdb1]# su oracle
[oracle@testserver newdb1]$ source ~/<target db name>.env
[oracle@testserver newdb1]$ sqlplus / as sysdba
SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54
2023
Version 19.19.0.0.0
Copyright (c) 1982, 2022, Oracle. All rights reserved.
Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 -
Production
Version 19.19.0.0.0
SQL> select name, db unique name, dbid from v$database;
                                 DBID
NAME
      DB UNIQUE NAME
- –
NEWDB1
                               3750832514
        newdbl uniq
SQL> select value from v$parameter where name='instance name';
VALUE
_____
_____
newdb11
SOL> 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]$
```

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 scagar06dvclu01-
scan1.us.oracle.com:1521/<source db unique name>.us.oracle.com --
targetPDBName <target pdb name>
Job id: 7d4f638a-1f3a-4219-a05a-0215588dcae8
Session log: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
dbaastools 2023-06-13 01-29-09-AM 179996.log
Enter REMOTE DB SYS PASSWORD:
Enter REMOTE DB SYS PASSWORD (reconfirmation):
Loading PILOT...
Session ID of the current execution is: 6857
```

```
Log file location: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
pilot 2023-06-13 01-29-21-AM 196991
_____
Running Plugin initialization job
Enter REMOTE DB SYS PASSWORD
* * * * * * * * * * * *
Completed Plugin initialization job
_____
Running Validate input params job
Completed Validate input params job
_____
Running Perform dbca prechecks job
Completed Perform dbca prechecks job
_____
Running PDB creation job
Completed PDB creation job
_____
Running Load pdb details job
Completed Load pdb details job
_____
```



Running Configure pdb service job Completed Configure pdb_service job _____ Running Configure tnsnames ora job Completed Configure tnsnames ora job Running Set pdb admin user profile job Completed Set pdb admin user profile job _____ Running Lock pdb admin user job Completed Lock pdb admin user job Running Register ocids job Skipping. Job is detected as not applicable. _____ Running Prepare blob for standby in primary job Skipping. Job is detected as not applicable. _____ Running Generate 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]#
```

Migrate to Exadata Database Service on Cloud@Customer

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

A recommended approach for migrating to Exadata 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 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. 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
- Hardware and Software Required for Out-of-Place Cloud Upgrade to New Exadata 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 Exadata Cloud@Customer Gen2 infrastructure.
- During Out-of-Place Cloud Upgrade to New Exadata Cloud@Customer Gen2 Infrastructure
- Post Out-of-Place Cloud Upgrade to New Exadata Cloud@Customer Gen2 Infrastructure The upgrade will move your resources into Exadata Cloud@Customer Gen 2 Cloud Control Plane and onto the new generation hardware.
- Best Practices for Out-of-Place Cloud Upgrade to New Exadata 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 2020octRU 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
- 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 Exadata 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.
 - 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 Exadata Database Service on Cloud@Customer Review OCI as well as the site, network and storage requirements to prepare and deploy 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 Exadata Cloud@Customer Gen2 Infrastructure For the numose of the ungrade, the recommended tool to use is Oracle Zero

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 Exadata 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
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= 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 NONCOBTOPDB 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 zl2tgtls -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"
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=
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

During Out-of-Place Cloud Upgrade to New Exadata Cloud@Customer Gen2 Infrastructure

Monitoring: Oracle will monitor the Exadata 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 Exadata 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 Exadata Cloud@Customer Gen2.

Post Out-of-Place Cloud Upgrade to New Exadata Cloud@Customer Gen2 Infrastructure

The upgrade will move your resources into Exadata Cloud@Customer Gen 2 Cloud Control Plane and onto the new generation hardware.

Use the Gen2 OCI Console to manage your Exadata 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
- About Oracle Managed Exadata Cloud@Customer Infrastructure Maintenance Updates Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud@Customer.



Best Practices for Out-of-Place Cloud Upgrade to New Exadata 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



6

Autonomous Database on Exadata Cloud@Customer

- What's New in ADB-D on Exadata Cloud@Customer Here's a summary of the noteworthy ADB-D on Exadata Cloud@Customer additions and enhancements.
- Introduction to ADB-D on Exadata Cloud@Customer
 Oracle Autonomous Database on Exadata Cloud@Customer combines the benefits of a self-driving, self-securing, and self-repairing database management system and the security and control offered by having it deployed securely on-premise behind your firewall.
- Managing Autonomous Exadata VM Clusters An Autonomous Exadata VM Cluster is a set of symmetrical VMs across all Compute nodes.
- Managing Encryption Keys on External Devices Learn how to store and manage database encryption keys.
- Managing Autonomous Container Databases
 Learn how you can create, view, move, change backup policies, manage maintenance
 schedules, and perform other Oracle Autonomous Container Database management.
- Managing Autonomous Databases An Autonomous Database resource is a user database. When you create an Autonomous Database, you choose the Autonomous Container Database for it and you specify "Data Warehouse" or "Transaction Processing" as its workload type to create an Autonomous Data Warehouse database or an Autonomous Transaction Processing database.
- Connecting to Autonomous Databases Applications and tools connect to an autonomous database using Oracle Net Services (also known as SQL*Net). Oracle Net Services enables a network session from a client application to an Oracle Database server.
- Authenticate Autonomous Database Users with Kerberos Describes how to configure Kerberos to authenticate Oracle Autonomous Database users.
- Patching ADB on Exadata Cloud@Customer Infrastructure
- Migrating On-Premises Databases to Oracle Autonomous Database on Exadata Cloud@Customer

Zero Downtime Migration supports migrations to Oracle Autonomous Database on Exadata Cloud@Customer from any on-premises Oracle Database, including existing Exadata Cloud@Customer systems, using logical migration method and NFS as a data transfer medium.



Using Autonomous Data Guard with Autonomous Database on Exadata Cloud@Customer

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

Using Performance Hub
 Use Performance Hub to monitor database activity, diagnose issues, and tune queries to improve the performance of Oracle databases.

What's New in ADB-D on Exadata Cloud@Customer

Here's a summary of the noteworthy ADB-D on Exadata Cloud@Customer additions and enhancements.

- Provision for Free Developer Database Instances
- Disable and Enable Automatic Backups
- Ability to Scale Up or Scale Down an Autonomous Exadata VM Cluster
- DST Time-Zone File Update for an Autonomous Container Database (ACD)
- Rolling and Non-Rolling Maintenance Methods to Update an Autonomous
 Container Database
- Interactive Visualization of Resource Usage Data
- ECPU Billing Model and CPU Allocation
- X10M System Support
- In-Memory Column Store Support
- Create Autonomous Data Guard Associations and Clone an Autonomous Database Across Control Plane Regions
- Set Compartment Quotas Using Policy Statements
- Base Shape support on Oracle Autonomous Database on Dedicated Exadata
 Infrastructure
- Bring Your Own Certificates (BYOC)
- Use Kerberos to Authenticate Autonomous Database Users
- Long-Term Backup
- Save as Stack
- Autonomous Exadata VM Cluster Node Subsetting
- Choose a Specific Oracle Database Version While Provisioning Autonomous Container Databases (ACDs)
- Autonomous Data Guard Enhancements
- Configure Non-Default SCAN Listener Ports
- Clone an Autonomous Database Across Exadata Infrastructure Within the Same Region
- Character Set Selection
- Enhanced Resource Tracking for Autonomous Database
- Longer Database Name



- Create a New Autonomous Database Instance from Backup
- Multiple Autonomous VM Cluster Support
- Automatic Failover with a Standby Autonomous Container Database
- X9M-2 System Support
- Fractional OCPU and GB Storage
- Autonomous Data Guard Enabled Autonomous Database and Oracle Key Vault (OKV)
 Integration
- Infrastructure Patching
- Access Control List (ACL) to Restrict Access to Autonomous Data Guard Enabled Autonomous Databases
- ADB-D on Exadata Cloud@Customer: Monitor Performance with Autonomous Database Metrics
- ADB-D on Exadata Cloud@Customer: Autonomous Data Guard
- Access Control List (ACL) to Restrict Access to Autonomous Databases
- Oracle Key Vault (OKV) Integration
- X8M-2 System Support
- Per-Second Billing for Autonomous Database OCPU Usage
- Oracle Autonomous Database on Oracle Exadata Database Service on Cloud@Customer

Provision for Free Developer Database Instances

- Services: Database
- Release Date: March 13, 2024

Customers who have subscribed to Oracle Exadata Database Service or Autonomous Database on either a Dedicated Exadata Infrastructure or Exadata Cloud@Customer can create and use Oracle Autonomous Database for Developers instances.

Autonomous Database for Developers instances are free Autonomous Databases that developers can use to build and test new applications.

With Autonomous Database for Developers instances, you can try new Autonomous Database features for free and apply them to ongoing or new development projects. Developer databases are limited in resources, so they are not suitable for large-scale testing and production deployments. When you need more compute or storage resources, you can transition to a paid database licensing by cloning your developer database into a regular Autonomous Database.

Related Topics

- Oracle Autonomous Database for Developers
 Oracle Autonomous Database for Developers instances are free Autonomous Databases that developers can use to build and test new applications.
- Create an Autonomous Database
- Create an Autonomous Data Guard Enabled Autonomous Database



- Scale the CPU Core Count or Storage of an Autonomous Database, or Enable/ Disable or Alter the Percentage of System Global Area (SGA) for IM Column Store
- Create a Long-Term Backup
- Clone an Autonomous Database
- Clone an Autonomous Database Backup
- Clone a Standby Database
- Clone a Standby Database Backup

Disable and Enable Automatic Backups

- Services: Database
- Release Date: Decmber 05, 2023

By default, automatic backups are enabled for an Autonomous Container Database (ACD). However, you can now disable them while provisioning the ACD, and enable them anytime later. Once the automatic backups are enabled you can not disable them for an ACD. You can set the backup retention period for the backups between 7 to 95 days.

Related Topics

- Create an Autonomous Data Guard Enabled Autonomous Container Database Follow these steps to create an Autonomous Data Guard Enabled Autonomous Container Database on an Oracle Exadata Cloud@Customer system.
- Edit Autonomous Container Database Backup Settings If automatic backups were disabled while provisioning an Autonomous Container Database (ACD), you can enable them later from the Oracle Cloud Infrastructure (OCI) console.

Ability to Scale Up or Scale Down an Autonomous Exadata VM Cluster

- Services: Database
- Release Date: Decmber 05, 2023

You can add resources to or remove them from an Autonomous Exadata VM Cluster (AVMC). This is called Autonomous Exadata VM Cluster vertical scaling.

This ability lets you scale up or scale down an AVMC by modifying the following resources:

- Number of CPUs per VM
- Number of Autonomous Container Databases
- Database storage

Related Topics

Scale Autonomous Exadata VM Cluster Resources



DST Time-Zone File Update for an Autonomous Container Database (ACD)

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

You can now choose to include or exclude a time-zone file update along with your quarterly automatic maintenance patches, either while provisioning an Autonomous Container Database (ACD) or for an existing ACD.

Updating the time-zone file would require complete downtime for the ACD and the associated Autonomous Databases. The downtime is dependent on the amount of data that is time-zone sensitive.

You can also schedule an on-demand maintenance to update RU (Release Update) along with the time-zone file or just the time-zone file for an ACD.

Related Topics

- Create an Autonomous Container Database
- View Details of an Autonomous Container Database
- Update Autonomous Container Database Maintenance Preferences
- Schedule a Quarterly Maintenance Update

Rolling and Non-Rolling Maintenance Methods to Update an Autonomous Container Database

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

You now have the option to choose between rolling or non-rolling maintenance methods to apply quarterly automatic maintenance patches. You can configure the maintenance method while provisioning an Autonomous Container Database (ACD) or for an existing ACD. The non-rolling maintenance method involves a full system downtime.

Related Topics

- Create an Autonomous Container Database
- View Details of an Autonomous Container Database
- Update Autonomous Container Database Maintenance Preferences
- Schedule a Quarterly Maintenance Update

Interactive Visualization of Resource Usage Data

- Services: Database
- Release Date: October 10, 2023



You can monitor and track your Exadata resource allocation and usage across Autonomous Exadata VM Clusters (AVMC) and Autonomous Container Databases (ACD) with enhanced granularity and new visuals on the Oracle Cloud Infrastructure (OCI) console, in real time.

This release lets you get a comprehensive and clear picture of the compute and storage resources' allocation and usage at the AVMC and ACD levels from their Details page on the OCI console. Depending on your choice, you can view this information in a graphical or table view.

Understanding and knowing these details can help you optimize your resource allocation to Autonomous Databases and efficiently forecast your capacity needs.

Related Topics

- Resource Terminology
- View a List of Autonomous Exadata VM Clusters
- View Details of an Autonomous Exadata VM Cluster
- View Details of an Autonomous Container Database
- Create an Autonomous Database
- Scale the CPU Core Count or Storage of an Autonomous Database, or Enable/ Disable or Alter the Percentage of System Global Area (SGA) for IM Column Store
- Clone an Autonomous Database
- Clone an Autonomous Database Backup
- Using the API to Manage Autonomous Exadata VM Clusters
- Using the API to Manage Autonomous Container Databases

ECPU Billing Model and CPU Allocation

- Services: Database
- Release Date: October 10, 2023

With this enhancement, you will be able to choose between ECPU or OCPU compute models while configuring your Autonomous Database resources.

During AVM creation, you will have the option to allocate CPUs to the AVM using either ECPU or OCPU. On ADB-C@C, where Multi VM is supported, you will also be able to size the AVM by assigning the number of ECPUs per Node and Memory per ECPU in the sizing widget. ACDs and ADBs will inherit the CPU type from their parent AVM. For example, if the AVM is configured using OCPU, the usage of the ACD and ADB resources in that AVM will be tracked using OCPU, such as Total/Available/ Reclaimable/Provisionable OCPUs, and so on. There is no mixing of ADB CPU types inside an AVM.

 ECPU: An ECPU is an abstracted measure of compute resources. ECPU is based on the number of cores elastically allocated from a pool of compute and storage servers.

While provisioning a new database, cloning an existing database, and scaling up or down the CPU resources of an existing database:

The CPU count defaults to 2 ECPUs.



- For databases that need 2 or more ECPUs, you must specify the number of assigned ECPUs in increment of 1.
- CPU over-provisioning is not supported on ECPUs. This applies to all ADB-C@C workloads.
- **OCPU**: An OCPU is a physical measure of compute resources. OCPUs are based on the physical core of a processor with hyper-threading enabled.

While provisioning a new database, cloning an existing database, and scaling up or down the CPU resources of an existing database:

- The CPU count defaults to 1 OCPU.
- For databases that do not need an entire OCPU, you can assign OCPU count from 0.1 to 0.9 in increments of 0.1 OCPU. This allows you to over-provision CPU at the system level and run more databases on each infrastructure instance.
- For databases that need more than 1 OCPU, you must specify the number of assigned cores as an integer. For example, you cannot assign 3.5 OCPUs to a database. The next available number of OCPUs above 3 is 4.

Related Topics

- Create an Autonomous Exadata VM Cluster
- Create an Autonomous Data Guard Enabled Autonomous Container Database Follow these steps to create an Autonomous Data Guard Enabled Autonomous Container Database on an Oracle Exadata Cloud@Customer system.
- Create an Autonomous Data Guard Enabled Autonomous Database
- Enable or Disable Auto Scaling for an Autonomous Database
- Create an Autonomous Database
- Clone an Autonomous Database
- Clone an Autonomous Database Backup
- Clone a Standby Database
- Clone a Standby Database Backup

X10M System Support

- Services: Database
- Release Date: September 13, 2023

Autonomous Database on Exadata Cloud@Customer has been extended to support X10M system.

Related Topics

Available Exadata Infrastructure Hardware Shapes

In-Memory Column Store Support

- Services: Database
- Release Date: August 22, 2023



The In-Memory Column Store (IM column store) stores tables and partitions in memory using a columnar format optimized for rapid scans. Oracle Database uses a sophisticated architecture to manage data in columnar and row formats simultaneously.

When you enable an IM column store, the SGA manages data in separate locations: the InMemory Area and the database buffer cache. For more information, see *In-Memory Column Store Architecture*.

With this enhancement, you can:

- Enable IM column store during creation or enable/disable IM column store on an existing ADB without any downtime or restarts as long as a minimum of 4 OCPUs are allocated.
- Select any memory ratio between 5%-70% and adjust this dynamically without any downtime.
- Alter or create new objects and specify they are stored in the IM column store.
- Manually scale up or down the OCPU for an ADB running In-Memory. The In-Memory size will adjust following the percentage specified and does not require an ADB restart

In-Memory Initialization Parameters

Several initialization parameters control the behavior of the IM column store.

- INMEMORY QUERY: Specifies whether In-Memory queries are allowed.
- OPTIMIZER_INMEMORY_AWARE: Controls the optimizer cost model enhancements for Database In-Memory.
- INMEMORY_OPTIMIZED_ARITHMETIC: Controls whether NUMBER columns are stored in an In-Memory optimized format.

For more information, see:

- In-Memory Initialization Parameters in the Oracle® Database In-Memory Guide
- Initialization Parameters in the Oracle® Database Reference

REST API Endpoint	Description
CreateAutonomousDatabase	Creates a new Autonomous Database.
UpdateAutonomousDatabase	Updates one or more attributes of the specified Autonomous Database. See UpdateAutonomousDatabaseDetails reference for a full list of attributes that can be updated.
GetAutonomousDatabase	Gets the details of the specified Autonomous Database.
ListAutonomousDatabases	Gets a list of Autonomous Databases based on the query parameters specified.

Related Topics

- In-Memory Column Store Architecture
- In-Memory Initialization Parameters
- Initialization Parameters
- CreateAutonomousDatabase



- UpdateAutonomousDatabase
- GetAutonomousDatabase
- ListAutonomousDatabases
- Create an Autonomous Database
- Scale the CPU Core Count or Storage of an Autonomous Database, or Enable/Disable or Alter the Percentage of System Global Area (SGA) for IM Column Store
- View Details of an Autonomous Database
- Create an Autonomous Data Guard Enabled Autonomous Container Database Follow these steps to create an Autonomous Data Guard Enabled Autonomous Container Database on an Oracle Exadata Cloud@Customer system.
- Clone an Autonomous Database
- Clone an Autonomous Database Backup

Create Autonomous Data Guard Associations and Clone an Autonomous Database Across Control Plane Regions

- Services: Database
- Release Date: July 11, 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.

With this enhancement, you can:

- create Autonomous Data Guard associations between Autonomous Container Databases in two different control plane regions
- clone Autonomous Databases across regions, for example, clone an Autonomous Database from an Autonomous Container Database in Region A to an Autonomous Container Database in Region B

Related Topics

- Create an Autonomous Data Guard Enabled Autonomous Container Database Follow these steps to create an Autonomous Data Guard Enabled Autonomous Container Database on an Oracle Exadata Cloud@Customer system.
- Clone an Autonomous Database
- Clone an Autonomous Database Backup

Set Compartment Quotas Using Policy Statements

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

With this enhancement, you can:

- view tenancy's limits, quotas, and usage in the OCI console for Oracle Autonomous Database on Exadata Cloud@Customer
- set Compartment quotas using Policy statements



Related Topics

- Manage Resource Availability with Compartment Quotas
- Overview of Compartment Quotas
- Quota Policy Quick Start
- Authentication and Authorization
- Service Limits
- Viewing Your Service Limits, Quotas, and Usage
- Requesting a Service Limit Increase
- Database Limits

Base Shape support on Oracle Autonomous Database on Dedicated Exadata Infrastructure

- Services: Database
- Release Date: April 18, 2023

Oracle Autonomous Database on Dedicated Exadata Infrastructure comes in different infrastructure shapes to support workloads of different sizes. In this release, the capability of Autonomous Database on Dedicated Exadata Infrastructure has been extended to support Base shape.

Related Topics

Available Exadata Infrastructure Hardware Shapes

Bring Your Own Certificates (BYOC)

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

This feature lets you use your CA-signed server-side Oracle REST Data Services (ORDS) or Secure Socket Layer (SSL) certificates while rotating the security certificates of your AVMC resource.

Choose a method to seed and rotate both database and ORDS SSL certificates with the following seeding and rotation options.

- Stick to the default service-provided self-signed certificates (default option).
- Integrate with OCI Certificate service to pick up CA, CA Bundle, and Certificate.

Related Topics

- Create an Autonomous Exadata VM Cluster
- Rotate Oracle Database TLS Certificate and Oracle REST Data Services (ORDS)
 TLS Certificate

Use Kerberos to Authenticate Autonomous Database Users

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

This enhancement enables you to use Kerberos to authenticate database users centrally along with your choice of directory service for user management.

Related Topics

 Authenticate Autonomous Database Users with Kerberos Describes how to configure Kerberos to authenticate Oracle Autonomous Database users.

Long-Term Backup

- Services: Database
- Release Date: March 07, 2023

With this enhancement, you can take an on-demand backup that lives outside of the defined ACD retention period. In a range of 90 days to 10 years, you can specify any number of days, months, or years.

- In Data Guard environments, long-term backups will happen on the primary or standby database wherever the action was initiated. So, you must make two separate requests to back up the primary and standby databases, if desired.
- A long-term backup will be available as long as the ADB is available either in a running or stopped state (not terminated).
- You can use a long-term backup only to create a new database. You cannot use it for an in-place/PITR restore.
- Long-term backups include all the necessary data files, archive logs, control file, and server parameter file (SPFILE) to restore the database on the same or different ADB-D infrastructure.
- You can delete a long-term backup whenever you want.
- Creating an ADB from a long-term backup will always be upgraded to the currently supported version (target ACD version).
 For example, a 5-year-old backup may not have 19c available and Oracle will upgrade the database to 23c. However, the ADB thus created will not guarantee that an application will work just because the data is available.

Related Topics

- Create a Long-Term Backup
- View Details and Edit Retention Period of a Long-Term Backup
- Delete a Long-Term Backup
- Clone an Autonomous Database Backup
- Terminate an Autonomous Database

Save as Stack

- Services: Database
- Release Date: February 15, 2023



Stack is a collection of Oracle Cloud Infrastructure resources corresponding to a given Terraform configuration. Each stack resides in the compartment you specify, in a single region; however, resources on a given stack can be deployed across multiple regions. For more information, see *stack*.

With this enhancement, while provisioning an Exadata Infrastructure, VM Cluster, Autonomous VM Cluster, Autonomous Container Database, and Autonomous Database, you can save resource configuration as a stack. Use the stack to install, configure, and manage the resource through the Resource Manager service. For requirements and recommendations for Terraform configurations used with Resource Manager, see *Terraform Configurations for Resource Manager*. To provision, the resources defined in your stack, follow the steps outlined in *Creating an Apply Job*.

Related Topics

- stack
- Creating an Apply Job
- Terraform Configurations for Resource Manager
- Using the Console to Create Exadata Cloud@Customer Infrastructure To create your Exadata Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- Using the Console to Create a VM Cluster
 To create your VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- Create an Autonomous Exadata VM Cluster
- Create an Autonomous Container Database
- Create an Autonomous Database

Autonomous Exadata VM Cluster Node Subsetting

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

Autonomous Exadata VM Cluster node subsetting enables you to allocate a subset of database servers to new Autonomous Exadata VM Clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.

With Autonomous Exadata VM Cluster node subsetting, you can:

 Create a smaller Autonomous Exadata 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.

Consider reviewing the points below that will assist you in subsetting Autonomous Exadata VM Cluster nodes.

- Autonomous Exadata VM Cluster node subsetting capability is available for new Autonomous Exadata VM Clusters in Gen2 Exadata Cloud@Customer service.
- All VMs across an Autonomous Exadata 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 Autonomous Exadata VM Cluster.



- Any Autonomous Exadata VM Cluster should have a minimum of 2 VMs even with the node subsetting capability. We currently do not support clusters with a single VM.
- Each Autonomous Exadata 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 Autonomous Exadata VM Cluster and is validated to ensure the IP addresses do not overlap with other cluster networks.

Maximum number of clusters across the infrastructure depends on the resources available per DB Server and is subject to the per DB Server maximum VM limit.

Related Topics

- Create an Autonomous Exadata VM Cluster
- View a List of DB Servers on an Exadata Infrastructure

Choose a Specific Oracle Database Version While Provisioning Autonomous Container Databases (ACDs)

- Services: Database
- Release Date: February 07, 2023

Autonomous Container Database (ACD) is provisioned with the latest Oracle Database version and it is constantly patched with production fixes and the latest Release updates (RU) every quarter. When you provision an ACD, the new version update that is pushed to production becomes the default. You will not have the option of choosing a database version while provisioning an ACD.

With this enhancement, while provisioning an ACD, you can choose a database version of your choice.

Additionally, you can create an ACD using the previous Oracle Database version. If the latest database version supported for creating ACD is 19.15, then you can select Oracle Database version 19.14.

The enhancement also implements Oracle Application Name Tags for database deployments. Oracle Database deployments for an Oracle application may require specific database versions. Setting an appropriate application tag key in the <code>Oracle-ApplicationName</code> tag namespace while deploying an Autonomous Database will ensure that the certified database version is deployed for that particular application.

Related Topics

- Create an Autonomous Data Guard Enabled Autonomous Container Database
 Follow these steps to create an Autonomous Data Guard Enabled Autonomous
 Container Database on an Oracle Exadata Cloud@Customer system.
- Configure Automatic Maintenance Schedule for a Data Guard Enabled Autonomous Container Database

Autonomous Data Guard Enhancements

- Services: Database
- Release Date: February 07, 2023

With these enhancements, you can:



- Change the protection mode from Maximum Performance to Maximum Availability or vice versa on ACDs.
- Change Fast-Start Failover (FSFO) lag limit in increments of 1. Minimum: 5 and Maximum: 3600 seconds. Default: 30 seconds
- · Convert physical standby to snapshot standby and vice versa

Related Topics

- Create an Autonomous Data Guard Enabled Autonomous Container Database Follow these steps to create an Autonomous Data Guard Enabled Autonomous Container Database on an Oracle Exadata Cloud@Customer system.
- View Details of a Data Guard Enabled Primary or Standby Autonomous Container Database

Follow these steps to view detailed information about a primary or standby Autonomous Container Database on an Oracle Exadata Database Service on Cloud@Customer system.

- Convert a Physical Standby ACD to Snapshot Standby ACD
 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 ACD to Physical Standby ACD A snapshot standby database will automatically convert back to a physical standby database after 7 days.
- Create an Autonomous Database
- Terminate an Autonomous Database
- Scale the CPU Core Count or Storage of an Autonomous Database, or Enable/ Disable or Alter the Percentage of System Global Area (SGA) for IM Column Store
- Perform a Failover to Standby Autonomous Container Database Initiate a failover operation by using the Data Guard association of the standby database.
- Perform a Switchover to Standby or Primary Autonomous Container Database Initiate a switchover operation by using the Data Guard association of the primary database.
- Restore an Autonomous Database

Configure Non-Default SCAN Listener Ports

- Services: Database
- Release Date: November 15, 2022

This feature enables you to:

- Choose a Single Client Access Name (SCAN) listener port for Transport Layer Security (TLS) and non-TLS from a range of available ports. For TLS, the default is 2484; for non-TLS, it is 1521.
- Choose between one-way TLS and mutual TLS (mTLS) for database client connections. ORDS is always one-way TLS as it currently is.

Related Topics

Create an Autonomous Exadata VM Cluster



Clone an Autonomous Database Across Exadata Infrastructure Within the Same Region

- Services: Database
- Release Date: August 09, 2022

With this enhancement, you can clone an ADB from an ACD in one Exadata Infrastructure to an ACD in another Exadata Infrastructure within the same region.

Related Topics

- Clone an Autonomous Database
- Clone an Autonomous Database Backup
- Clone a Standby Database
- Clone a Standby Database Backup

Character Set Selection

- Services: Database
- Release Date: May 17, 2022

This feature enhancement enables you to specify the database character set and the national character set while creating an ADB.

- While creating an ADB, you can choose only the character sets supported. To get the complete list of supported character sets, use *ListAutonomousDatabaseCharacterSets* API.
- You cannot change the character set after creating an ADB.
- You can only set the character sets while creating an ADB. Character set selection is not supported for cloning (meta or full) or creating from a backup.
- The CDB will remain using AL32UTF8 for the database character set and AL16UTF16 for the national character set.
- Each ADB in an ACD can have different combinations of database/national character sets.

Related Topics

- Create an Autonomous Database
- API to Manage Autonomous Databases

Enhanced Resource Tracking for Autonomous Database

- Services: Database
- Release Date: May 10, 2022

Identify where the reclaimable OCPUs are located to reclaim them to create new database resources or for scaling automation when a certain threshold is reached.



You can create multiple AVMCs on an Exadata Infrastructure resource. The OCPUs that you allocate while provisioning the AVMC resource will be the **Total OCPUs** available for its Autonomous Databases. When you create multiple AVMCs, each AVMC can have its own value for total OCPUs.

At an AVMC or ACD level, the total number of OCPUs available for creating databases is called **Available OCPUs**.

This feature enhancement provides better resource usage tracking:

• Reclaimable OCPUs:

When an Autonomous Database is terminated or scaled down, the number of OCPUs allocated to it is not immediately returned to the available OCPUs at its parent AVMC level for the overall deployment. They continue to be included in the count of OCPUs available to its parent container database until that parent container database is restarted. These OCPUs are called **reclaimable OCPUs**. Reclaimable OCPUs at the parent AVMC level are the sum of reclaimable OCPUs of all its ACDs. When an ACD is restarted, it returns all its reclaimable OCPUs to the available OCPUs at its parent AVMC level.

• Failure free ADB Provisioning with Provisionable OCPUs:

Based on the resource utilization on each node; not all the values of the available OCPUs can be used to provision or scale Autonomous Databases. For example, suppose you have 20 OCPUs available at the AVMC level, not all the values from 1 to 20 OCPUs can be used to provision or scale Autonomous Databases depending on the resource availability at the node level. The list of OCPU values that can be used to provision or scale an Autonomous Database is called **Provisionable OCPUs**.

On the console, when you try to provision or scale an Autonomous Database, the OCPU count will be validated against the list of provisionable OCPUs, and if the value is not provisionable, you will be provided with the two nearest provisionable OCPU values. Alternatively, if you want to see the complete list of provisionable OCPU values for an Autonomous Exadata VM Cluster, you can use the following API:

GetAutonomousContainerDatabase returns a list of provisionable OCPU values that can be used to create a new Autonomous Database in the given Autonomous Container Database. See *GetAutonomousContainerDatabase* for more details.

GetAutonomousDatabase returns a list of provisionable OCPU values that can be used for scaling a given Autonomous Database. See *GetAutonomousDatabase* for more details.

Autonomous VM Clusters

As a Fleet Administrator, you will be able to identify:

- **Total OCPUs:** The number of CPU cores allocated to the VM Cluster.
- Exadata Storage (TB): The storage allocated to the VM Cluster in TBs.
- Total Autonomous Database Storage (TB): Total Autonomous Database Storage storage allocated to the Autonomous VM Cluster.
- Available Autonomous Database Storage (TB): Storage available to create Autonomous Databases in the Autonomous VM Cluster
- Total Memory: The memory allocated to the VM Cluster in GB.
- Total Local Storage (GB): Total Local Storage in the Exadata Infrastructure.



- Available OCPUs: CPU cores available for allocation to ADBs.
- **Reclaimable OCPUs:** Sum of all the reclaimable OCPUs in all the ACDs in the Autonomous VM Cluster.
- Available ACDs: The number of ACDs you can create in the AVM using the available resources.
- Total ACDs: Number of ACDs customers want to create in the AVM.
- Which ACDs have a lot of reclaimable OCPUs and restart them when needed.

Autonomous Container Database

As a Database Administrator, you will be able to identify:

- Total OCPUs: The number of CPU cores allocated to the VM cluster.
- Available OCPUs: Sum of OCPU available on the Autonomous VM Cluster + Reclaimable OCPU within the same ACD.
- Reclaimable OCPUs: When an Autonomous Database is terminated or scaled down, the number of OCPUs allocated to it is not immediately returned to the available OCPUs at its parent AVMC level for the overall deployment. They continue to be included in the count of OCPUs available to its parent container database until that parent container database is restarted. These OCPUs are called reclaimable OCPUs. Reclaimable OCPUs at the parent AVMC level are the sum of reclaimable OCPUs of all its ACDs. When an ACD is restarted, it returns all its reclaimable OCPUs to the available OCPUs at its parent AVMC level.

Note:

You cannot create an ACD if 2 OCPUs are not available in the VM Cluster or Exadata Infrastructure.

Autonomous Database

As a Database Administrator, you will be able to identify the OCPU count that you can use to provision or scale an ADB.

- OCPU: The number of OCPU cores to be made available to the database.
- **Storage:** The quantity of storage available to store data in the database, in terabytes.
- While provisioning an ADB, if you specify an OCPU count that the service cannot provision in the existing ACD, then the service displays an error message and suggests 2 values in close proximity to the value that you have specified.
 For example, assume that you are creating an ADB with 15 OCPUs on a Quarter Rack Exadata Infrastructure that has 20 OCPUs available. However, there are only 10 OCPUs available on each node. In this case, the service will not be able to provision the ADB because 15 OCPUs is less than the split threshold of 16 OCPUs. Therefore, the nearest possible values are 17 and 18.
- While scaling an ADB, if you specify an OCPU count that the service cannot use to scale the ADB, then the service displays an error message and suggests 2 values in close proximity to the value that you have specified.



REST API Endpoint	Description
GetAutonomousVmCluster	View a list of available OCPU, reclaimable OCPU, available Autonomous Database storage, and available ACDs.
GetAutonomousContainerDatabase	 View a list of available OCPU, total OCPU, and reclaimable OCPU. View an array of provisionable OCPUs.
GetAutonomousDatabase	View an array of scalable OCPUs.

Table 6-1 REST API Endpoints to Track and Manage Resource Usage

Related Topics

- View a List of Autonomous Exadata VM Clusters
- View Details of an Autonomous Exadata VM Cluster
- Restart an Autonomous Container Database
- Create an Autonomous Data Guard Enabled Autonomous Container Database Follow these steps to create an Autonomous Data Guard Enabled Autonomous Container Database on an Oracle Exadata Cloud@Customer system.
- Create an Autonomous Data Guard Enabled Autonomous Database
- Scale the CPU Core Count or Storage of an Autonomous Database, or Enable/ Disable or Alter the Percentage of System Global Area (SGA) for IM Column Store
- Clone an Autonomous Database
- Clone an Autonomous Database Backup
- GetAutonomousContainerDatabase
- GetAutonomousDatabase

Longer Database Name

- Services: Database
- **Release Date**: May 03, 2022

The length of the database name has been extended from 14 characters to 30 characters. 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 30 characters
- contain only alphanumeric characters
- begin with an alphabetic character
- must not contain spaces

Create a New Autonomous Database Instance from Backup

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

Restore an Autonomous Database (ADB) backup to an Autonomous Container Database (ACD) on the same or a different Autonomous Virtual Machine (AVM) running on the existing or different Autonomous Exadata Cloud@Customer system.

Prerequisites and Limitations

- If you are using customer-managed keys, then the target AVM/ACD will require access to the source Oracle Key Vault (OKV) for the keys.
- The target AVM must have access to the backup destination of the source for the restore to be possible.
- You can only use the **Full Clone** option to create a database clone.
- You cannot use disk-based backups to create ADB instances from backups.
- Target ACD must be on the same or higher version as the source.
- One ADB restore per AVM. The limit applies only to the target AVM.

Related Topics

- Clone an Autonomous Database
- Clone an Autonomous Database Backup
- Clone a Standby Database
- Clone a Standby Database Backup

Multiple Autonomous VM Cluster Support

- Services: Database
- Release Date: March 15, 2022

Multi-VM Cluster supports heterogeneous computing environments in which Autonomous and Non-Autonomous VM clusters can coexist on an Exadata Infrastructure.

With Multiple VM Cluster support you can:

- Create multiple Autonomous VM Clusters on an Exadata Infrastructure
- Schedule separate maintenance runs for each Autonomous VM Cluster
- Use different license models for each Autonomous Database

Related Topics

- Create an Autonomous Exadata VM Cluster
- View Details of an Autonomous Exadata VM Cluster
- Schedule Oracle-Managed Infrastructure Updates
 Exadata Cloud Service updates are released on a quarterly basis. You can set a
 maintenance window to determine the time your quarterly infrastructure maintenance will
 begin.
- View Details of an Autonomous Container Database
- Create an Autonomous Database
- View Details of an Autonomous Database



Clone an Autonomous Database

Automatic Failover with a Standby Autonomous Container Database

- Services: Database
- Release Date: January 18, 2022

With Fast-Start Failover (FSFO), the system automatically detects the failure of the primary Autonomous Container Database and then fails over to a designated standby Autonomous Container Database.

Automatic failover is optional while configuring Autonomous Data Guard. You can enable or disable automatic failover after configuring Autonomous Data Guard.

Related Topics

Managing a Standby Autonomous Container Database

Enabling Autonomous Data Guard on an Autonomous Container Database creates a standby (peer) Autonomous Container Database that provides data protection, high availability, and facilitates disaster recovery for the primary database.

- Perform a Failover to Standby Autonomous Container Database Initiate a failover operation by using the Data Guard association of the standby database.
- Perform a Switchover to Standby or Primary Autonomous Container Database Initiate a switchover operation by using the Data Guard association of the primary database.
- Reinstate Data Guard Enabled Standby Autonomous Container 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.

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 Available Exadata Infrastructure Hardware Shapes.

Related Topics

Available Exadata Infrastructure Hardware Shapes

Fractional OCPU and GB Storage

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



Create Autonomous Databases with less than 1 OCPU count using fractional units from 0.1 to 0.9 OCPU and GB sizing between 32 GB and the maximum usable storage for your Exadata shape.

Related Topics

- Create an Autonomous Database
- Scale the CPU Core Count or Storage of an Autonomous Database, or Enable/Disable or Alter the Percentage of System Global Area (SGA) for IM Column Store
- Clone an Autonomous Database

Autonomous Data Guard Enabled Autonomous Database and Oracle Key Vault (OKV) Integration

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

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

Related Topics

- Managing Encryption Keys on External Devices Learn how to store and manage database encryption keys.
- Create an Autonomous Data Guard Enabled Autonomous Container Database Follow these steps to create an Autonomous Data Guard Enabled Autonomous Container Database on an Oracle Exadata Cloud@Customer system.
- Create an Autonomous Data Guard Enabled Autonomous Database
- Operations Performed Using the APIs Learn how to use the API to manage Autonomous Data Guard Enabled Autonomous Container Database.
- Rotate CDB Encryption Key
- Rotate ADB Encryption Key

Infrastructure Patching

- Services: Database
- Release Date: May 04, 2021

ADB-Dedicated maintenance involves patching Exadata Infrastructure(EI) Autonomous VM Cluster, and Autonomous Container Database (ACD).

Related Topics

Patching ADB on Exadata Cloud@Customer Infrastructure

Access Control List (ACL) to Restrict Access to Autonomous Data Guard Enabled Autonomous Databases



- Services: Database
- Release Date: January 26, 2021

An access control list (ACL) provides additional protection to your database by allowing only the clients with specific IP addresses to connect to the database. You can add IP addresses individually, or in CIDR blocks.

Related Topics

- Create an Autonomous Data Guard Enabled Autonomous Database
- Manage Access Control List of an Autonomous Database

ADB-D on Exadata Cloud@Customer: Monitor Performance with Autonomous Database Metrics

- Services: Database
- Release Date: December 15, 2020

You can monitor the health, capacity, and performance of your Autonomous Databases with metrics, alarms, and notifications. You can use Oracle Cloud Infrastructure console or Monitoring APIs to view metrics.

Related Topics

Monitor Performance with Autonomous Database Metrics

ADB-D on Exadata Cloud@Customer: Autonomous Data Guard

- Services: Database
- Release Date: November 17, 2020

Enabling Autonomous Data Guard on an Autonomous Container Database on dedicated Exadata infrastructure creates a standby (peer) Autonomous Container Database that provides data protection, high availability, and facilitates disaster recovery for the primary database.

Related Topics

 Using Autonomous Data Guard with Autonomous Database on Exadata Cloud@Customer

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

Access Control List (ACL) to Restrict Access to Autonomous Databases

- Services: Database
- Release Date: November 10, 2020



An access control list (ACL) provides additional protection to your database by allowing only the clients with specific IP addresses to connect to the database. You can add IP addresses individually, or in CIDR blocks.

Related Topics

- Manage Access Control List of an Autonomous Database
- Create an Autonomous Database
- Clone an Autonomous Database

Oracle Key Vault (OKV) Integration

- Services: Database
- Release Date: October 27, 2020

Integrate your on-premises Oracle Key Vault (OKV) with Autonomous Database on Exadata 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.

X8M-2 System Support

- Services: Database
- Release Date: September 25, 2020
- Release Notes: ADB-D on Exadata Cloud@Customer: X8M-2 System Support

Per-Second Billing for Autonomous Database OCPU Usage

- Services: Database
- Release Date: July 21, 2020
- Release Notes: Exadata Cloud@Customer: Per-Second Billing for Autonomous
 Database OCPU Usage

Oracle Autonomous Database on Oracle Exadata Database Service on Cloud@Customer

- Services: Database
- Release Date: June 23, 2020
- Release Notes: Exadata Cloud@Customer: Oracle Autonomous Database



Introduction to ADB-D on Exadata Cloud@Customer

Oracle Autonomous Database on Exadata Cloud@Customer combines the benefits of a self-driving, self-securing, and self-repairing database management system and the security and control offered by having it deployed securely on-premise behind your firewall.

After purchasing Autonomous Database on Exadata Cloud@Customer and creating, provisioning and activating its Exadata Infrastructure hardware and Oracle Cloud resource, several additional resource types become available in the **Exadata Cloud@Customer** section of the Oracle Cloud Infrastructure console: Autonomous Exadata VM Clusters, Autonomous Container Databases and Autonomous Databases. You use these resources to create and manage your secure, on-premise deployment of Oracle Autonomous Database.

- Database System Architecture Overview
- User Roles
- Available Exadata Infrastructure Hardware Shapes

Database System Architecture Overview

Oracle Autonomous Database on Oracle Exadata Database Service on Cloud@Customer has a four-level database architecture model that makes use of Oracle multitenant database architecture.

- Resource Types
- Deployment Order

Resource Types

Each level of the architecture model corresponds to one of the following resources types:

 Oracle Exadata Cloud@Customer infrastructure: Hardware rack that includes compute nodes and storage servers, tied together by a high-speed, low-latency InfiniBand network and intelligent Exadata software.

Oracle Exadata Cloud@Customer infrastructure is common for both Autonomous and Non-Autonomous resources.

For a list of the hardware and Oracle Cloud resource characteristics of Oracle Exadata Cloud@Customer infrastructure resources that support Autonomous Databases, see *Available Exadata Infrastructure Hardware Shapes*.

- Only the Oracle Exadata Cloud@Customer infrastructures deployed before Oracle announced support for Autonomous Databases on Oracle Exadata Cloud@Customer do not support Autonomous resources listed below. Please contact your Oracle sales representative to understand the infrastructure upgrades required for supporting Oracle Autonomous Databases.
- You can create only one Autonomous VM cluster in an Exadata Infrastructure.
- Autonomous VM clusters on Oracle Exadata Cloud@Customer infrastructure: VM cluster is a set of symmetrical VMs across all Compute nodes. Autonomous



Container and Database run all the VMs across all nodes enabling high availability. It consumes all the resources of the underlying Exadata Infrastructure.

Before you can create any Autonomous Databases on your Oracle Exadata Cloud@Customer infrastructure, you must create an Autonomous VM cluster network, and you must associate it with a VM cluster.

- Autonomous Container Database: Provides a container for multiple Autonomous Databases.
- Autonomous Database: You can create multiple autonomous databases within the same autonomous container database. You can configure Oracle Autonomous Database for either transaction processing or data warehouse workloads.

Related Topics

Available Exadata Infrastructure Hardware Shapes

Deployment Order

You must create the dedicated Exadata infrastructure resources in the following order:

- **1. Exadata Infrastructure.** For more information, see *Preparing for Exadata Cloud@Customer* and *Provisioning Exadata Cloud@Customer System.*
- 2. Autonomous Exadata VM cluster. For more information, see Managing Autonomous Exadata VM Clusters.
- **3. Autonomous Container Database.** For more information, see *Managing Autonomous Container Databases.*
- 4. Autonomous Database. For more information, see Managing Autonomous Databases.

Related Topics

- Preparing for Exadata Database Service on Cloud@Customer Review OCI as well as the site, network and storage requirements to prepare and deploy Exadata Database Service on Cloud@Customer in your data center.
- Using the Console to Provision Exadata Database Service on Cloud@Customer Learn how to provision an Exadata Database Service on Cloud@Customer system.
- Managing Autonomous Exadata VM Clusters An Autonomous Exadata VM Cluster is a set of symmetrical VMs across all Compute nodes.
- Managing Autonomous Container Databases
 Learn how you can create, view, move, change backup policies, manage maintenance
 schedules, and perform other Oracle Autonomous Container Database management.
- Managing Autonomous Databases
 An Autonomous Database resource is a user database. When you create an
 Autonomous Database, you choose the Autonomous Container Database for it and you
 specify "Data Warehouse" or "Transaction Processing" as its workload type to create an
 Autonomous Data Warehouse database or an Autonomous Transaction Processing
 database.

User Roles

Your organization may choose to split the administration of the Oracle Autonomous Database on Oracle Exadata Database Service on Cloud@Customer into the following roles:



- Fleet Administrator. Fleet administrators create, monitor and manage Autonomous Exadata Infrastructure and Autonomous Container Database resources. They must also setup customer managed Backup Destinations, such as Recovery Appliance and NFS to be used by Autonomous Databases. A fleet administrator must have permissions for using the networking resources required by the Oracle Exadata Database Service on Cloud@Customer infrastructure, and permissions to manage the infrastructure and container database resources.
- Database Administrator. Database administrators create, monitor and manage Autonomous Databases. They also create and manage users within the database. Database administrators must have permissions for using container databases, for managing autonomous databases and backups, and for using the related networking resources. At the time of provisioning an Autonomous Database, the administrator provides user credentials for the automatically created ADMIN account, which provides administrative rights to the new database.
- **Database User.** Database users are the developers who write applications that connect to and use an Autonomous Database to store and access the data. Database users do not need Oracle Cloud Infrastructure accounts. They gain network connectivity to and connection authorization information for the database from the database administrator.

Available Exadata Infrastructure Hardware Shapes

Resource Limits

The following table lists the resource limits for Autonomous Database on Dedicated Exadata Infrastructure deployments on Oracle Public Cloud and Exadata Cloud@Customer.

Resource Limits (Maximum)

• The number of Provisioned CPUs determines the maximum number of Autonomous Databases that you can create. The smallest Autonomous Database can be created with 0.1 CPUs. Therefore the maximum number of Autonomous Databases that can be created is ten times the amount of Provisioned CPUs.

For example, an X9M Quarter Rack system with 252 Provisioned OCPUs allows a maximum of 2520 Autonomous Databases.

• The maximum number of Autonomous Container Databases per Autonomous Exadata VM Cluster (AVMC) is 16.

Recommended Resource Limits (Maximum)

Table 6-2 Recommended Resource Limits (Maximum)

Resource	Recommended Limit
Autonomous Databases per Autonomous Container Database	200
Autonomous Databases per Autonomous Container Database with Autonomous Data Guard Configured	25



Note:

It is possible to provision more Autonomous Databases than those mentioned in the above recommended limits table, especially with CPU overprovisioning. However, this implies compromising the Service Level Objectives (SLOs) to return an application online following an unplanned outage or a planned maintenance activity. To know the SLO details for Autonomous Database on Dedicated Exadata Infrastructure deployments, see Availability Service Level Objectives (SLOs).

Limits for Autonomous Exadata VM Clusters

You can create multiple Autonomous Exadata VM Clusters (AVMCs) on an Exadata Infrastructure resource. There are no hard limits on the number of AVMCs or Autonomous Container Databases (ACDs) you can provision on your Exadata Infrastructure. The AVMCs and ACDs have a minimum resource requirement, and you can create them as long as the minimum amount of resources is available.

To create an Autonomous Exadata VM Cluster, the minimum resources required per node are 5 OCPUs, 100 GB Memory, 150 GB Local Storage, and 3.25 TB Exadata Storage. Similarly, the minimum resources needed per node to create an ACD are 2 OCPUs and 50 GB Local Storage. As long as your Exadata Infrastructure has these minimum resources available, an AVMC and ACD can be created.

Exadata System Shapes

Autonomous Database on Dedicated Exadata Infrastructure is offered in the following Exadata System Shapes:

- **Base System:** Contains two database servers and three Oracle Exadata Storage Servers. A Base System is an entry-level configuration. Compared to other configurations, a Base System contains Oracle Exadata Storage Servers with significantly less storage capacity, and database servers with significantly less memory and processing power.
- **Quarter Rack:** Contains two database servers and three Oracle Exadata Storage Servers.
- Half Rack: Contains four database servers and six Oracle Exadata Storage Servers.
- Full Rack: Contains eight database servers and 12 Oracle Exadata Storage Servers.
- **X10M:** Starts with two standard memory (1390 GB) database servers and three Oracle Exadata Storage Servers, and can be expanded up to a maximum of 32 database servers and 64 storage servers.
- **X10M-Large:** Starts with two large memory (2090 GB) database servers and three Oracle Exadata Storage Servers, and can be expanded up to 32 large memory database servers and 64 storage servers.
- X10M-Extra Large: Starts with two large memory (2800 GB) database servers and three Oracle Exadata Storage Servers, and can be expanded up to 32 large memory database servers and 64 storage servers.

Each Exadata System Shape is equipped with a fixed amount of memory, storage, and network resources. All Shapes are based on Oracle Exadata X10M, X9M-2, X8M-2, X8-2, or X7-2 System Models.



Oracle Exadata X10M System Model Specifications

Review the technical specifications of X10M system model.

Specification	Q Quarter Rack X10M-L u (min config) a r t e r R a c k X 1 0 M (m i n c c o n f i g)	Quarter Rack X10M-XL (min config)
Minimum Number of X10M DB servers	2 2	2
Minimum Number of X10M Storage Servers	3 3	3
Total Number of Usable DB node OCPUs (assuming 2 cores per host)	3 380 8 0	380
Total Memory Available for VMs (GB)	2 4180 7 8 0	5600
Max /u02 capacity (GB) per VM	9 900 0 0	900
VM Image overhead per VM (GB)	1 184 8 4	184
Max local storage (/u02 + VM Image capacity) (GB) per infrastructure	2 2243 2 4 3	2243
Total number of VM clusters*	8 8	8
Minimum CPUs per VM	2 2	2
Minimum memory per VM (GB)	3 30 0	30

Table 6-3 Oracle Exadata Cloud@Customer X10M System Specifications



Specification	Q Quarter Rack X10M-L u (min config) a r t e r R a c k X 1 0 M (m i n c c o n f i g)	Quarter Rack X10M-XL (min config)
Min Local File System Storage per VM (GB)	6 60 0	60
Min Exadata Storage per VM Cluster	2 2 TB T B	2 TB
Total Usable ASM storage capacity (TB) per Cell	8 80 TB 0 T B	80 TB

Table 6-3 (Cont.) Oracle Exadata Cloud@Customer X10M System Specifications

Specification	Q Quarter Rack X10M-L u (min config) a r t e r R a c k X 1 0 M (m i n c c o n f f i g)	Quarter Rack X10M-XL (min config)
Exadata RMDA Memory (XRMEM)	1 1.25 TB per storage server 2 5 T B p e r s t t o r a g e s s t v e r v e r	1.25 TB per storage server

Table 6-3 (Cont.) Oracle Exadata Cloud@Customer X10M System Specifications

* A maximum of 16 clusters can be created on a single system.

Oracle Exadata X9M-2 System Model Specifications



Specification	E Exadata X9M-2 x Quarter Rack a d a t a X 9 M - 2 B a a s e R a c k	Exadata X9M-2 Half Rack	Exadata X9M-2 Full Rack
Number of Compute Nodes	22	4	8
Total Maximum Number of Enabled CPU Cores	4 124 8	248	496
Total RAM Capacity	6 2780 GB 5 6 G B	5560 GB	11120 GB
Total Persistent Memory Capacity	0 4.5 TB	9.0 TB	18.0 TB
Number of Exadata Storage Servers	33	6	12
Maximum Database Size, No Local Backup	5 152 TB 9 T B	305 TB	610 TB
Maximum Database Size, Local Backup (Exadata Cloud@Customer only)	2 76 TB 9 T B	152 TB	305 TB

Table 6-4 Oracle Exadata Cloud@Customer X9M-2 System Specifications

For more information, see Oracle Exadata Database Service on Exadata Cloud@Customer X9M datasheet.

Oracle Exadata X8M-2 System Model Specifications



Specification	X Exadata X8M-2 8 Quarter Rack M - 2 B a s e R a c k	Exadata X8M-2 Half Rack	Exadata X8M-2 Full Rack
Number of Compute Nodes	22	4	8
Total Maximum Number of Enabled CPU Cores	4 100 8	200	400
Total RAM Capacity	6 2780 GB 5 6 G B	5560 GB	11120 GB
Persistent Memory	0 4.5 TB	9.0 TB	18.0 TB
Number of Exadata Storage Servers	33	6	12
Maximum Database Size, No Local Backup	5 119 TB 9 T B	239 TB	479 TB
Maximum Database Size, Local Backup (Exadata Cloud@Customer only)	2 59 TB 9 T B	119 TB	239 TB

Table 6-5 Oracle Exadata Cloud@Customer X8M-2 System Specifications

Oracle Exadata X8-2 System Model Specifications



Specification	E Exadata X8-2 x Quarter Rack a d a t a X 8 - 2 B a s e R a s e R a k	Exadata X8-2 Half Rack	Exadata X8-2 Full Rack
Shape Name	E Exadata.Quarter3.1 x 00 a d a t a	Exadata.Half3.200	Exadata.Full3.400
	B a s e 3 4		
Number of Compute Nodes	8 22	4	8
Total Maximum Number of Enabled CPU Cores	4 100 8	200	400
Total RAM Capacity	7 1440 GB 2 0 G B	2880 GB	5760 GB
Number of Exadata Storage Servers	33	6	12
Maximum Database Size, No Local Backup	5 119 TB 9 T B	238 TB	476 TB
Maximum Database Size, Local Backup (Exadata Cloud@Customer only)	2 59 TB 9 T B	119 TB	239 TB

Table 6-6 Oracle Exadata Cloud@Customer X8-2 System Specifications



Oracle Exadata X7-2 System Model Specifications

Table 6-7	Oracle Exadata	Cloud@Customer X7-2 S	vstem Specifications

Specification	E Exadata X7-2 x Quarter Rack a d a t a X 7 - 2 B a s e R a s e R a k	Exadata X7-2 Half Rack	Exadata X7-2 Full Rack
Shape Name	E Exadata.Quarter2.9 x 2 a d a t a t B a s e 2 4	Exadata.Half2.184	Exadata.Full2.368
Number of Compute Nodes	22	4	8
Total Maximum Number of Enabled CPU Cores	4 92 4	184	368
Total RAM Capacity	4 1440 GB 8 0 G B	2880 GB	5760 GB
Number of Exadata Storage Servers	33	6	12
Maximum Database Size, No Local Backup	5 85.5 TB 9 T B	171.1 TB	342.1 TB

Specification	E Exadata X7-2 x Quarter Rack a d a t a X 7 - 2 B a s e R a c k	Exadata X7-2 Half Rack	Rack
Maximum Database Size, Local Backup (Exadata Cloud@Customer only)	2 42.8 TB 9 T B	85.5 TB	171.1 TB

Table 6-7 (Cont.) Oracle Exadata Cloud@Customer X7-2 System Specifications

Related Topics

- Availability Service Level Objectives (SLOs)
- Oracle Exadata Database Service on Exadata Cloud@Customer X9M

Managing Autonomous Exadata VM Clusters

An Autonomous Exadata VM Cluster is a set of symmetrical VMs across all Compute nodes.

Autonomous Container and Database run all the VMs across all nodes enabling high availability. It consumes all the resources of the underlying Exadata Infrastructure.

After you have created the Autonomous Exadata VM Cluster, you can create up to 12 Autonomous Container Database resources on it, depending on the capacity of your Exadata Infrastructure hardware, as described in *Available Exadata Infrastructure Hardware Shapes*.

- About Autonomous Exadata VM Clusters
- Resource Terminology
- Create an Autonomous Exadata VM Cluster
- View a List of DB Servers on an Exadata Infrastructure
- View a List of Autonomous Exadata VM Clusters
- View Details of an Autonomous Exadata VM Cluster
- Scale Autonomous Exadata VM Cluster Resources



- Schedule Oracle-Managed Infrastructure Updates Exadata Cloud Service updates are released on a quarterly basis. You can set a maintenance window to determine the time your quarterly infrastructure maintenance will begin.
- Change the License Type on an Autonomous VM Cluster
- Move an Autonomous Exadata VM Cluster to Another Compartment
- Terminate an Autonomous Exadata VM Cluster
- Using the API to Manage Autonomous Exadata VM Clusters
- Rotate Oracle Database TLS Certificate and Oracle REST Data Services (ORDS)
 TLS Certificate

Related Topics

• Available Exadata Infrastructure Hardware Shapes

About Autonomous Exadata VM Clusters

- Please contact your Oracle sales representative to understand the infrastructure upgrades required for supporting Oracle Autonomous Databases.
- Create multiple Autonomous Exadata VM Clusters on a single Exadata Infrastructure resource.
- Create both Autonomous Exadata VM Clusters and Exadata VM Clusters on the same Exadata Infrastructure.
- Support for multiple VM Clusters lets you:
 - Schedule separate maintenance runs for each Autonomous VM Cluster on the same Exadata Infrastructure.
 - Choose different license models for Autonomous Databases on the same Exadata Infrastructure.
 - Create and test Autonomous Data Guard between Autonomous Exadata VM Clusters on the same Exadata Infrastructure.
 - Customize compute, storage, and memory of each Autonomous Exadata VM Cluster configuration for the intended workload.

Resource Terminology

It is important to understand the various terms used with resource allocation and usage. So, let's look at the important terms you see on the Oracle Cloud Infrastructure (OCI) console and understand what they mean:

- Available CPUs: CPUs available for allocation to provision new Autonomous Databases or scale existing Autonomous Databases.
- Provisioned CPUs: Total CPUs allocated for all the Autonomous Database within the Autonomous VM Cluster.
- **Reserved CPUs**: Total CPUs reserved to support auto-scaling of Autonomous Databases, Autonomous Database failover on node failure, and lifecycle management of empty Autonomous Container Databases.



- Reclaimable CPUs: Total CPUs from all terminated and scaled-down Autonomous Databases in all the Autonomous Container Databases within the Autonomous VM Cluster.Reclaimable CPUs are not returned to the Available state until Autonomous Container Database is restarted.
- **Provisionable ACDs:** Number of Autonomous Container Databases that can be created in the Autonomous VM Cluster.
- **Provisioned ACDs:** Number of Autonomous Container Databases that have been created in the Autonomous VM Cluster.
- **Non-provisionable ACDs:** Number of Autonomous Container Databases that cannot be created because of a shortage of available CPUs in the Autonomous VM Cluster.
- **Total Storage (in TBs):** Total storage allocated to the AVMC.
- Available Storage (in TBs): Storage available for Autonomous Databases' use in this AVMC.
- Used Storage (in TBs): Storage currently used by Autonomous Database(s) in this AVMC.
- Memory per CPU (in GBs): Memory allocated to the AVMC per CPU.

Create an Autonomous Exadata VM Cluster

Follow these steps to create an Autonomous Exadata VM cluster on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Exadata VM Clusters.
- 3. Click Create Autonomous Exadata VM Cluster.
- 4. In the Create Autonomous Exadata VM Cluster dialog, enter the following general information:
 - **Compartment**: Specify the compartment in which the Autonomous Exadata VM Cluster will be created.
 - **Display Name**: A user-friendly description or other information that helps you easily identify the infrastructure resource. The display name does not have to be unique. Avoid entering confidential information.
 - **Exadata Infrastructure**: Select an Exadata Infrastructure.
 - VM Cluster Network: Select a VM Cluster Network.

Configure Autonomous VM Cluster Resources All DB Servers that have the minimum resources to create an Autonomous VM Cluster are selected by default.

- a. Click Edit DB Server selection for VM placement to allocate VM resources.
- **b.** In the resulting Add Virtual Machines dialog, select a minimum of two database servers for VM placement. Maximum resources available for allocation per VM are based on the number of database servers selected.



Note:

If the DB server is not added to the VM Cluster Network resources, then that DB server cannot be selected.

c. Click Save Changes.

Note:

The minimum and maximum values for these parameters change in relation to each other, for example, the OCPU count allocation will impact the number of ACDs you can create.

Compute Model: Choose a compute model for your Autonomous Exadata VM Cluster resource.

The default model is ECPU. ECPU is based on the number of cores elastically allocated from a pool of compute and storage servers.

Click **Change compute model** if you wish to select OCPU. OCPU compute model is based on the physical core of a processor with hyper-threading enabled.

Note:

The compute model chosen here will apply to all the Autonomous Container Databases and Autonomous Database instances created in this Autonomous Exadata VM Cluster resource.

- VM Count: Number of database servers selected for the Autonomous Exadata VM Cluster. You cannot modify the VM count.
- Maximum number of Autonomous Container Databases for the Autonomous VM Cluster: The number of ACDs specified represents the upper limit on ACDs. These ACDs must be created separately as needed. ACD creation also requires 2 available OCPUs or 8 available ECPUs per node.
- CPU count per VM: Specify the CPU count for each individual VM. The minimum value is 5 OCPUs or 20 ECPUs per VM.
- Database memory per CPU (GB): The memory per OCPU allocated for the Autonomous Databases in the Autonomous Exadata VM Cluster.
- Allocate Storage for Local Backups: Check this option to configure the Exadata storage to enable local database backups.
- Autonomous Database storage for the Autonomous VM Cluster (TB): Data storage allocated for Autonomous Database creation in the Autonomous VM Cluster.
- 5. Configure automatic maintenance.
 - a. Click Edit Maintenance Preferences.



On the Edit Maintenance Preferences page do the following:

- **No preference:** The system assigns a date and start time for infrastructure maintenance.
- **Specify a schedule:** Choose your preferred month, week, weekday, start time, and lead time for infrastructure maintenance.
 - Under Maintenance months, specify at least one month for each quarter during which Exadata infrastructure maintenance will take place. You can select more than one month per quarter. If you specify a long lead time for advanced notification (for example, 4 weeks), you may wish to specify 2 or 3 months per quarter during which maintenance runs can occur. This will ensure that your maintenance updates are applied in a timely manner after accounting for your required lead time. Lead time is discussed in the following steps.
 - 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 Start hour, 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 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.
- b. Click Save Changes.
- 6. Choose the license type you wish to use.

Your choice affects metering for billing. You have the following options:

- **Bring your own license**: If you choose this option, make sure you have proper entitlements to use for new service instances that you create.
- License included: With this choice, the cost of the cloud service includes a license for the Database service.
- 7. The following Advanced Options are available:
 - **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 browserdetected 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**.

• **Listener**: VM Cluster Networks are created with the default ports (non-TLS: 1521 and TLS: 2484). You can select a non-default SCAN listener port for both TLS and non-TLS connections within the permissible range of ports.

Note:

- You can configure SCAN listener ports and TLS/mTLS authentication mode only when you provision a new VM Cluster.
- You cannot change the settings after provisioning the VM Cluster.
- Once configured, the configuration applies to all ACDs in the cluster.
- Non-TLS port:
 - * Default: 1521
 - * **Permissible range:**1024 8999
 - * Exceptions:
 - * Agent port: 7070
 - * Admin port: 7879
 - * Connect to agent port: 7060
 - * Oracle notification service: 6100, 6200
- TLS port:
 - * Default: 2484
 - * Permissible range:1024 8999
 - * Exceptions:
 - * Agent port: 7070
 - * Admin port: 7879
 - * Connect to agent port: 7060
 - * Oracle notification service: 6100, 6200
- Enable mutual TLS (mTLS) authentication: Select or deselect this check box to choose between one-way TLS and mutual TLS for database SSL certificates.



- **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.
- 8. 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.
- 9. Click Create Autonomous Exadata VM Cluster.

The state of the Autonomous VM Cluster changes to **Provisioning**. Likewise, the newly added Autonomous Virtual Machines will be in the **Provisioning** state. Upon successful completion of the operation, the state of the Autonomous Virtual Machine Cluster and the Autonomous Virtual Machines will change to **Available**.

Related Topics

- Resource Tags
- Resource Manager and Terraform

View a List of DB Servers on an Exadata Infrastructure

Follow these steps to view a list of database server hosts on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle 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.

View a List of Autonomous Exadata VM Clusters

Follow these steps to view a list of autonomous Exadata VM clusters on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Exadata VM Clusters.

Autonomous Exadata VM Clusters view page lists AVMCs in the chosen Exadata Infrastructure and displays details that include:

- Name: Name of AVMC
- **State:** Lifecycle state of AVMC such as Updating, Provisioning, Available, and so on.
- **CPU(%):** Available and total CPUs and percentage of used CPUs represented as a color-coded bar chart. The colors on this bar represent the following:
 - Less than 70%: Green
 - Between 70 to 90%: Yellow
 - Greater than 90%: Red
- **Reclaimable CPUs**: Total CPUs from all terminated and scaled-down Autonomous Databases in all the Autonomous Container Databases within the Autonomous VM Cluster. Reclaimable CPUs are not returned to the Available state until Autonomous Container Database is restarted.
- Storage (TB) (%): Available and total storage in TB and the percentage of used storage are represented as a color-coded bar chart. The colors on this bar represent the following:
 - Less than 70%: Green
 - Between 70 to 90%: Yellow
 - Greater than 90%: Red
- **Provisionable ACDs:** The number of Autonomous Container Databases that can be created in the Autonomous VM Cluster.
- Memory per CPU: Represents the size of memory in GB per CPU.
- **Created:** Date and time of creation of AVMC.

View Details of an Autonomous Exadata VM Cluster

Follow these steps to view detailed information about an autonomous Exadata VM cluster on an Oracle Exadata Database Service on Cloud@Customer system.

1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.



- 2. Click Autonomous Exadata VM Clusters.
- In the list of Autonomous Exadata VM Clusters, click the display name of the Exadata VM cluster you wish to view details.

(or)

- a. Click Exadata Infrastructure.
- In the list of Exadata Infrastructure, click the display name of the Exadata Infrastructure you wish to view details.
 Infrastructure Details page is displayed.
- c. Under Resources, click Autonomous Exadata VM Clusters.
- d. In the list of Autonomous Exadata VM Clusters, click the display name of the Autonomous Exadata VM Clusters you wish to view details. Or, click the action menu (three dots), and then select View Details. Autonomous Exadata VM Clusters Details page is displayed.

The **Resource allocation** section provides an overview of the resources allocated.

- **CPU(%):** Available and total CPUs and percentage of used CPUs represented as a color-coded bar chart. The colors on this bar represent the following:
 - Less than 70%: Green
 - Between 70 to 90%: Yellow
 - Greater than 90%: Red
- ADB Storage (TB) (%): Available and total storage in TB and the percentage of used storage are represented as a color-coded bar chart. The colors on this bar represent the following:
 - Less than 70%: Green
 - Between 70 to 90%: Yellow
 - Greater than 90%: Red
- Database memory per CPU: This represents the size of memory in GB per CPU.
- **Provisionable ACDs:** The number of Autonomous Container Databases that can be created in the Autonomous VM Cluster.

Click View details to view resource allocation details.

The resulting **Resource allocation details** page has two tabs: **Autonomous Exadata VM Cluster** to view details of the resources allocated to AVMC and **Autonomous Container Database** to view details of the resources allocated to ADB.

Autonomous Exadata VM Cluster:

- **Total resources allocated:** This section lists the latest values for the following resources allocated to this AVMC:
 - CPUs
 - Exadata Storage in TB
 - Local Storage in GB
 - Memory in GB
 - Maximum number of ACDs



- Autonomous Database Storage in TB
- Autonomous Database Memory per CPU in GB.
- You can also see if the Storage for local backups is enabled or disabled.
- **Resource usage visualization:** This section has graphical and tabular representations of AVMC's resource usage.

Note: You can choose to see this information either in the graphical or tabular view by selecting Chart view or Table view from the drop-down list at the top-right corner of this section.

- Chart View: Chart view is the default view. In this view, you can see 4 graphical visualizations that provide usage details for different resources, as:
 - * **CPU usage:** Depicts the Total number of CPUs allocated to this AVMC and how many of those CPUs are reclaimable, available, provisioned, and reserved. This is a doughnut chart with the total number of CPUs shown in the center of the chart.
 - * CPU usage at VM level: This is a horizontal bar graph giving a breakdown of the CPU usage for each VM in the cluster. Each bar shows the number of reclaimable, available, provisioned, and reserved CPUs for that VM with color coding. Hovering on each colored part of the horizontal bar displays the number of reclaimable, available, provisioned, and reserved CPUs for that specific VM. Clicking the reclaimable, provisioned, and reserved bars will open a new panel with the breakdown of those CPU components by ACDs.
 - * Autonomous Container Database (ACD) usage: Depicts the Total number of ACDs that can be created in this AVMC along with a break-up of Provisionable ACDs, Provisioned ACDs, and Nonprovisionable ACDs. See Resource Terminology to understand what each of these means. This is a doughnut chart with the total number of ACDs shown in the center of the chart.
 - * Autonomous Database (ADB) storage (in TBs) usage: This is a doughnut chart depicting the available, used, and total Autonomous Database storage in TBs. The total storage value is shown in the center of the chart with available and used storage values shown on the chart in different colors.
- Table View: To see the resource usage details of an AVMC in the table view, select Table View from the drop-down list at the top-right corner of the Resource usage visualizations section. The table view shows the exact same details as the chart view, in the form of tables. The four tables you can see are:
 - * **CPU usage:** Lists the number of total, available, provisioned, reserved, and reclaimable CPUs in this AVMC.
 - CPU usage at VM level: Lists the number of available, provisioned, reserved, and reclaimable CPUs for each VM in this VM cluster.
 - * Autonomous Container Database (ACD) usage: Lists the number of provisionable, provisioned, and non-provisionable ACDs in this AVMC.



* Autonomous Database (ADB) storage (in TBs) usage: Shows the available and used Autonomous Database storage in TBs.

Autonomous Container Database: This tab lists the following details for all the ACDs, in the selected AVMC, created in any compartment in the tenancy:

- ACD's display name
- CPU value of the largest provisionable Autonomous Database.
- The number of CPUs provisioned to Autonomous Databases.
- The number of reserved CPUs.
- The number of reclaimable CPUs.

Scale Autonomous Exadata VM Cluster Resources

Follow these steps to scale Autonomous Exadata VM Cluster resources.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Exadata VM Clusters.
- 3. In the list of Autonomous Exadata VM Cluster, click the Autonomous Exadata VM Cluster that you want to scale.
- 4. On the Autonomous Exadata VM Clusters details page, click Scale Autonomous VM Cluster.
- 5. In the resulting Scale Autonomous VM Cluster panel, adjust the sliders to increase or decrease the following resources:
 - CPU count per VM
 - Maximum number of Autonomous Container Databases
 - Database storage (TB)

The minimum and maximum values on the sliders are the smallest and largest values the resources can scale to.

Note:

Modifying CPU count per VM or the maximum number of Autonomous Container Databases of an AVMC triggers a rolling restart of the AVMC. This results in restarting all the ACDs and Autonomous Databases created in that AVMC.

- 6. Click Save Changes.
- 7. To confirm a rolling restart, enter the AVMC name in the **Confirm rolling restart** dialog, and click **Confirm**.



Note:

In case of an ongoing maintenance activity on this AVMC or the ACDs and Autonomous Databases within it, your scale request fails with an appropriate message.

Schedule Oracle-Managed Infrastructure Updates

Exadata Cloud Service updates are released on a quarterly basis. You can set a maintenance window to determine the time your quarterly infrastructure maintenance will begin.

You can also view scheduled maintenance runs and the maintenance history of your Autonomous Exadata VM Cluster in the Oracle Cloud Infrastructure Console.

- Set the Automatic Maintenance Schedule for Autonomous Exadata VM Cluster
- View or Edit the Time of the Next Scheduled Maintenance for Autonomous Exadata VM Cluster
- View the Maintenance History of Autonomous Exadata VM Cluster

Set the Automatic Maintenance Schedule for Autonomous Exadata VM Cluster

Learn how to set the maintenance schedule for an Autonomous Exadata VM Cluster on an Oracle Exadata Cloud@Customer system.

- Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Exadata VM Clusters.
- 3. In the list of Autonomous Exadata VM Clusters, find the Autonomous Exadata VM Cluster you want to set the maintenance window for and click its highlighted name.
- 4. On the Autonomous Exadata VM Cluster Details page, under Maintenance, click the edit link in the Maintenance Details field.
- 5. In the Edit Automatic Maintenance page, select Specify a schedule.
- 6. Under Maintenance months, specify at least one month for each quarter during which Autonomous Exadata VM Cluster maintenance will take place.

You can select more than one month per quarter. If you specify a long lead time for advanced notification (for example, 4 weeks), you may wish to specify 2 or 3 months per quarter during which maintenance runs can occur. This will ensure that your maintenance updates are applied in a timely manner after accounting for your required lead time. Lead time is discussed in the following steps.

7. *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.



8. *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.

- 9. Optional. **Under Start hour**, 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.
- **10. Under 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.

11. Click **Save Changes**.

View or Edit the Time of the Next Scheduled Maintenance for Autonomous Exadata VM Cluster

Learn how to view and edit the time of the next scheduled maintenance.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Exadata VM Clusters.
- In the list of Autonomous Exadata VM Clusters, find the Autonomous Exadata VM Cluster you want to set the maintenance window for and click its highlighted name.
- 4. On the Autonomous Exadata VM Cluster details page, under Maintenance, click the view link in the Next Maintenance field.
- 5. On the Maintenance page, scheduled maintenance events are listed.
- 6. *Optional*. To change the time of the next scheduled maintenance, click the **Edit** link in the **Scheduled Start Time** field.
- 7. In the Edit Infrastructure Maintenance Scheduled Start Time page, enter a date and time in the Scheduled Start time field.

The following restrictions apply:

- You can reschedule the infrastructure maintenance to a date no more than 180 days from the prior infrastructure maintenance. 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. You cannot reschedule the maintenance if the current time is within 2 hours of the scheduled maintenance start time.
- Oracle reserves certain dates each quarter for internal maintenance operations, and you cannot schedule your maintenance on these dates.

View the Maintenance History of Autonomous Exadata VM Cluster

Learn how to view the maintenance history for an Autonomous Exadata VM Cluster.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Exadata VM Clusters.



- 3. In the list of Autonomous Exadata VM Clusters, find the Autonomous Exadata VM Cluster you want to set the maintenance window for and click its highlighted name.
- 4. On the Autonomous Exadata VM Cluster details page, under Maintenance, click the view link in the Next Maintenance field.
- 5. Click **Maintenance History** to see a list of past maintenance events including details on their completion state.

Change the License Type on an Autonomous VM Cluster

Follow these steps to update the license type of an autonomous Exadata VM cluster on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Exadata VM Clusters.
- 3. In the list of Autonomous Exadata VM Clusters, click the display name of the Exadata VM cluster you wish to administer.
- 4. Click Update License Type.
- 5. On the Update License Type dialog box, choose one of the following license types.
 - 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 interrupt the operation of the VM cluster.

6. Click Save Changes.

Move an Autonomous Exadata VM Cluster to Another Compartment

Follow these steps to move an autonomous Exadata VM cluster on an Oracle Exadata Database Service on Cloud@Customer system from one compartment to another compartment.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Exadata VM Clusters.
- 3. In the list of Autonomous Exadata VM Clusters, click the display name of the Exadata VM cluster you wish to administer.
- 4. Click Move Resource.
- 5. Select the new compartment.
- 6. Click Move Resource.

Terminate an Autonomous Exadata VM Cluster

Follow these steps to terminate an autonomous Exadata VM cluster on an Oracle Exadata Database Service on Cloud@Customer system.



- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. In the list of Autonomous Exadata VM Clusters, click the display name of the Exadata VM cluster you wish to administer.
- 3. Click Terminate.
- 4. Confirm that you wish to terminate your Autonomous Exadata VM Cluster in the confirmation dialog.
- 5. Click Terminate VM Cluster.

Using the API to Manage Autonomous Exadata 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.

The following table lists the REST API endpoints to manage Autonomous Exadata VM Clusters.

Operation	REST API Endpoint
Create an Autonomous Exadata VM Cluster	CreateAutonomousVmCluster
View a list of Autonomous Exadata VM Clusters	ListAutonomousVmClusters
View details of an Autonomous Exadata VM Cluster	GetAutonomousVmCluster
Change the license type of an Autonomous VM Cluster	UpdateAutonomousVmCluster
Move an Autonomous Exadata VM Cluster to another compartment	ChangeAutonomousVmClusterCompartment
Terminate an Autonomous Exadata VM Cluster	DeleteAutonomousVmCluster
Get the resource usage statistics of an Autonomous Container Database	GetAutonomousContainerDatabaseResourceUsa ge

Rotate Oracle Database TLS Certificate and Oracle REST Data Services (ORDS) TLS Certificate

To rotate the database TLS certificate or ORDS TLS certificate, use this procedure.

Note:

Rotating the Database TLS certificate or ORDS TLS certificate is a disruptive operation. When a database TLS certificate is rotated, the listener is restarted, disrupting database availability. During ORDS TLS certificate rotation, ORDS restarts, disrupting application connectivity. It is recommended that you rotate your certificates at an appropriate time.

1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.



- 2. Click Autonomous Exadata VM Clusters.
- 3. In the list of Autonomous Exadata VM Clusters, click the display name of the Exadata VM cluster for which you want to rotate certificates.

The Network section provides an overview of the SCAN listener ports (**TLS**, **non-TLS**), authentication mode (**One way TLS**, **Mutual TLS**), and the expiry dates for the Database and ORDS TLS certificates.

Note:

- Six weeks before certificates expire, a warning banner displays a list of certificates to rotate. Until the certificates are refreshed, the banner will continue to display.
- The system displays an error banner if the certificate has already expired, indicating that the database on this VM cluster cannot be accessed. Also, it suggests that you can access the database using 1521, a non-TLS port.
- In the Autonomous Exadata VM Clusters Details page, click Manage Certificates. Manage Certificates page is displayed.

Note:

- Six weeks before certificates expire, a warning banner displays a list of certificates to rotate. Until the certificates are refreshed, the banner will continue to display.
- The system displays an error banner if the certificate has already expired, indicating that the database on this VM cluster cannot be accessed. Also, it suggests that you can access the database using 1521, a non-TLS port.
- 5. Select a certificate type to manage.

Database TLS Certificate: Select this option to manage Autonomous Database client connections.

A warning banner is displayed six weeks prior to the certificate expires.

- a. Select a certificate generation type.
 - **System generated**: Select this option if you want to use Oracle provided certificate.
 - **Bring your own certificate**: Select this option to choose your own certificate.
 - Certificate source: Defaults to VM Cluster managed certificate. You cannot edit this field.
 - **Certificate**: Select a certificate from a compartment of your choice.
 - **Specify CA certificate**: Select this check box to provide CA details.



- * **Certificate Authority**: Select a CA from the compartment of your choice.
- * **CA Bundle**: Select a CA bundle from the compartment of your choice.
- b. Click Save changes.

Confirm Database TLS certificate update dialog is displayed.

- c. Review the information on the banner.
- d. Click Update DB TLS certificate.

The status of Autonomous Exadata VM Cluster changes to **Updating** and the status changes to **Available** after updating the changes successfully.

ORDS TLS Certificate: Select this option to rotate the TLS certificate for APEX application.

A warning banner is displayed six weeks prior to the certificate expires.

- a. Select a certificate generation type.
 - **System generated**: Select this option if you want to use Oracle provided certificate.
 - Bring your own certificate: Select this option to choose your own certificate.
 - **Certificate source**: You cannot edit this field.
 - **Certificate**: Select a certificate from a compartment of your choice.
 - Specify CA certificate: Select this check box to provide CA details.
 - * **Certificate Authority**: Select a CA from the compartment of your choice.
 - * **CA Bundle**: Select a CA bundle from the compartment of your choice.

b. Click Save changes.

Confirm ORDS TLS certificate update dialog is displayed.

c. Review the information on the banner.

d. Click Update ORDS TLS certificate.

The status of Autonomous Exadata VM Cluster changes to **Updating** and the status changes to **Available** after updating the changes successfully.

Related Topics

- Resource Tags
- Resource Manager and Terraform

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 autonomous databases on Exadata Cloud@Customer:

- 1. In the Guest VM on the Exadata Infrastructure.
- 2. On an external key management device. Oracle Key Vault is the currently supported device.
- 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 vaults from one compartment to another.
- Setting Up Your Exadata Cloud@Customer to Work With Oracle Key Vault
- Managing Your Key Store

About Oracle Key Vault

Oracle Key Vault is a full-stack, security-hardened software appliance built to centralize the management of keys and security objects within the enterprise.

Note:

The Oracle Key Vault is a customer-provisioned and managed system and it is not part of Oracle Cloud Infrastructure managed services.

Related Topics

• Oracle Key Vault

Overview of Key Store

Integrate your on-premises Oracle Key Vault (OKV) with customer-managed database cloud services to secure your critical data on-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 Autonomous Databases 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.

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 vaults from one compartment to another.

After you move a vault to a new compartment, inherent policies apply immediately and affect access to the vault. Moving a vault doesn't affect access to any keys or secrets that the vault contains. You can move a key or secret from one compartment to another independently of moving the vault it's associated with. For more information, see "Managing Compartments".

Related Topics

• Managing Compartments

Setting Up Your Exadata 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 to access the OKV server.
- 2. Ensure that the REST interface is enabled from the OKV user interface.
- Create "OKV REST Administrator" user. You can use any qualified username of your choice, for example, "okv_rest_user". For ADB-C@C and ExaDB-C@C, use the same or different REST users. Those databases can be key-managed in the same or different on-prem OKV clusters. ExaDB-C@C needs REST user with create endpoint privilege. ADB-C@C needs REST user with create endpoint and create endpoint group privileges.
- Gather OKV administrator credentials and IP address, which is required to connect to OKV.
- 5. Open the ports 443, 5695, and 5696 for egress on the client network to access the OKV server.

For more information, see *Network Port Requirements*, *Managing Oracle Key Vault Users*, and *Managing Administrative Roles and User Privileges*

- Step 1: Create a Vault in OCI Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password
- Step 2: Create a Dynamic Group and a Policy Statement for Key Store to Access Secret in OCI Vault
- Step 3: Create a Dynamic Group and a Policy Statement for Exadata Infrastructure to Key Store
- Step 4: Create a Policy Statement for Database Service to Use Secret from OCI Vault Service
- Step 5: 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 Dynamic Group and a Policy Statement for Key Store to Access Secret in OCI Vault

To grant your Key Store resources permission to access Secret in OCI Vault, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the Secret you created in the OCI Vaults and Secrets.

When defining the dynamic group, you identify your Key Store resources by specifying the OCID of the compartment containing your Key Store.

1. Copy the OCID of the compartment containing your Key Store resource.

You can find this OCID on the Compartment Details page of the compartment.

2. Create a dynamic group by following the instructions in "To create a dynamic group" in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

ALL {resource.compartment.id ='<compartment-ocid>'}

where **<compartment-ocid>** is the OCID of the compartment containing your Key Store resource.

3. After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your vaults and secrets. Then, add a policy statement of this format:

allow dynamic-group **<dynamic-group>** to use secret-family in compartment **<vaults-and-secrets-compartment>**



where *<dynamic-group>* is the name of the dynamic group you created and *<vaults-and-secrets-compartment>* is the name of the compartment in which you created your vaults and secrets.

Related Topics

• To create a dynamic group

Step 3: Create a Dynamic Group and a Policy Statement for Exadata Infrastructure to Key Store

To grant your Exadata infrastructure resources permission to access Key Store, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the Key Store you created.

When defining the dynamic group, you identify your Exadata infrastructure resources by specifying the OCID of the compartment containing your Exadata infrastructure.

1. Copy the OCID of the compartment containing your Exadata infrastructure resource.

You can find this OCID on the Compartment Details page of the compartment.

2. Create a dynamic group by following the instructions in "To create a dynamic group" in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

ALL {resource.compartment.id = '<compartment-ocid>'}

where **<compartment-ocid>** is the OCID of the compartment containing your Exadata infrastructure resource.

3. After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your Key Store. Then, add a policy statement of this format:

Allow dynamic-group **<dynamic-group>** to use keystores in compartment **<key-store-compartment>**

where *<dynamic-group>* is the name of the dynamic group you created and *<key-store-compartment>* is the name of the compartment in which you created your Key Store.

Step 4: Create a Policy Statement for Database Service to Use Secret from OCI Vault Service

To grant the Autonomous Database service permission to use the secret in OCI Vault to log in to the OKV REST interface, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your OCI Vaults and Secrets. Then, add a policy statement of this format:

allow service database to read secret-family in compartment <vaultsand-secrets-compartment>



where <*vaults-and-secrets-compartment>* is the name of the compartment in which you created your OCI Vaults and Secrets.

Once the OCI Vault is set up and the IAM configuration is in place, you are now ready to deploy your Oracle Key Vault 'Key Store' in OCI and associate it with your Exadata Cloud@Customer VM Cluster.

Step 5: Create Key Store

Follow these steps to create a Key Store to connect to an on-premises encryption key appliance such as Oracle Key Vault (OKV).

- 1. Open the navigation menu. Under Oracle Database, click Exadata 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 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 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 Key Store.
- 7. Ensure that you use the same "*okv_rest_user*" user credentials, while provisioning Autonomous Database.

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 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.
- 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 Database, click Exadata 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.

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 Database, click Exadata 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 Cloud@Customer system from one compartment to another compartment.

- 1. Open the navigation menu. Under **Oracle Database**, click **Exadata 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 Move Resource.
- 6. On the Move Resource to a Different Compartment page, select the new compartment.
- 7. Click Move Resource.

Delete a Key Store

You can delete a Key Store only if it is not associated with any CDBs.

- 1. Open the navigation menu. Under Oracle Database, click Exadata 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 Delete.
- 6. On the Delete Key Store dialog, click Delete.

View Key Store Associated Container Database Details

Follow these steps to view details of the container database associated with a Key Store.

- 1. Open the navigation menu. Under Oracle 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 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
Choose between customer-managed and Oracle-managed encryption	CreateAutonomousContainerDatabase
Get the Key Store (OKV or Oracle-managed) and OKV wallet name	GetAutonomousContainerDatabase
Rotate OKV and Oracle-managed key	RotateAutonomousContainerDatabaseKey
Get the Key store (OKV or Oracle-managed) and OKV wallet name	GetAutonomousDatabase
Rotate OKV and Oracle-managed key	RotateAutonomousDatabaseKey
Get the Key Store (OKV or Oracle-managed) and OKV wallet name	GetAutonomousDatabaseBackup

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface

Managing Autonomous Container Databases

Learn how you can create, view, move, change backup policies, manage maintenance schedules, and perform other Oracle Autonomous Container Database management.

An Autonomous Container Database resource provides a container for your Autonomous Databases. You can create multiple Autonomous Container Database resources in a single Autonomous Exadata VM Cluster resource, but you must create at least one before you can create any Autonomous Databases.

- Create an Autonomous Container Database
- View a List of Autonomous Container Databases
- View Details of an Autonomous Container Database
- Schedule a Quarterly Maintenance Update
- Rotate CDB Encryption Key
- Change the Backup Retention Policy of an Autonomous Container Database
- Change the Maintenance Schedule of an Autonomous Container Database
- Restart an Autonomous Container Database
- Move an Autonomous Container Database to Another Compartment
- Terminate an Autonomous Container Database



Using the API to Manage Autonomous Container Databases

Create an Autonomous Container Database

Follow these steps to create an autonomous container database on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. Click Create Autonomous Container Database.

The Create Autonomous Container Database page is displayed.

- 4. Provide the following basic information:
 - **Compartment**: Choose the compartment in which your autonomous container database will be created.
 - **Display Name**: Enter a user-friendly description or other information that helps you easily identify the autonomous container database. The display name does not have to be unique. Avoid entering confidential information.
- 5. Select the Autonomous Exadata VM Cluster you wish to use to create your autonomous container database.
- 6. Optionally, you can configure an automatic maintenance schedule.
 - a. Click Modify Maintenance.
 - b. In the resulting dialog enter the following details:
 - **Configure maintenance method**: Choose between Rolling or Non-rolling maintenance methods. By default, it's always Rolling with Time Zone File update disabled.
 - Rolling: The Rolling method updates the ACD, one node at a time, with no downtime for the Autonomous Databases.
 - Non-rolling: The Non-rolling method shuts down and updates the ACD in parallel across all nodes. This method minimizes maintenance time but it requires full downtime for the ACD and all associated Autonomous Databases.

Optionally, you can also select **Enable time-zone update**. Updating the time-zone file would require complete downtime for the ACD and the associated Autonomous Databases. The downtime is dependent on the amount of data that is time-zone sensitive.

Note:

All the RUs that include time-zone updates will be patched using the non-rolling maintenance method (with full system downtime) only. So, if you configure your maintenance to the rolling method and enable time-zone update, all the RUs that include a time-zone update are applied in the non-rolling method only. Only those RUs without a time-zone update are applied in the rolling method.

• Configure container database maintenance version: Choose Next release update (RU) or Latest release update (RU) to configure the next maintenance

version of the Autonomous Container Database. Accordingly, the autonomous container database will be updated to the next immediate release or the latest release in the upcoming maintenance window.

- **Configure automatic maintenance schedule:** You can change the maintenance schedule settings, as desired:
 - Select No preference to permit Oracle to schedule maintenance as needed based on the maintenance schedule of the Exadata Infrastructure resource hosting the Autonomous Container Database.
 - Select Custom schedule to restrict when Oracle can schedule maintenance. You can choose your preferred month, week, weekday, start time, and buffer period between primary and standby maintenance execution for the Autonomous Container Database maintenance.
- To configure the Custome schedule, select **Specify a schedule**. Choose your preferred month, week, weekday, and start time for autonomous container database maintenance.
 - 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.
 - Under Day of the week, specify the day of the week on which the maintenance will occur.
 - Under **Hour of the day**, specify the hour during which the maintenance run will begin.
- c. Click Save Changes.
- 7. Select a Backup Destination Type:
 - **Object Storage**: Stores backups in an Oracle-managed object storage container on Oracle Cloud Infrastructure. Optionally, you can use this field to specify your corporate HTTP proxy. Oracle recommends using an HTTP proxy when possible for enhanced security.
 - **Network File System (NFS)**: Stores backups in one of your previously defined backup destinations that uses Network File System (NFS) storage.
 - Select a Backup Destination.
 - Recovery Appliance: Stores backups in one of your previously defined backup destinations that uses Oracle Zero Data Loss Recovery Appliance (ZDLRA).
 - Select a Backup Destination.
 - Provide a unique name for the database.
 - Provide a VPC user name and password.

To use ZDLRA as the Backup Destination, before creating the ACD, add the database to the Recovery Appliance, and assign a protection policy to the database.



a. 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.

- b. Create a VPC user on the Recovery Appliance or use an existing user. For more information about creating VPC user accounts, see *Creating Virtual Private Catalog Accounts*.
- c. Add the database to the Recovery Appliance using the intended unique name of the database (db unique name) and intended protection policy.

EXEC dbms_ra.add_db(db_unique_name=>'<DB_UNIQUE_NAME of the database to add>', protection_policy_name=>'<name of the protection policy to assign to the database>', reserved_space=>'<guaranteed disk space available for the protected database>');

For example:

```
EXEC dbms_ra.add_db(db_unique_name=>'ACDNAME',
protection policy name=>'BRONZE', reserved space=>'1024G');
```

For more information about enrolling the protected database, see *Enrolling the Protected Database with Recovery Appliance (Command Line)*.

d. Grant access to this database to the Recovery Appliance user account.

```
EXEC dbms_ra.grant_db_access(db_unique_name=>'<DB_UNIQUE_NAME of the
database>', username=>'<Recovery Appliance user account>');
```

For example:

```
EXEC dbms_ra.grant_db_access(db_unique_name=>'ACDNAME',
username=>'RAVPC1');
```

- 8. The following Advanced Options are available:
 - Backup retention period: Customize the retention period for automatic backups. For Recovery Appliance, you cannot select the retention period.
 - Encryption Key: Choose an encryption option, Encrypt using Oracle-managed keys or Encrypt using customer-managed keys. The default option is Oraclemanaged keys.

To use customer-managed keys, select the **Encrypt using customer-managed keys** option, 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.



- Autonomous Container Databases and Autonomous Databases 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 ACD. 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 remains in OKV and will not be deleted.
- **Tags**: Optionally, you can 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, skip this option (you can apply tags later) or ask your administrator. Avoid entering confidential information.
- 9. 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.
 - 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.



10. Click Create Autonomous Container Database.

Related Topics

- Creating Virtual Private Catalog Accounts
- Enrolling the Protected Database with Recovery Appliance (Command Line)
- Resource Tags
- Resource Manager and Terraform

View a List of Autonomous Container Databases

There are two ways to view a list of autonomous container databases on an Oracle Exadata Database Service on Cloud@Customer system.

- View the List of Autonomous Container Databases in an Autonomous Exadata VM Cluster
- View the List of Autonomous Container Databases in a Compartment

View the List of Autonomous Container Databases in an Autonomous Exadata VM Cluster

Follow these steps to view a list of an autonomous container databases in a given autonomous Exadata VM cluster on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Exadata VM Clusters.
- 3. Click the display name of the Autonomous Exadata VM Cluster that you interested in.

In the Autonomous Exadata VM Clusters Details page, a list of Autonomous Container Databases is displayed under **Resources**.

You can also filter out to view Autonomous Container Databases in a particular compartment. Under **List Scope**, select a compartment from the **Compartment** drop-down list.

View the List of Autonomous Container Databases in a Compartment

Follow these steps to view a list of an autonomous container databases in a given compartment on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. Under List Scope, select a compartment from the Compartment drop-down list.

View Details of an Autonomous Container Database

Follow these steps to view detailed information about an autonomous container database on an Oracle Exadata Database Service on Cloud@Customer system.



- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the database you wish to view details.

On the Autonomous Container Database Details page encryption details are displayed under **Encryption**.

If you have chosen customer-managed keys while creating the database, then you will see a link to the **Encryption Key Store** and **OKV Wallet Name**.

Note:

OKV Wallet Name represents the name of the wallet in which keys for this CDB are generated on the OKV.

Click the Key Store link to view details.

If you have chosen Oracle-managed keys while creating the database, then you will not see the link to **Encryption Key Store** and **OKV Wallet Name**.

Resource usage: For more information, see Resource Terminology.

In the **General Information** section, the **Time zone file version** field displays the current time zone file version.

• Update Autonomous Container Database Maintenance Preferences

Update Autonomous Container Database Maintenance Preferences

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the database you wish to view details.
- 4. Click More actions, and then click Edit Maintenance Preferences.
- 5. In the resulting Edit automatic Maintenance panel, enter the following details:
 - Configure maintenance method:
 - Rolling: The Rolling method updates the ACD, one node at a time, with no downtime for the Autonomous Databases.
 - Non-rolling The Non-rolling method shuts down and updates the ACD in parallel across all nodes. This method minimizes maintenance time but it requires full downtime for the ACD and all associated Autonomous Databases.
 - Optionally, you can also select **Enable time-zone update**. Time-zone files can only be updated using the non-rolling configuration method.



All the RUs that include time-zone updates will be patched using the non-rolling maintenance method (with full system downtime) only. So, if you configure your maintenance to the rolling method and enable time zone update, all the RUs that include a time-zone update are applied in the non-rolling method only. Only those RUs without a time-zone update are applied in the rolling method.

- Configure container database maintenance version: Choose Next release update (RU) or Latest release update (RU) to configure the next maintenance version of the Autonomous Container Database. Accordingly, the autonomous container database will be updated to the next immediate release or the latest release in the upcoming maintenance window.
- **Configure automatic maintenance schedule**: You can change the maintenance schedule settings, as desired:
 - Select No preference to permit Oracle to schedule maintenance as needed based on the maintenance schedule of the Exadata Infrastructure resource hosting the Autonomous Container Database.
 - Select Custom schedule to restrict when Oracle can schedule maintenance.
 You can choose your preferred month, week, weekday, start time, and buffer period between primary and standby maintenance execution for the Autonomous Container Database maintenance.
- 6. Click Save Changes.

Schedule a Quarterly Maintenance Update

You can schedule a one-time maintenance activity to update time zone file for an Autonomous Container Database. For ACDs with Oracle Home version prior to the latest release, you can also choose to update to the latest RU along with the time-zone file update in this schedule.

Note:

For on-demand maintenance updates with a time-zone file update, the standby ACD will be patched 3 days before the primary ACD in autonomous Data Guardenabled ACDs.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the database you wish to view details.
- 4. Click More actions, and then click Schedule maintenance.
- 5. In the resulting Schedule Maintenance panel, enter the following details:
 - Configure update type



- Release Update (RU): Click to schedule a release update (RU) along with a time zone file update. This option is disabled if your ACD is already patched to the latest RU.
- **Time-zone file update:** Click to schedule time zone file update only.

Time zone files can only be updated using the non-rolling configuration method that involves full system downtime.

- Configure maintenance method:
 - Rolling: The Rolling method updates the ACD, one node at a time, with no downtime for the Autonomous Databases.
 - Non-rolling: The Non-rolling method shuts down and updates the ACD in parallel across all nodes. This method minimizes maintenance time but it requires full downtime for the ACD and all associated Autonomous Databases.

• Enable time-zone update

Optionally, you can also select **Enable time-zone update**. Time-zone files can only be updated using the non-rolling configuration method.

If you configure your maintenance to the **rolling** method and **enable timezone update**, all the RUs that include a time-zone update are applied in the non-rolling method only. Only those RUs without a time-zone update are applied in the rolling method.

Schedule start time

Using the calendar tool, choose a date and time for the maintenance update.

6. Click Schedule maintenance.

Note:

The maintenance settings selected from the Schedule maintenance dialog apply to the maintenance update scheduled from this dialog only. These changes do not apply to the maintenance preferences defined at the ACD level.

Rotate CDB Encryption Key

Follow these steps to rotate the TDE Master key. On key rotation, the ACD life cycle goes through the regular updating state and returns to available.

You can rotate the TDE Master key as many times as you want. The new TDE Master Key is stored in the same wallet in which the previous key was stored. Rotating the TDE Master Key leads to the new key being generated in OKV and assigned to this database. You can view all of the keys in OKV.



You can rotate both Oracle-managed and customer-managed encryption keys.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the database you wish to view details.
- 4. On the Autonomous Container Database Details page, click Rotate Encryption Key.
- 5. On the Rotate Encryption Key dialog, click **Rotate Encryption Key**.

Change the Backup Retention Policy of an Autonomous Container Database

Follow these steps to update the backup retention policy of an autonomous container database on an Oracle Exadata Database Service on Cloud@Customer system.

Note:

By default, database backups are retained for 30 days if you have chosen Object Storage or NFS as a backup destination. You have the option of retaining backups for 7, 15, 30, or 60 days. If you have chosen Local storage as a backup destination, then by default, database backups are retained for a maximum of 7 days. If you have chosen Recovery Appliance as a backup destination, then you cannot update the backup retention policy.

The current backup retention policy for an Autonomous Container Database is displayed on the Autonomous Container Database details page.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the container database you are interested in.
- 4. On the Autonomous Container Database details page, under **Backup**, click the **Edit** link in the **Backup retention policy** field.
- 5. Specify a backup retention period from the list of choices.
- 6. Click Update.

Change the Maintenance Schedule of an Autonomous Container Database

Follow these steps to change the maintenance schedule of an autonomous container database on an Oracle Exadata Database Service on Cloud@Customer system.



- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the container database you are interested in.
- 4. On the Autonomous Container Database details page, under **Maintenance**, click the edit link in the **Maintenance Details** field.
- 5. In the Edit Automatic Maintenance dialog that opens, configure automatic maintenance schedule.
 - **No preference**: The system assigns a date and start time for container database maintenance.
 - **Specify a schedule**: Choose your preferred month, week, weekday, and start time for container database maintenance.
- 6. Click Save Changes.

Restart an Autonomous Container Database

Follow these steps to restart an autonomous container database on an Oracle Exadata Database Service on Cloud@Customer system.

The restart of an autonomous container database occurs in a rolling fashion, first stopping and starting one of the container database's database instances and then stopping and starting its other database instance.

Note:

You cannot restart an autonomous container database if a backup is in progress on any of its autonomous databases.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the container database you are interested in.

Autonomous Container Database details page is displayed.

Autonomous Container Database details page displays details that include total OCPU, available OCPU, reclaimable OCPU, and database memory per OCPU in GB.

- 4. On the Autonomous Container Database details page, click Restart.
- 5. In the confirmation dialog, type the name of the Autonomous Container Database.
- 6. Click Restart.
 - (or)
 - a. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.



- b. Click Autonomous Container Databases.
- c. In the list of Autonomous Container Databases, click the **Actions** icon (three dots) for the container database you are interested in, and then select **Restart**.
- d. In the confirmation dialog, type the name of the Autonomous Container Database.
- e. Click Restart.

Move an Autonomous Container Database to Another Compartment

Follow these steps to move an autonomous container database on an Oracle Exadata Database Service on Cloud@Customer system from one compartment to another compartment.

Note:

- To move an autonomous container database you must have the right to manage it in its current compartment and in the compartment you are moving it to.
- As soon as you move an autonomous container database to a different compartment, the policies that govern the new compartment apply immediately and affect access to the autonomous container database. Therefore, both your and other Oracle Cloud users' access to it may change, depending on the policies governing the user account's access to resources. For example, a user may lose the ability to create autonomous databases in the autonomous container database, given its new compartment.
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the container database you wish to move.
- 4. Click Move Resource.
- 5. Select the new compartment.
- 6. Click Move Resource.

Terminate an Autonomous Container Database

Follow these steps to terminate an autonomous container database on an Oracle Exadata Database Service on Cloud@Customer system.

Note:

You must terminate all Autonomous Databases within a container database before you can terminate the container database itself.



- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the infrastructure resource you are interested in.
- 4. Click Terminate.
- 5. In the confirmation dialog, type the name of the Autonomous Container Database, and then click **Terminate Autonomous Container Database**.

Using the API to Manage Autonomous Container 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.

The following table lists the REST API endpoints to manage Autonomous Container Databases.

Operation	REST API Endpoint
Create an Autonomous Container Database	CreateAutonomousContainerDatabase
View a list of Autonomous Container Databases	ListAutonomousContainerDatabases
View details of an Autonomous Container Database	GetAutonomousContainerDatabase
Change the backup retention policy of an Autonomous Container Database	UpdateAutonomousContainerDatabase
Change the maintenance schedule of an Autonomous Container Database	UpdateAutonomousContainerDatabase
Restart an Autonomous Container Database	RestartAutonomousContainerDatabase
Move an Autonomous Container Database to another compartment	ChangeAutonomousContainerDatabaseComp artment
Rotate an Autonomous Container Database key	RotateAutonomousContainerDatabaseEncrypt ionKey
Terminate an Autonomous Container Database	TerminateAutonomousContainerDatabase
Get the resource usage statistics of an Autonomous Container Database	GetAutonomousContainerDatabaseResource Usage

Managing Autonomous Databases

An Autonomous Database resource is a user database. When you create an Autonomous Database, you choose the Autonomous Container Database for it and you specify "Data Warehouse" or "Transaction Processing" as its workload type to create an Autonomous Data Warehouse database or an Autonomous Transaction Processing database.

You can create hundreds of Autonomous Databases on an Exadata Infrastructure. As described in *Available Exadata Infrastructure Hardware Shapes*, the maximum is determined by the capacity of your Exadata Infrastructure hardware.



- Oracle Autonomous Database for Developers
 Oracle Autonomous Database for Developers instances are free Autonomous Databases
 that developers can use to build and test new applications.
- Create an Autonomous Database
- Manage Access Control List of an Autonomous Database
- View a List of Autonomous Databases
- View Details of an Autonomous Database
- Rotate ADB Encryption Key
- Set the Password of an Autonomous Database's ADMIN User
- Scale the CPU Core Count or Storage of an Autonomous Database, or Enable/Disable or Alter the Percentage of System Global Area (SGA) for IM Column Store
- Enable or Disable Auto Scaling for an Autonomous Database
- Move an Autonomous Database to Another Compartment
- Stop or Start an Autonomous Database
- Restart an Autonomous Database
- Back Up an Autonomous Database Manually
- Create a Long-Term Backup
- View Details and Edit Retention Period of a Long-Term Backup
- Delete a Long-Term Backup
- Restore an Autonomous Database
- Clone an Autonomous Database
- Clone an Autonomous Database Backup
- Clone a Standby Database
- Clone a Standby Database Backup
- Terminate an Autonomous Database
- API to Manage Autonomous Databases
- Monitor Performance with Autonomous Database Metrics

Related Topics

Available Exadata Infrastructure Hardware Shapes

Oracle Autonomous Database for Developers

Oracle Autonomous Database for Developers instances are free Autonomous Databases that developers can use to build and test new applications.

With Autonomous Database for Developers instances, you can try new Autonomous Database features for free and apply them to ongoing or new development projects. Developer databases are limited in resources, so they are not suitable for large-scale testing and production deployments. When you need more compute or storage resources, you can transition to a paid database licensing by cloning your developer database into a regular Autonomous Database.



Requirements

To create an Autonomous Database for Developers instance, you must have access to Oracle Exadata Database Service or Autonomous Database on either a Dedicated Exadata Infrastructure or Exadata Cloud@Customer. In other words, only those customers with active subscriptions to any of the following service platforms can create developer databases:

- Autonomous Database on Dedicated Exadata Infrastructure
- Exadata Database Service on Dedicated Exadata Infrastructure
- Autonomous Database on Exadata Cloud@Customer
- Exadata Database Service on Cloud@Customer

There is no limit on the number of free developer databases; it's limited by the capacity of your Exadata infrastructure.

Provisioning Workflow

You can provision an Autonomous Database for Developers instance from the Oracle Cloud Infrastructure (OCI) console or using API. To create a developer database, you need an ACD without an Autonomous Data Guard in an ECPU-based AVMC. If you do not have these resources provisioned already, create the ECPU-based AVMC first and then create an ACD without disaster recovery (Autonomous Data Guard) using that AVMC.

After creating or identifying (if they already exist) the AVMC and ACD, you can create an Autonomous Database for Developers using them. Provisioning a developer database using the OCI console follows the same workflow as creating a regular Autonomous Database, as described in *Create Autonomous Database*. Once created, the Autonomous Database for Developers instances appear with a Developer label in the list of Autonomous Databases on the OCI console.

Specifications

Each developer database comes with the following specifications:

- **Compute:** Fixed 4 ECPUs, with no CPU scaling
- Storage: Fixed 32 GB (~ 20 GB of DATA)
- Session limits: 30 concurrent database sessions
- Workload Type: Data Warehouse, Transaction Processing

Excluded Features

Autonomous Database for Developers supports all the features offered by a regular Autonomous Database except those listed below. These limitations are in place to ensure that the developer databases are optimally used as a development sandbox.

Developer database instances:

- Do not support Autonomous Data Guard. Hence, they can only be provisioned in an ACD without Autonomous Data Guard.
- Support ECPU only. Therefore, you can provision them only on an ECPU based ACD.



- Come with fixed compute and storage sizing, do not support manual or auto-scaling and storage scaling.
- Can not have long-term backups.
- Do not provide Database In-memory.

Supported Features

- **Cloning:** Autonomous Database for Developers offers fewer resources and features than a regular autonomous database. For non-development use, such as load/stress testing and production, or to get access to all features, users can use cloning to clone from a developer database to a regular Autonomous database. You can also clone a regular database into a developer database; however, to successfully clone a regular database into a developer database, the source database's actual used space, rounded up to the next GB, must be <= 32GB.
- **Backup and Recovery:** You can enable automatic backups or trigger manual backups of your developer database, as needed. If the backup destination is an Object Storage and Recovery Service, the backups will be billed.
- Service Maintenance: Developer databases follow the same patching schedule as regular Autonomous Database; however, there will be no support for critical one-off patches.
- Database Application Development and Developer Tools: With Autonomous Database for Developers, you can use all the developer-related features and built-in tools an Autonomous Database offers.

Autonomous Database for Developers comes with a service level objective (SLO) of 99.5% and you can log service requests (SR) to Oracle Support for assistance. However, there is no Severity 1 SR support for developer databases. See Create a Service Request in My Oracle Support to learn how to contact Oracle Support for assistance.

Create an Autonomous Database

Follow these steps to create an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system.

Note:

- If the standby ACD is in snapshot standby mode, then you cannot create an ADB in the primary ACD.
- For better management and sharing of the underlying SGA/memory resources, Oracle recommends that all Autonomous Databases configured for In-Memory be in the same Autonomous Container Database.
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. Click Create Autonomous Database.
- 4. In the Create Autonomous Database dialog, enter the following:



Basic Database Information

- **Compartment**: Select the compartment of the Autonomous Database.
- **Display Name**: A user-friendly description or other information that helps you easily identify the resource. The display name does not have to be unique. Avoid entering confidential information.
- **Database Name**: The database name must consist of letters and numbers only, starting with a letter. The maximum length is 14 characters. Avoid entering confidential information.

Workload Type

Select the desired workload type.

- Data Warehouse
- Transaction Processing

See *About Autonomous Database* for information about each workload type. **Autonomous Container Database**: Select an Autonomous Container Database.

Compartment: Specify the compartment containing the Autonomous Container Database you wish to use.

Configure the database: Free Instance:

Note:

As developer database instances can only be created on ECPU-based ACDs without Autonomous Data Guard, the Free instance toggle button is disabled for ACDs with OCPU, Autonomous Data Guard, or both.

Toggle the **Free instance** button on, if you want to create an Autonomous Database for Developers instance.

ECPU count and **Storage (GB)** are auto-populated with 4 and 32 respectively because Autonomous Database for Developers comes fixed at 4 ECPUs and 32GB storage

Compute auto-scaling is disabled because developer database instances do not support manual or auto-scaling.

Configure the database

• **CPU Count**: Select the number of CPUs for your database from the list of provisionable CPUs. The CPU type, that is, OCPU or ECPU is determined by the parent

Autonomous Exadata VM Cluster resource's compute type.

This value defaults to 1 OCPU.

For ECPUs, this value defaults to 2 ECPUs.

You can also select a fractional OCPU value for databases that do not need an entire OCPU. This allows you to overprovision CPU and run more databases on each infrastructure instance. Refer to CPU Overprovisioning for more details.



CPU Overprovisioning is not allowed with ECPUs.

Databases with CPU over-provisioning can only connect using tp and low services.

Auto scaling: Enable or disable auto-scaling, which permits Autonomous Database to automatically use up to three times the allocated CPUs as the workload on the database increases.

- Storage (GB): Specify the storage you wish to make available to your Autonomous Database, in GB. The available storage depends on the infrastructure shape and what is already consumed by other Autonomous Databases.
 - Default: 1024 GB
 - Minimum: 32 GB
 - Increment: 1 GB

Administrator Credentials

Set the password for the Autonomous Database Admin user by entering a password that meets the following criteria. Use this password when accessing the Autonomous Database service console and when using an SQL client tool.

- Contains from 12 to 30 characters
- Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

Configure network access You can optionally create an ACL during database provisioning, or at any time thereafter.

- Select the Enable database level access control checkbox.
- Click Access Control Rule.

Note:

The database-level access control will be enabled without any IP addresses in the access control list. Enabling an access control list with an empty list of IP addresses makes the database inaccessible to all clients

- Specify the following types of addresses in your list by using the IP notation type drop-down selector:
 - IP Address allows you to specify one or more individual public IP addresses. Use commas to separate your addresses in the input field.
 - CIDR Block allows you to specify one or more ranges of public IP addresses using CIDR notation. Use commas to separate your CIDR block entries in the input field.



Advanced Options:

- Encryption Key: ADB inherits encryption settings from the parent ACD. If the
 parent ACD is configured for customer-managed OKV-based encryption, then
 the child ADB will also have a TDE Master Key generated and managed in the
 same OKV wallet used to store ACD master keys. Additionally, any backups
 taken on the Autonomous Database will have the OKV-based key associated
 with it.
- Database In-memory:
 - Enable database In-memory: It requires at least four OCPUs and a
 percentage of the System Global Area (SGA) to enable in-memory. If you
 enable In-memory, select the percentage of SGA to allocate to the IM
 Column Store. In-memory may have an impact on the autonomous
 database's performance if a large amount of memory is allocated or if it is
 disabled.
- **Management**: Choose a Character Set and National Character from the dropdown list.
- **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.
- 5. 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.
- 6. Click Create Autonomous Database.



The following naming restrictions apply to Autonomous Transaction Processing and Autonomous Data Warehouse databases:

- Names associated with databases terminated within the last 60 days cannot be used when creating a new database.
- A database name cannot be used concurrently for both an Autonomous Data Warehouse and an Autonomous Transaction Processing database.

Related Topics

- About Autonomous Database
- GetAutonomousContainerDatabase
- Resource Tags
- Resource Manager and Terraform

Manage Access Control List of an Autonomous Database

An access control list (ACL) provides additional protection to your database by allowing only the clients with specific IP addresses to connect to the database. You can add IP addresses individually, or in CIDR blocks.

You can optionally create an ACL during database provisioning, or at any time thereafter. You can also edit an ACL at any time. Enabling an access control list with an empty list of IP addresses makes the database inaccessible.

Note the following about using an ACL with your Autonomous Database:

- The Autonomous Database Service console is not subject to ACL rules.
- Oracle Application Express (APEX), RESTful services, and SQL Developer Web are not subject to ACLs. Choosing to enable an ACL disables these features automatically.
- Performance Hub is not subject to ACL rules.
- While creating a database, if setting ACL fails, then provisioning the database also fails.
- Updating ACL is allowed if the database is in Available and AvailableNeedsAttention states.
- Restoring a database does not overwrite the existing ACLs.
- Cloning a database, full and metadata, will have the same access control settings as the source database. You can make changes as necessary.
- All CDB operations are allowed during ACL update. However, ADB operations are not allowed during ACL update.
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Choose your Compartment.
- 3. In the list of Autonomous Databases, click the display name of the database you want to administer.



4. Under **Network** in the database details, find the **Access Control List** field and click **Edit** to enable or disable your database-level access control and make changes to the ACL rules.

Note:

Autonomous Data Guard enabled Automonous databases:

- You can only view ACLs for standby databases.
- You can reset ACL for both the primary and standby databases from the primary database details page. You cannot configure ACL from the standby database details page.
- 5. In the Access Control List dialog, add or modify entries, as applicable.

If you are editing an ACL, the ACL's existing entries display in the Access Control List dialog. Do not overwrite the existing values unless you intend to replace one or more entries. To add new ACL entries, click **+** Access Control Rule.

You can specify the following types of addresses in your list by using the IP notation type drop-down selector:

- IP Address allows you to specify one or more individual public IP addresses. Use commas to separate your addresses in the input field.
- CIDR Block allows you to specify one or more ranges of public IP addresses using CIDR notation. Use commas to separate your CIDR block entries in the input field.

Click + Access Control Rule to add additional access rules to your list.

To remove an access control rule, simply delete the entry from the list. Deleting all access control rules from the ACL will render the database inaccessible because the allow list is empty.

To disable the database-level access control configuration, clear the **Enable database level access control** checkbox. Once ACL is disabled and the configuration is saved, all the access control rules are removed from the ACL and no longer applicable.

6. Click Save Changes.

If the Lifecycle State is Available when you click Save, the Lifecycle State changes to Updating until the ACL update is complete. The database is still up and accessible, there is no downtime. When the update is complete the Lifecycle State returns to Available and the network ACL rules from the access control list are in effect.

View a List of Autonomous Databases

Follow these steps to view a list of autonomous databases on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.



View Details of an Autonomous Database

Follow these steps to view detailed information about an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to view details.

In the resulting Autonomous Database details page

- Encryption details are displayed in the Encryption section.
 - If you have chosen customer-managed keys while creating the database, then you will see a link to the Encryption Key Store and OKV Wallet Name. Click the Key Store link to view details.
 - If you have chosen Oracle-managed keys while creating the database, then you will not see the link to Encryption Key Store and OKV Wallet Name.
- In-memory details are displayed in the **Resource allocation** section.
 - If you have not enabled In-memory, then the system displays the Enable link. Click it to enable In-memory.
 - If you have enabled and wish to modify the settings, then click **Edit**.

Rotate ADB Encryption Key

Follow these steps to rotate the TDE Master key. On key rotation, the ADB life cycle goes through the regular updating state and returns to available.

You can rotate the TDE Master key as many times as you want. The new TDE Master Key is stored in the same wallet in which the previous key was stored. Rotating the TDE Master Key leads to the new key being generated in OKV and assigned to this database. You can view all of the keys in OKV.

Note:

You can rotate both Oracle-managed and customer-managed encryption keys.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to view details.
- 4. On the Autonomous Database Details page, from the **More Actions** drop-down list, select **Rotate Encryption Key**.
- 5. On the Rotate Encryption Key dialog, click Rotate Encryption Key.



Set the Password of an Autonomous Database's ADMIN User

Follow these steps to set the ADMIN database user's password for an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to administer.
- 4. From the More Actions drop-down list, select Admin Password.

The Admin Password dialog opens.

5. Enter a password for the Autonomous Database.

The password must meet the following criteria:

- Contains from 12 to 30 characters
- Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of case
- Is not one of the last four passwords used for the database
- Is not a password you previously set within the last 24 hours
- 6. Enter the password again in the **Confirm Password** field.
- 7. Click Update.

Scale the CPU Core Count or Storage of an Autonomous Database, or Enable/Disable or Alter the Percentage of System Global Area (SGA) for IM Column Store

Follow these steps to scale the CPU core count or storage of an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system up or down.

Note:

- If the standby ACD is in snapshot standby mode, then you cannot scale an ADB in the primary ACD.
- For better management and sharing of the underlying SGA/memory resources, Oracle recommends that all Autonomous Databases configured for In-Memory be in the same Autonomous Container Database.



- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to view details.
- 4. Click Scale Up/Down.

This option is not enabled for an Autonomous Database for Developers instance.

- 5. Enter a new value for CPU Core Count or Storage.
 - OCPU Count: Select the number of CPUs for your database from the list of provisionable CPUs.

Based on the resource utilization on each node; not all the values of the available CPUs can be used to provision or scale Autonomous Databases. For example, suppose you have 20 CPUs available at the AVMC level, not all the values from 1 to 20 CPUs can be used to provision or scale Autonomous Databases depending on the resource availability at the node level. The list of CPU values that can be used to provision or scale an Autonomous Database is called **Provisionable CPUs**.

On the console, when you try to provision or scale an Autonomous Database, the CPU count will be validated against the list of provisionable CPUs, and if the value is not provisionable, you will be provided with the two nearest provisionable CPU values. Alternatively, if you want to see the complete list of provisionable CPU values for an Autonomous Exadata VM Cluster, you can use the following API:

GetAutonomousContainerDatabase returns a list of provisionable CPU values that can be used to create a new Autonomous Database in the given Autonomous Container Database. See *GetAutonomousContainerDatabase* for more details.

GetAutonomousDatabase returns a list of provisionable CPU values that can be used for scaling a given Autonomous Database. See *GetAutonomousDatabase* for more details.

You can also select a fractional OCPU value for databases that do not need an entire OCPU. This allows you to overprovision CPU and run more databases on each infrastructure instance. Refer to CPU Overprovisioning for more details.

- Auto scaling: Enable or disable auto-scaling, which permits Autonomous Database to automatically use up to three times the allocated CPUs as the workload on the database increases.
- Storage (GB): Specify the storage you wish to make available to your Autonomous Database, in GB. The available storage depends on the infrastructure shape and what is already consumed by other Autonomous Databases.
 - Default: Current value
 - Minimum: 32 GB
 - Increment: 1 GB
- Enable/Disable or Alter Percentage of System Global Area (SGA): It requires at least four OCPUs and a percentage of the System Global Area (SGA) to enable inmemory. If you enable In-memory, select the percentage of SGA to allocate to the IM



Column Store. In-memory may have an impact on the autonomous database's performance if a large amount of memory is allocated or if it is disabled.

6. Click Update.

You can find the memory allocated to In-memory in the **Resource allocation** section on the Autonomous Database details page.

- Click Edit to modify the settings.
- If you have not enabled In-memory, then the system displays the **Enable** link. Click it to enable In-memory.

Related Topics

- Service Limits
- GetAutonomousContainerDatabase
- GetAutonomousDatabase

Enable or Disable Auto Scaling for an Autonomous Database

Oracle Autonomous Database on Oracle Exadata Cloud@Customer systems provides an auto-scaling feature that automatically increases the number of CPUs in an autonomous database during periods of increased demand and, as demand returns to normal, automatically decreases the number of cores down to the databases's base number.

Note the following points regarding the auto-scaling feature:

- With auto-scaling enabled, the database can use up to three times more CPU and IO resources than specified by the number of CPUs currently shown in the Scale Up/Down dialog.
- If auto-scaling is disabled while more CPU cores are in use than the database's currently assigned number of cores, then Autonomous Database scales the number of CPU cores in use down to the assigned number.
- Enabling auto scaling does not change the concurrency and parallelism settings for the predefined services.

Follow these steps to enable or disable auto-scaling for an autonomous database.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to view details.
- 4. Click Scale Up/Down.
- 5. Check **Auto Scaling** to enable the auto-scaling feature, or uncheck **Auto Scaling** to disable the feature.
- 6. Click Update.



Fip: You can view the number of CPUs the database is currently using by running the following SQL statements:
 ECPU: SELECT AVG_RUNNING_SESSIONS FROM V\$RSRCPDBMETRIC;

 OCPU: SELECT AVG_RUNNING_SESSIONS / 2 FROM V\$RSRCPDBMETRIC;

Move an Autonomous Database to Another Compartment

Follow these steps to move an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system from one compartment to another compartment.

Note:

- To move an autonomous database you must have the right to manage it in its current compartment and in the compartment you are moving it to.
- As soon as you move an autonomous database to a different compartment, the
 policies that govern the new compartment apply immediately and affect access
 to the autonomous database. Therefore, both your and other Oracle Cloud
 users' access to it may change, depending on the policies governing the user
 account's access to resources. For example, a user may lose the ability to
 manage the autonomous databae, given its new compartment.
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to move.
- 4. From the More Actions drop-down list, select Move Resource.
- 5. Select the new compartment.
- 6. Click Move Resource.

Stop or Start an Autonomous Database

Follow these steps to stop or start an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.



- 3. In the list of Autonomous Databases, click the display name of the database you wish to view details.
- 4. Click Stop (or Start).

When you stop your Autonomous Database, billing stops for CPU usage. Billing for storage continues when the database is stopped.

5. Confirm that you want to stop or start your Autonomous Database in the confirmation dialog.

Note:

Stopping your database has the following consequences:

- On-going transactions are rolled back.
- You will not be able to connect to your database using database clients or tools.

Restart an Autonomous Database

To resolve some autonomous database issues with minimal downtime on Exadata Database Service on Cloud@Customer systems, you can restart the database.

Restarting an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system is equivalent to manually stopping and then starting the database. Using restart allows you to minimize downtime and requires only a single action.

Follow these steps to restart an autonomous database.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- In the list of Autonomous Databases, click the display name of the database you wish to restart.
- 4. Click Restart.
- 5. Confirm that you want to restart your Autonomous Database in the confirmation dialog.

The system stops and then immediately starts your database.

Back Up an Autonomous Database Manually

Oracle Autonomous Database automatically backs up autonomous databases on an Oracle Exadata Database Service on Cloud@Customer system. In addition, you can manually back up an autonomous database should the need arise.



During the backup operation, your autonomous database remains available. However, lifecycle management operations such as stopping it, scaling it, or terminating it are disabled.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to back up.
- 4. On the Details page, under **Resources**, click **Backups**.
- 5. Click Create Manual Backup.
- 6. In the Create Manual Backup dialog, enter a name for your backup. Avoid entering confidential information.
- 7. Click Update.

The backup operation begins. This operation may take several hours to complete, depending on the size of the database.

Optionally, you can check the state of your backup in the list of backups on the database details page. For some states, an information icon is displayed to provide additional details regarding the state or ongoing operations like deletions. The backup has one of the following states:

- Creating
- Active
- Deleting
- Deleted
- Failed

Create a Long-Term Backup

To create a long-term backup, use this procedure.

Note:

Long-term backups are not available with Autonomous Database for Developers instances. See *Oracle Autonomous Database for Developers* for more details.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to create a long-term backup.



Autonomous Details page is displayed.

- 4. Under **Resources**, click **Backups**.
- 5. In the Backups section, click **Create long-term backup**.
- 6. In the resulting window, enter the following details:
 - **Name**: Enter a user-friendly description or other information that helps you easily identify the backup.
 - **Backup destination type**: Network File System (NFS) is selected by default. In this release only NFS is supported, so you cannot change it no matter what destination type (Object Storage, Network File System, or Oracle Zero Data Loss Recovery Appliance) you have chosen for the ACD.
 - **Backup destinations**: Specify an NFS destination. You can use an existing NFS destination or create one for this long-term backup.
 - To choose an existing NFS destination:
 - * Click Backup Destinations under Infrastructure.
 - * Choose an NFS destination from the list of NFS backup destinations in the chosen compartment.
 - To create an NFS backup destination, see Using the Console to Create a Backup Destination.
 - **Retention period**: Set the retention period.
- 7. Click Create

Related Topics

• Using the Console to Create a Backup Destination

View Details and Edit Retention Period of a Long-Term Backup

To view the details of a long-term backup and edit the retention period, use this procedure.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- **3.** In the list of Autonomous Databases, click the display name of the database you wish the details of a long-term backup.
- 4. In the resulting Autonomous Details page, under **Resources**, click **Backups**.
- 5. In the Backups section, identify the backup, and review the details.
- 6. To edit the retention period, click the action icon (three dots) and select **Edit** retention period.
- 7. In the resulting window, set the retention period.
- 8. Click Save.

Delete a Long-Term Backup

To delete a long-term backup, use this procedure.



- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to view the details of a long-term backup.
- 4. In the resulting Autonomous Details page, under **Resources**, click **Backups**.
- 5. In the Backups section, identify the backup, click the action icon (three dots), and then select **Delete**.
- 6. In the resulting window, click **Delete** if you are sure about deleting it.

Restore an Autonomous Database

You can use any existing manual or automatic backup to restore and recover an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system, or you can restore and recover the database to any point in time during the retention period of its automatic backups.

Note:

- Restoring an autonomous database puts the database in an unavailable state during the restore operation. You cannot connect to a database in this state. The only lifecycle management operation supported in the unavailable state is terminate.
- You cannot perform a restore operation on a primary ADB if the standby database is in snapshot standby mode. Convert your standby ACD to physical standby mode to restore an Autonomous Database.
- Restore from a Backup
- Restore to a Point in Time

Restore from a Backup

Follow these steps to restore an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system from a specific backup.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you want to clone.
- 4. From the More Actions drop-down list, select Restore.
- 5. Specify the date range for a list of backups to display.
- 6. Select the backup.
- 7. Click Restore.



Restore to a Point in Time

Follow these steps to restore an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system to a specific point in time.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you want to restore.
- 4. From the More Actions drop-down list, select Restore.
- 5. Click Specify Timestamp.
- 6. Enter a timestamp.

Your Autonomous Database decides which backup to use for faster recovery. The timestamp input allows you to specify precision to the seconds level (YYYY-MM-DD HH: MM: SS GMT).

7. Click Restore.

Clone an Autonomous Database

Follow these steps to clone an autonomous database on an Oracle Exadata Cloud@Customer system.

You can use the cloning feature to create a point-in-time copy of your Autonomous Database for purposes such as testing, development, or analytics. To clone only the database schema of your source database, choose the metadata clone option.

Note:

If IM is enabled, the source In-Memory Column Store settings or parameters will not be applied to the clone. However, you can enable In-Memory Column Store like a normal ADB creation flow.

Clone Types

The clone feature offers the following two types of Autonomous Database clones:

- The full-clone option creates a database that includes the metadata and data from the source database.
- The metadata-clone option creates a database that includes only the metadata from the source database.

Steps

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you want to clone.



- 4. From the More Actions drop-down list, select Create Clone.
- 5. On the Create Autonomous Database Clone page, provide the following information:

In the **Clone Type** section, select the type of clone you want to create. Choose either **Full Clone** or **Metadata Clone**.

Clone Source: The clone source selection allows you to specify whether the clone is created from a running database or from a database backup. Select one of the following options:

- **Clone from a database instance**: Creates a clone of a running database as it exists at the current moment.
- **Clone from a backup**: Creates a clone from a database backup. If you choose this option, select one of the following options:
 - **Specify a timestamp**: Creates a point-in-time clone. The timestamp has to be between the first and latest backups of the database.
 - Select from a list of backups: Creates a clone using all data from the specified backup. To limit your list of backups to a specific date range, enter the starting date in the From field and the ending date in the To field.

Provide basic information for the Autonomous Database.

- **Choose a compartment**: Your current compartment is the default selection but you can select a different compartment in which to create the clone from the drop-down list.
- **Source database name**: The name of the source database displays in the read-only Source database name field.
- **Display name**: Enter a description or other information to identify the database clone. You can change the display name any time and it does not have to be unique. Avoid entering confidential information.
- **Database name**: Enter a database name for the clone that contains only letters and numbers, begins with a letter. Avoid entering confidential information.
- Three additional fields are displayed if you opt to clone from a backup.
 - **Region**: Choose your preferred region to place the clone database.
 - Exadata Infrastructure: You can choose to create the database clone in the same Exadata Infrastructure where the source database resides, or you can choose a different compartment by clicking CHANGE COMPARTMENT and choosing one from the drop-down list.
 - Autonomous Exadata VM Cluster: You can choose to create the database clone in the same Autonomous Exadata VM Cluster where the source database resides, or you can choose a different compartment by clicking CHANGE COMPARTMENT and choosing one from the drop-down list.
- Autonomous Container Database: You can choose to create the database clone in the same compartment and container database as the source database, or you can choose a different compartment by clicking CHANGE COMPARTMENT, and a different container database by choosing one from the drop-down list.
- Configure the database: Free Instance: Toggle the Free instance button on, if you want to create an Oracle Autonomous Database for Developers instance. ECPU count and Storage (GB) are auto-populated with 4 and 32 respectively because Oracle Autonomous Database for Developers comes fixed at 4 ECPUs and 32GB



storage Compute auto-scaling is disabled because developer database instances do not support manual or auto-scaling.

Note:

- As developer database instances can only be created on ECPUbased ACDs without Autonomous Data Guard, the Free instance toggle button is disabled for ACDs with OCPU, Autonomous Data Guard, or both.
- To successfully clone a backup to an Oracle Autonomous Database for Developers instance, the current storage allocation of the backup database must be 32GB. In case, this condition is not met, you can try cloning the database instance to a developer database, as long as its actual used space, rounded up to the next GB, must be <= 32GB.

Configure the database

CPU Count: Select the number of CPUs for your clone database from the list of provisionable CPUs.
 The CPU type, that is, OCPU or ECPU is determined by the parent Autonomous Exadata VM Cluster resource's compute type.

This value defaults to 1 OCPU.

You can also select a fractional OCPU value for databases that do not need an entire OCPU. This allows you to overprovision CPU and run more databases on each infrastructure instance. Refer to CPU Overprovisioning for more details.

Note:

CPU Overprovisioning is not allowed with ECPUs.

Databases with CPU over-provisioning can only connect to the tp and low services for the Autonomous Database for Transaction Processing and Mixed Workloads workloads. In the case of an Autonomous Database for Analytics and Data Warehousing workloads, you can only connect to the low services when create on over-provisioned CPUs.

For ECPUs, this value defaults to 2 ECPUs. For databases that need 2 or more ECPUs, you must specify the number of assigned ECPUs in increment of 1.

- **Storage (GB)**: Specify the amount of storage, in GB, that you want to make available to your cloned Autonomous Database, and it depends on the storage available to use. For full clones, the size of the source database determines the minimum amount of storage you can make available.
 - Default: 1024 GB
 - Minimum: 32 GB
 - Increment: 1 GB



 Auto scaling: Enable or disable auto scaling, which permits Autonomous Database to automatically use up to three times the allocated CPUs as the workload on the database increases.

Create administrator credentials

Set the password for the Autonomous Database administrator user by entering a password that meets the following criteria.

- Password cannot be one of the three most recently used passwords of the source database
- Between 12 and 30 characters long
- Contains at least one lowercase letter
- · Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

Use this password when accessing the service console and when using a SQL client tool.

Configure Network Access

You can change the access control list to enable or disable database-level access control or add or modify entries to the access control list.

- Click Modify Access Control.
- Select the Enable database level access control check box.
- Click Access Control Rule.

Note: The database-level access control will be enabled without any IP addresses in the access control list. Enabling an access control list with an empty list of IP addresses makes the database inaccessible to all clients.

- Specify the following types of addresses in your list by using the IP notation type drop-down selector:
 - IP Address allows you to specify one or more individual public IP addresses. Use commas to separate your addresses in the input field.
 - CIDR Block allows you to specify one or more ranges of public IP addresses using CIDR notation. Use commas to separate your CIDR block entries in the input field.

Advanced Options:

- Encryption Key:
 - Clone from a database instance: The source and the target ACD must be the same Keystore type. When the source is OKV, the target must also be the same OKV destination.
 - Clone from a backup: The source and the target ACDs can be different Keystore types. When the source is OKV, the target must also be the same OKV destination.
- Database In-memory:
 - Enable database In-memory: It requires at least four OCPUs and a percentage of the System Global Area (SGA) to enable in-memory. If you enable In-memory, select the percentage of SGA to allocate to the IM Column Store. In-memory may



have an impact on the autonomous database's performance if a large amount of memory is allocated or if it is disabled.

- Management: Choose a Character Set and National Character from the drop-down list.
- **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 Autonomous Database Clone.

The Console displays the details page for the new clone of your database and the service begins provisioning the Autonomous Database. Note the following:

- The new clone displays the **Provisioning** lifecycle state until the provisioning process completes.
- The source database remains in the **Available** lifecycle state.
- Backups associated with the source database are not cloned for either the fullclone or the metadata-clone option.

The **Clone source** is displayed in the **General Information** section of the cloned database details page. Click the name to view details of the source database. Note that if the source database is deleted, then this key/value pair is not displayed.

Related Topics

- GetAutonomousContainerDatabase
- Resource Tags

Clone an Autonomous Database Backup

Follow these steps to clone an autonomous database on an Oracle Exadata Cloud@Customer system.

You can use the cloning feature to create a point-in-time copy of your Autonomous Database for purposes such as testing, development, or analytics. To clone only the database schema of your source database, choose the metadata clone option.

Note:

If IM is enabled, the source In-Memory Column Store settings or parameters will not be applied to the clone. However, you can enable In-Memory Column Store like a normal ADB creation flow.

Clone Types

The clone feature offers the following two types of Autonomous Database clones:

• The full-clone option creates a database that includes the metadata and data from the source database.



• The metadata-clone option creates a database that includes only the metadata from the source database.

Steps

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you want to clone.
- 4. Under Resources, click Backups.
- 5. In the list of backups, find the backup that you want to clone, click the action icon (three dots), and then click **Create Clone**.
- 6. On the Create Autonomous Database Clone page, provide the following information:

Provide basic information for the Autonomous Database.

- **Choose a compartment**: Your current compartment is the default selection but you can select a different compartment in which to create the clone from the drop-down list.
- **Source database name**: The name of the source database displays in the read-only Source database name field.
- **Display name**: Enter a description or other information to identify the database clone. You can change the display name any time and it does not have to be unique. Avoid entering confidential information.
- **Database name**: Enter a database name for the clone that contains only letters and numbers, begins with a letter. Avoid entering confidential information.
- **Region**: Choose your preferred region to place the clone database.
- **Exadata Infrastructure**: You can choose to create the database clone in the same Exadata Infrastructure where the source database resides, or you can choose a different compartment by clicking **CHANGE COMPARTMENT** and choosing one from the drop-down list.
- Autonomous Exadata VM Cluster: You can choose to create the database clone in the same Autonomous Exadata VM Cluster where the source database resides, or you can choose a different compartment by clicking CHANGE COMPARTMENT and choosing one from the drop-down list.
- Autonomous Container Database: You can choose to create the database clone in the same compartment and container database as the source database, or you can choose a different compartment by clicking CHANGE COMPARTMENT, and a different container database by choosing one from the drop-down list.

Note:

When the target Autonomous Exadata VM Cluster is the same as the source, then the database name cannot be the same as the source database name.

• **Configure the database: Free Instance**: Toggle the **Free instance** button on, if you want to create an Oracle Autonomous Database for Developers instance. **ECPU**

count and **Storage (GB)** are auto-populated with 4 and 32 respectively because Oracle Autonomous Database for Developers comes fixed at 4 ECPUs and 32GB storage Compute auto-scaling is disabled because developer database instances do not support manual or auto-scaling.

Note:

- As developer database instances can only be created on ECPUbased ACDs without Autonomous Data Guard, the Free instance toggle button is disabled for ACDs with OCPU, Autonomous Data Guard, or both.
- To successfully clone a backup to an Oracle Autonomous Database for Developers instance, the current storage allocation of the backup database must be 32GB. In case, this condition is not met, you can try cloning the database instance to a developer database, as long as its actual used space, rounded up to the next GB, must be <= 32GB.

Configure the database

• **CPU Count**: Select the number of CPUs for your clone database from the list of provisionable CPUs.

After the clone, you can resize to a lower value if needed. You can even resize the CPU count to less than 1 OCPU (0.1 to 0.9 in increments of 0.1 OCPUs) to databases that do not need a full CPU. This allows you to overprovision CPU and run more databases on each infrastructure instance. Note that fractional CPU applies to OCPU only.

There is a minimum requirement of 1 OCPU or 4 ECPUs for an Autonomous Database clone from Backup.

The total number of CPUs available to all databases within the Autonomous Exadata VM Cluster depends on the infrastructure shape and what is already allocated to other Autonomous Databases.

The CPU type, that is, OCPU or ECPU is determined by the parent Autonomous Exadata VM Cluster resource's compute type.

The time taken to clone an Autonomous Database depends on the CPU Count and the network bandwidth between the Backup Destination and the target Autonomous Container Database.

You can also select a fractional OCPU value for databases that do not need an entire OCPU. This allows you to overprovision CPU and run more databases on each infrastructure instance. Refer to CPU Overprovisioning for more details.

For databases that need 2 or more ECPUs, you must specify the number of assigned ECPUs in increment of 1.

Note:

CPU Overprovisioning is not allowed with ECPUs.

Databases with CPU over-provisioning can only connect to the tp and low services for the Autonomous Database for Transaction Processing and Mixed Workloads workloads. In the case of an Autonomous Database for Analytics and Data Warehousing workloads, you only connect to the low services when created on over-provisioned CPUs.

- **Storage (GB)**: Specify the amount of storage, in GB, that you want to make available to your cloned Autonomous Database, and it depends on the storage available to use.
 - **Default/Minimum:** Allocated storage of the source database
 - Increment: 1 GB
- Auto scaling: Enable or disable auto-scaling, which permits Autonomous Database to automatically use up to three times the allocated CPUs as the workload on the database increases.

Create administrator credentials

Set the password for the Autonomous Database administrator user by entering a password that meets the following criteria.

- Password cannot be one of the three most recently used passwords of the source database
- Between 12 and 30 characters long
- Contains at least one lowercase letter
- · Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

Use this password when accessing the service console and when using a SQL client tool.

Configure Network Access

You can change the access control list to enable or disable database-level access control or add or modify entries to the access control list.

- Click Modify Access Control.
- Select the Enable database level access control check box.
- Click Access Control Rule.

Note: The database-level access control will be enabled without any IP addresses in the access control list. Enabling an access control list with an empty list of IP addresses makes the database inaccessible to all clients.

- Specify the following types of addresses in your list by using the IP notation type drop-down selector:
 - IP Address allows you to specify one or more individual public IP addresses. Use commas to separate your addresses in the input field.
 - CIDR Block allows you to specify one or more ranges of public IP addresses using CIDR notation. Use commas to separate your CIDR block entries in the input field.

Advanced Options:



- Encryption Key:
 - Clone from a database instance: The source and the target ACD must be the same Keystore type. When the source is OKV, the target must also be the same OKV destination.
 - Clone from a backup: The source and the target ACDs can be different Keystore types. When the source is OKV, the target must also be the same OKV destination.
- Database In-memory:
 - Enable database In-memory: It requires at least four OCPUs and a
 percentage of the System Global Area (SGA) to enable in-memory. If you
 enable In-memory, select the percentage of SGA to allocate to the IM
 Column Store. In-memory may have an impact on the autonomous
 database's performance if a large amount of memory is allocated or if it is
 disabled.
- Management: Choose a Character Set and National Character from the drop-down list.
- **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.

7. Click Create Autonomous Database Clone.

The Console displays the details page for the new clone of your database and the service begins provisioning the Autonomous Database. Note the following:

- The new clone displays the **Provisioning** lifecycle state until the provisioning process completes.
- The source database remains in the **Available** lifecycle state.

Related Topics

- GetAutonomousContainerDatabase
- Resource Tags

Clone a Standby Database

Follow these steps to clone a standby autonomous database on an Oracle Exadata Cloud@Customer system.

You can use the cloning feature to create a point-in-time copy of your Autonomous Database for purposes such as testing, development, or analytics.

Clone Types: The clone feature offers the full-clone option to create a database that includes the metadata and data from the source database.

Steps

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.



- 3. In the list of Autonomous Databases, click the display name of the primary database.
- 4. Under Resources, click Autonomous Data Guard.
- 5. In the list of standby databases, find the database that you want to clone, and then click the display name to view details.
- 6. From the More Actions drop-down list, select Create Clone.
- 7. On the Create Autonomous Database Clone page, provide the following information:

In the Clone Type section, select Full Clone.

Clone Source: You can clone the standby database only from a backup.

- **Clone from a backup**: Creates a clone from a database backup. If you choose this option, select one of the following options:
 - **Specify a timestamp**: Creates a point-in-time clone.
 - Select from a list of backups: Creates a clone using all data from the specified backup. To limit your list of backups to a specific date range, enter the starting date in the From field and the ending date in the To field.

Provide basic information for the Autonomous Database.

- **Choose a compartment**: Your current compartment is the default selection but you can select a different compartment in which to create the clone from the drop-down list.
- **Source database name**: The name of the source database displays in the read-only Source database name field.
- **Display name**: Enter a description or other information to identify the database clone. You can change the display name any time and it does not have to be unique. Avoid entering confidential information.
- **Database name**: Enter a database name for the clone that contains only letters and numbers, begins with a letter. Avoid entering confidential information.
- Three additional fields are displayed if you opt to clone from a backup.
 - Exadata Infrastructure: You can choose to create the database clone in the same Exadata Infrastructure where the source database resides, or you can choose a different compartment by clicking CHANGE COMPARTMENT and choosing one from the drop-down list.
 - Autonomous Exadata VM Cluster: You can choose to create the database clone in the same Autonomous Exadata VM Cluster where the source database resides, or you can choose a different compartment by clicking CHANGE COMPARTMENT and choosing one from the drop-down list.
- Autonomous Container Database: You can choose to create the database clone in the same compartment and container database as the source database, or you can choose a different compartment by clicking CHANGE COMPARTMENT, and a different container database by choosing one from the drop-down list.
- Configure the database: Free Instance: Toggle the Free instance button on, if you want to create an Oracle Autonomous Database for Developers instance. ECPU count and Storage (GB) are auto-populated with 4 and 32 respectively because Oracle Autonomous Database for Developers comes fixed at 4 ECPUs and 32GB storage Compute auto-scaling is disabled because developer database instances do not support manual or auto-scaling.



- As developer database instances can only be created on ECPUbased ACDs without Autonomous Data Guard, the Free instance toggle button is disabled for ACDs with OCPU, Autonomous Data Guard, or both.
- To successfully clone a backup to an Oracle Autonomous
 Database for Developers instance, the current storage allocation of the backup database must be 32GB.

Configure the database

CPU Count: There is a minimum requirement of 1 OCPU or 4 ECPUs for an Autonomous Database clone from Backup.
Specify the number of OCPU for your database. The total number of cores available to all databases within the Autonomous Exadata Infrastructure depends on the infrastructure shape and what is already allocated to other Autonomous Databases.

The time taken to clone an Autonomous Database depends on the CPU Count and the network bandwidth between the Backup Destination and the target Autonomous Container Database.

The CPU type, that is, OCPU or ECPU is determined by the parent Autonomous Exadata VM Cluster resource's compute type.

For databases that need 2 or more ECPUs, you must specify the number of assigned ECPUs in increment of 1.

Note:

CPU Overprovisioning is not allowed with ECPUs.

For OCPUs, you can assign a fractional OCPU value from 0.1 to 0.9 (in increments of 0.1 OCPU) to databases that do not need a full OCPU. For databases that need 1 or more OCPUs, you must specify the number of assigned OCPUs as an integer. For example, you cannot assign 3.5 OCPUs to a database. The next available number of OCPUs above 3 is 4.

Databases with CPU over-provisioning can only connect to the tp and low services for the Autonomous Database for Transaction Processing and Mixed Workloads workloads. In the case of an Autonomous Database for Analytics and Data Warehousing workloads, you only connect to the low services when created on over-provisioned CPUs.

- **Storage (GB)**: Specify the amount of storage, in GB, that you want to make available to your cloned Autonomous Database, and it depends on the storage available to use. For full clones, the size of the source database determines the minimum amount of storage you can make available.
 - Default: 1024 GB
 - Minimum: 32 GB



- Increment: 1 GB
- **Auto scaling**: Enabling auto-scaling allows the system o automatically use up to three times more CPU and I/O resources to meet the workload demand.

Create administrator credentials

Set the password for the Autonomous Database administrator user by entering a password that meets the following criteria.

- Password cannot be one of the three most recently used passwords of the source database
- Between 12 and 30 characters long
- · Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

Use this password when accessing the service console and when using a SQL client tool.

Configure Network Access

You can change the access control list to enable or disable database-level access control or add or modify entries to the access control list.

- Click Modify Access Control.
- Select the Enable database level access control check box.
- Click Access Control Rule.

Note: The database-level access control will be enabled without any IP addresses in the access control list. Enabling an access control list with an empty list of IP addresses makes the database inaccessible to all clients.

- Specify the following types of addresses in your list by using the IP notation type drop-down selector:
 - IP Address allows you to specify one or more individual public IP addresses. Use commas to separate your addresses in the input field.
 - CIDR Block allows you to specify one or more ranges of public IP addresses using CIDR notation. Use commas to separate your CIDR block entries in the input field.

Advanced Options:

- Encryption Key:
 - Clone from a database instance: The source and the target ACD must be the same Keystore type. When the source is OKV, the target must also be the same OKV destination.
 - Clone from a backup: The source and the target ACDs can be different Keystore types. When the source is OKV, the target must also be the same OKV destination.
- Management: Choose a Character Set and National Character from the drop-down list.



• **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.

8. Click Create Autonomous Database Clone.

The Console displays the details page for the new clone of your database and the service begins provisioning the Autonomous Database. Note the following:

- The new clone displays the **Provisioning** lifecycle state until the provisioning process completes.
- The source database remains in the **Available** lifecycle state.
- Backups associated with the source database are not cloned for either the fullclone or the metadata-clone option.

The **Clone source** is displayed in the **General Information** section of the cloned database details page. Click the name to view details of the source database. Note that if the source database is deleted, then this key/value pair is not displayed.

Related Topics

Resource Tags

Clone a Standby Database Backup

Follow these steps to clone an autonomous database backup on an Oracle Exadata Database Service on Cloud@Customer system.

You can use the cloning feature to create a point-in-time copy of your Autonomous Database for purposes such as testing, development, or analytics. To clone only the database schema of your source database, choose the metadata clone option.

Clone Types

The clone feature offers the following two types of Autonomous Database clones:

- The full-clone option creates a database that includes the metadata and data from the source database.
- The metadata-clone option creates a database that includes only the metadata from the source database.

Steps

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- **3.** In the list of Autonomous Databases, click the display name of the primary database.
- 4. Under Resources, click Autonomous Data Guard.
- 5. In the list of standby databases, find the database that you want to clone, and then click the display name to view details.
- 6. Under Resources, click Backups.



- 7. In the list of backups, find the backup that you want to clone, click the action icon (three dots), and then click **Create Clone**.
- 8. On the Create Autonomous Database Clone page, provide the following information:

In the Clone Type section, select Full Clone.

Clone Source: The clone source section displays the source backup details.

Provide basic information for the Autonomous Database.

- **Choose a compartment**: Your current compartment is the default selection but you can select a different compartment in which to create the clone from the drop-down list.
- **Source database name**: The name of the source database displays in the read-only Source database name field.
- **Display name**: Enter a description or other information to identify the database clone. You can change the display name any time and it does not have to be unique. Avoid entering confidential information.
- **Database name**: Enter a database name for the clone that contains only letters and numbers, begins with a letter. Avoid entering confidential information.
- **Exadata Infrastructure**: You can choose to create the database clone in the same Exadata Infrastructure where the source database resides, or you can choose a different compartment by clicking **CHANGE COMPARTMENT** and choosing one from the drop-down list.
- Autonomous Exadata VM Cluster: You can choose to create the database clone in the same Autonomous Exadata VM Cluster where the source database resides, or you can choose a different compartment by clicking CHANGE COMPARTMENT and choosing one from the drop-down list.
- Autonomous Container Database: You can choose to create the database clone in the same compartment and container database as the source database, or you can choose a different compartment by clicking CHANGE COMPARTMENT, and a different container database by choosing one from the drop-down list.

Note:

When the target Autonomous Exadata VM Cluster is the same as the source, then the database name cannot be the same as the source database name.

Configure the database: Free Instance: Toggle the Free instance button on, if you want to create an Oracle Autonomous Database for Developers instance. ECPU count and Storage (GB) are auto-populated with 4 and 32 respectively because Oracle Autonomous Database for Developers comes fixed at 4 ECPUs and 32GB storage Compute auto-scaling is disabled because developer database instances do not support manual or auto-scaling.



- As developer database instances can only be created on ECPUbased ACDs without Autonomous Data Guard, the Free instance toggle button is disabled for ACDs with OCPU, Autonomous Data Guard, or both.
- To successfully clone a backup to an Oracle Autonomous
 Database for Developers instance, the current storage allocation of the backup database must be 32GB.

Configure the database

• **CPU Count**: Specify the number of CPUs for your clone database. The total number of CPUs available to all databases within the Autonomous Exadata Infrastructure depends on the infrastructure shape and what is already allocated to other Autonomous Databases.

After the clone, you can resize it to a lower value if needed. You can even resize to a fractional OCPU value to databases that do not need a full CPU. This allows you to overprovision CPU and run more databases on each infrastructure instance. Note that fractional CPU applies to OCPU only.

Note:

The time taken to clone an ADB depends on the CPU Count and the network bandwidth between the Backup Destination and the target ACD.

The selected CPU count is validated against a list of provisionable CPUs, and if the database can not be scaled up to the chosen CPU count, you will be provided with the two nearest provisionable CPU values.

You can use the *GetAutonomousContainerDatabase* API to get a complete list of provisionable CPU values.

There is a minimum requirement of 1 OCPUs or 4 ECPUs for an Autonomous Database clone from Backup.

The CPU type, that is, OCPU or ECPU is determined by the parent Autonomous Exadata VM Cluster resource's compute type.

For databases that need 2 or more ECPUs, you must specify the number of assigned ECPUs in increment of 1.

Note:

CPU Overprovisioning is not allowed with ECPUs.

For OCPUs, you can assign a fractional OCPU value from 0.1 to 0.9 (in increments of 0.1 OCPU) to databases that do not need a full OCPU. For databases that need 1 or more OCPUs, you must specify the number of



assigned OCPUs as an integer. For example, you cannot assign 3.5 OCPUs to a database. The next available number of OCPUs above 3 is 4.

Databases with CPU over-provisioning can only connect to the tp and low services for the Autonomous Database for Transaction Processing and Mixed Workloads workloads. In the case of an Autonomous Database for Analytics and Data Warehousing workloads, you only connect to the low services when created on overprovisioned CPUs.

- Storage (GB): Specify the amount of storage, in GB, that you want to make available to your cloned Autonomous Database, and it depends on the storage available to use.
 - Default/Minimum: Allocated storage of the source database
 - Increment: 1 GB
- **Auto scaling**: Enabling auto-scaling allows the system o automatically use up to three times more CPU and I/O resources to meet the workload demand.

Create administrator credentials

Set the password for the Autonomous Database administrator user by entering a password that meets the following criteria.

- Password cannot be one of the three most recently used passwords of the source database
- Between 12 and 30 characters long
- Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

Use this password when accessing the service console and when using a SQL client tool.

Configure Network Access

You can change the access control list to enable or disable database-level access control or add or modify entries to the access control list.

- Click Modify Access Control.
- Select the Enable database level access control check box.
- Click Access Control Rule.

Note: The database-level access control will be enabled without any IP addresses in the access control list. Enabling an access control list with an empty list of IP addresses makes the database inaccessible to all clients.

- Specify the following types of addresses in your list by using the IP notation type drop-down selector:
 - IP Address allows you to specify one or more individual public IP addresses. Use commas to separate your addresses in the input field.
 - CIDR Block allows you to specify one or more ranges of public IP addresses using CIDR notation. Use commas to separate your CIDR block entries in the input field.

Advanced Options:



- Encryption Key:
 - Clone from a database instance: The source and the target ACD must be the same Keystore type. When the source is OKV, the target must also be the same OKV destination.
 - Clone from a backup: The source and the target ACDs can be different Keystore types. When the source is OKV, the target must also be the same OKV destination.
- Management: Choose a Character Set and National Character from the drop-down list.
- **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.

9. Click Create Autonomous Database Clone.

The Console displays the details page for the new clone of your database and the service begins provisioning the Autonomous Database. Note the following:

- The new clone displays the **Provisioning** lifecycle state until the provisioning process completes.
- The source database remains in the **Available** lifecycle state.

Related Topics

- GetAutonomousContainerDatabase
- Resource Tags

Terminate an Autonomous Database

Follow these steps to terminate an Autonomous Database on an Oracle Exadata Database Service on Cloud@Customer system.

Note:

If the standby ACD is in snapshot standby mode, then you cannot delete an ADB in the primary ACD.

WARNING:

Terminating an Autonomous Database permanently deletes it. The database data will be lost when the system is terminated. However, automatic backups are not deleted if you have chosen Recovery Appliance or NFS as a backup destination. You can delete automatic backups directly from the Recovery Appliance or NFS.



- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to terminate.
- 4. From the More Actions drop-down list, select Terminate.
- 5. Confirm that you wish to terminate your Autonomous Database in the confirmation dialog.
- 6. Click Terminate Autonomous Database.

API to Manage Autonomous 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.

The following table lists the REST API endpoints to manage Autonomous Databases.

Operation	REST API Endpoint
Create an Autonomous Database	CreateAutonomousDatabase
View a list of Autonomous Databases	ListAutonomousDatabases
View details of an Autonomous Database	GetAutonomousDatabase
View a list of character sets supported by Autonomous Database.	ListAutonomousDatabaseCharacterSets
Set the password of an Autonomous Database's ADMIN user	UpdateAutonomousDatabase
Scale the CPU core count or storage of an Autonomous Database	UpdateAutonomousDatabase
Enable or disable auto scaling for an Autonomous Database	UpdateAutonomousDatabase
Move an Autonomous Database to another compartment	ChangeAutonomousDatabaseCompartment
Stop or start an Autonomous Database	StartAutonomousDatabase
Stop or start an Autonomous Database	StopAutonomousDatabase
Restart an Autonomous Database	RestartAutonomousDatabase
Back up an Autonomous Database manually	CreateAutonomousDatabaseBackup
View the list of Autonomous Database backups	ListAutonomousDatabaseBackups
Restore an Autonomous Database	RestoreAutonomousDatabase
Clone an Autonomous Database	CreateAutonomousDatabase
Terminate an Autonomous Database	DeleteAutonomousDatabase

Monitor Performance with Autonomous Database Metrics

You can monitor the health, capacity, and performance of your Autonomous Databases with metrics, alarms, and notifications. You can use Oracle Cloud Infrastructure console or Monitoring APIs to view metrics.



- View Top Six Metrics for an Autonomous Database Displays the top six metrics that are available in the metrics section on the Autonomous Database details page.
- View Aggregated Metrics for Autonomous Databases in a Compartment Learn to view aggregated metrics for Autonomous Databases in a compartment.
- Autonomous Database Metrics and Dimensions
 You can limit the instances where you see metrics with dimensions. The available
 dimensions include: workload type, instance display name, region, and the
 instance OCID.

View Top Six Metrics for an Autonomous Database

Displays the top six metrics that are available in the metrics section on the Autonomous Database details page.

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 other tools). See Getting Started with Policies for information on policies.

Perform the following prerequisite steps as necessary:

- Open the Oracle Cloud Infrastructure console by clicking the hamburger menu next to Oracle Cloud.
- From the Oracle Cloud Infrastructure left navigation list click Oracle Databases > Exadata Cloud@Customer.
- On the Autonomous Databases page select an Autonomous Database from the links under the **Name** column.

To view metrics for an Autonomous Database instance:

- 1. On the Autonomous Database Details page, under Resources, click Metrics.
- 2. There is a chart for each metric. In each chart, you can select the **Interval and Statistic** or use the default values.
- 3. To create an alarm on a metric, click **Options** and select **Create an Alarm on this Query**.

See Managing Alarms for information on setting and using alarms.

For more information about metrics see Database Metrics.

You can also use the Monitoring API to view metrics. See Monitoring API for more information.

View Aggregated Metrics for Autonomous Databases in a Compartment

Learn to view aggregated metrics for Autonomous Databases in a compartment.

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 other tool). See Getting Started with Policies for information on policies

Perform the following prerequisite steps as necessary:

• Open the Oracle Cloud Infrastructure console by clicking the hamburger menu next to Oracle Cloud.



 From the left navigation list click Solutions and Platform > Monitoring > Service Metrics.

To use the metrics service to view Autonomous Database metrics:

- 1. On the Service Metrics page, under **Compartment** select your compartment.
- 2. On the Service Metrics page, under Metric Namespace select oci_autonomous_database.
- 3. If there are multiple Autonomous Databases in the compartment you can show metrics aggregated across the Autonomous Databases by selecting **Aggregate Metric Streams**.
- 4. If you want to limit the metrics you see, next to **Dimensions** click **Add** (click **Edit** if you have already added dimensions).
 - a. In the **Dimension Name** field select a dimension.
 - **b.** In the **Dimension Value** field select a value.
 - c. Click Done.
 - **d.** In the Edit dimensions dialog click +Additional Dimension to add an additional dimension. Click x to remove a dimension.

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.

Autonomous Database Metrics and Dimensions

You can limit the instances where you see metrics with dimensions. The available dimensions include: workload type, instance display name, region, and the instance OCID.

Use dimensions by selecting values in the Oracle Cloud Infrastructure Console Service Metrics page or by setting dimension values with the API. See View Aggregated Metrics for Autonomous Databases in a Compartment to view metrics and to select metric dimensions.

See Database Metrics for a list of the database metrics and dimensions.

Connecting to Autonomous Databases

Applications and tools connect to an autonomous database using Oracle Net Services (also known as SQL*Net). Oracle Net Services enables a network session from a client application to an Oracle Database server.

When a network session is established, Oracle Net Services acts as the data courier for both the client application and the database. It is responsible for establishing and maintaining the connection between the client application and the database, as well as exchanging messages between them. It supports a variety of connection types to autonomous databases, including:

- Oracle Call Interface (OCI), which is used by many applications written in C language. Examples include Oracle utilities such as Oracle SQL*Plus, SQL*Loader, and Oracle Data Pump.
- ODBC drivers, which can be used by applications running on Microsoft Windows, are layered over Oracle Call Interface (OCI).
- JDBC OCI, which can be used by Java language applications, is layered over Oracle Call Interface (OCI). The Oracle SQLcl command-line interface uses JDBC OCI.



 JDBC Thin Driver, also for Java applications, is a pure Java driver. Oracle SQL Developer supports JDBC Thin Driver connections.

Third-party products and custom applications can use any of these connection types.

Oracle Autonomous Database provides several pairs of database services to use when connecting to autonomous databases. In each pair, one of the pair provides a secure TCP (TCPS) connection using the TLS protocol, and the other provides a TCP connection. In all other respects, the two members of a pair are the same. To ensure security of data in transit, Oracle strongly recommends that you use a secure connection, even if the database is only available through a private network. If you are familiar with using an Oracle Database within your own data center, you may not have previously used these secure connections.

To provide the secure TCPS connection, certification authentication uses an encrypted key stored in a *wallet* on both the client (where the application is running) and the server (where the autonomous database is running). The key on the client must match the key on the server to make a connection. A wallet contains a collection of files, including the key and other information needed to connect to your database . All communications between the client and the server are encrypted.

- Download the Wallet for an Autonomous Database
- Get the APEX and SQL Developer Web URLs for an Autonomous Database

Download the Wallet for an Autonomous Database

Follow these steps to download the wallet for an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- **3.** In the list of Autonomous Databases, click the display name of the database whose wallet you wish to download.
- 4. Click DB Connections.
- 5. Select the **DB Connection** option.
- 6. In the Download Wallet dialog, enter a wallet password in the **Password** field and confirm the password in the **Confirm Password** field.

The password must be at least 8 characters long and must include at least 1 letter and either 1 numeric character or 1 special character.

Note:

This password protects the downloaded Client Credentials wallet. This wallet is not the same as the Transparent Data Encryption (TDE) wallet for the database; therefore, use a different password to protect the Client Credentials wallet.

7. Click **Download** to save the client security credentials zip file.

By default the filename is Wallet_databasename.zip. You can save this file as any filename you want.

You must protect this file to prevent unauthorized database access.

The zip file includes the following:

- tnsnames.ora and sqlnet.ora: Network configuration files storing connect descriptors and SQL*Net client side configuration.
- cwallet.ora and ewallet.p12: Auto-open SSO wallet and PKCS12 file. PKCS12 file is protected by the wallet password provided in the UI.
- keystore.jks and truststore.jks: Java keystore and truststore files. They are protected by the wallet password provided while downloading the wallet.
- ojdbc.properties: Contains the wallet related connection property required for JDBC connection. This should be in the same path as tnsnames.ora.

Note:

Wallet files, along with the Database user ID and password, provide access to data in your autonomous database. Store wallet files in a secure location. Share wallet files only with authorized users. If wallet files are transmitted in a way that might be accessed by unauthorized users (for example, over public email), transmit the wallet password separately and securely.

Get the APEX and SQL Developer Web URLs for an Autonomous Database

Follow these steps to get the URLs to use to connect to APEX (Oracle Application Express) and Oracle SQL Developer Web in an autonomous database on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- In the list of Autonomous Databases, click the display name of the database whose APEX and SQL URLs you wish to get.
- 4. Click DB Connections.
- 5. Select the Application Connection option.
- 6. Application URLs are displayed in plain text in the **Application URL** field. Copy the URL string using the **Copy** link.

Paste the URL into a browser running on a system with network access to your autonomous database.

Authenticate Autonomous Database Users with Kerberos

Describes how to configure Kerberos to authenticate Oracle Autonomous Database users.

- About Kerberos Authentication
- Components of the Kerberos Authentication System
- Notes about Kerberos Authentication on Autonomous Database



- Enable Kerberos Authentication on Autonomous Database
- Disable Kerberos Authentication on Autonomous Database

About Kerberos Authentication

You can configure Autonomous Database on Dedicated Exadata Infrastructure to use Kerberos network authentication protocol to authenticate database users. Kerberos is a strong network authentication protocol. It uses secret-key cryptography to enable strong authentication by providing user-to-server authentication.

- Autonomous Database on Dedicated Exadata Infrastructure support for Kerberos provides the benefits of single sign-on and centralized authentication of Oracle users. Kerberos is a trusted third-party authentication system that relies on shared secrets. It presumes that the third party is secure, and provides single sign-on capabilities, centralized password storage, database link authentication, and enhanced PC security. It does this through a Kerberos authentication server.
- The Kerberos system revolves around the concept of a ticket. A ticket is a set of electronic information that identifies a user or a service. A ticket identifies you and your network access privileges.
- In Kerberos-based authentication, you transparently send a request for a ticket to a Key Distribution Center (KDC). The Key Distribution Center authenticates you and grants you a ticket to access the database.

Components of the Kerberos Authentication System

Provides an overview of the Kerberos authentication system.

- A realm establishes an authentication administrative domain. Each realm has its own Kerberos database which contains the users and services for that particular administrative domain.
- Tickets are issued by the Key Distribution Center (KDC). Clients present tickets to the Database Server to demonstrate the authenticity of their identity. Each ticket has expiration and a renewal time.
- Keytabs stores long-term keys for one or more principals. A keytab file is generated by invoking the tool kadmin.local (for MIT Key Distribution Center) or ktpass (for Active Directory Key Distribution Center).
- Principals are the entries in the Key Distribution Center database. Each user, host
 or service is given a principal. A principal is a unique identity to which the Key
 Distribution Center can assign tickets.
- Kerberos support in Autonomous Database uses these values for various components that make up a service principal's name:

Components of Service Principal	Value in Autonomous Database
kinstance	The value of the host parameter found in the TNS connect string.



Components of Service Principal	Value in Autonomous Database
kservice	You can retrieve GUID for your Autonomous Database using a query on the GUID column of V\$PDBS. It must be in uppercase when used as the value of kservice.
REALM	Any realm supported by your KDC. REALM must always be in uppercase.

To enable Kerberos authentication for your Autonomous Database, you must keep your Kerberos configuration files (krb.conf) and service key table file (v5srvtab) ready. For more information on these files and steps to obtain them, please see About Kerberos Configuration Files.

Notes about Kerberos Authentication on Autonomous Database

Before proceeding with Kerberos authentication on Autonomous Database on Dedicated Exadata Infrastructure, review the following notes:

- You can use Kerberos authentication only with Autonomous Database versions 19.18 or later.
- If you enable Kerberos authentication for your Autonomous Database, you can still use password-based database authentication for your database.
- Only one external authentication method can be used for your Autonomous Database at any time. That is, you can only have Oracle Cloud Infrastructure (IAM), Centrally Managed User with Active Directory (CMU-AD), Azure AD, or Kerberos authentication schemes enabled at any time.

Note:

The only exception is that Kerberos authentication can be configured on top of CMU-AD to provide CMU-AD Kerberos authentication for Microsoft Active Directory users.

- Kerberos authentication is not supported for the following tools:
 - Oracle Database API for MongoDB
 - Oracle REST Data Services
 - Oracle Machine Learning
 - APEX
 - Oracle Graph Studio
 - Oracle Database Actions
- You can enable Kerberos authentication to authenticate the ADMIN user. You can use the Reset Password functionality on the Oracle Cloud Infrastructure (OCI) Console to reset the ADMIN user's password and regain access if a corrupted keytab file causes ADMIN user's authentication to fail.
- Kerberos authentication is supported only with TCPS protocol.



 Kerberos authentication is not supported with DB_LINKs and databases with Autonomous Data Guard.

Enable Kerberos Authentication on Autonomous Database

Follow the below-listed steps to enable Kerberos authentication for Autonomous Database on Dedicated Exadata Infrastructure:

 As there can only be one external authentication scheme enabled for an Autonomous Database at any given time, run the DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION procedure to disable any external authentication scheme that is already enabled for your database.

To run the procedure, you must be logged in as ADMIN user or have the EXECUTE privilege on DBMS CLOUD ADMIN.

```
BEGIN
    DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION;
END;
/
```

- 2. Obtain the Kerberos configuration files: krb.conf and the service key table file v5srvtab. For more information on these files and steps required to obtain them, see Configuring Kerberos Authentication.
- 3. Copy the Kerberos configuration files krb.conf and v5srvtab to a bucket in your Object Store.

Note:

Oracle recommends that you store the Kerberos configuration files in a private bucket in your Object Store.

If you are using Oracle Cloud Infrastructure Object Store, see Putting Data into Object Storage for details on uploading files.

4. To enable Kerberos as the external authentication for your Autonomous Database, run the DBMS_CLOUD_ADMIN.ENABLE_EXTERNAL_AUTHENTICATION procedure and pass in a location URI with the params JSON argument. You must place the configuration files krb.conf and v5srvtab in the Object Storage location specified in the location_uri parameter.

For example:



```
END;
/
```

To provide CMU-AD Kerberos authentication for Microsoft Active Directory users, you can enable Kerberos authentication on top of CMU-AD by setting type to CMU in the above example.

In this example, namespace-string is the Oracle Cloud Infrastructure object storage namespace and bucketname is the bucket name. See Understanding Object Storage Namespaces for more information.

The credential name you use in this step is the credentials for the Object Store.

If the location_uri is a pre-authenticated URL then supplying a credential_name is not required.

This creates a directory object named KERBEROS_DIR in your database and uses the credential to download the Kerberos configuration files from the Object Store location to the directory object.

5. After you enable Kerberos authentication, remove the configuration krb.conf and v5srvtab from Object Store. You can use local Object Store methods to remove these files or use DBMS CLOUD.DELETE OBJECT to delete the files from Object Store.

See Navigate to Oracle Cloud Infrastructure Object Storage and Create Bucket for more information on Object Storage.

See ENABLE_EXTERNAL_AUTHENTICATION Procedure for more information.

Disable Kerberos Authentication on Autonomous Database

Before enabling any other external authentication scheme on your database, you must disable the Kerberos authentication running the DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION procedure.

To run this procedure, you must be logged in as ADMIN user or have the EXECUTE privilege on DBMS CLOUD ADMIN.

```
BEGIN
    DBMS_CLOUD_ADMIN.DISABLE_EXTERNAL_AUTHENTICATION;
END;
/
```

See DISABLE_EXTERNAL_AUTHENTICATION Procedure for more information.



To provide CMU-AD Kerberos authentication for the Microsoft Active Directory users, you must continue to configure CMU-AD authentication without disabling Kerberos authentication.

Patching ADB on Exadata Cloud@Customer Infrastructure

- Overview of ADB on Exadata Cloud@Customer Infrastructure Patching
- Specifying When Maintenance Can Occur
- Specifying What Kind of Patches to Apply

Overview of ADB on Exadata Cloud@Customer Infrastructure Patching

ADB-Dedicated maintenance involves patching Exadata Infrastructure(EI) Autonomous VM Cluster, and Autonomous Container Database (ACD).

- Autonomous Exadata Infrastructure/Exadata Infrastructure (EI): IB Switches, Storage cells, Dom0 operating system, and so on.
- Autonomous VM clusters (AVM): Grid Infrastructure (GI), and Guest VM operating system patching.
- Autonomous Container Database (ACD): Oracle Database patching.

The patching is done every quarter. Although there is no order dependency between EI/AVM/ACD, it is preferred to patch in the sequence EI, AVM, and then ACD.

Oracle schedules and performs all patching and other maintenance operations on all dedicated Exadata infrastructure resources. You can also specify when such maintenance operations can occur, and what kind of patching is performed.

Specifying When Maintenance Can Occur

In general, Oracle schedules and performs entire fleet maintenance spread throughout each quarter. You can let Oracle handle maintenance scheduling, or you can set a specific maintenance window when Oracle can begin maintenance operations. You set this maintenance window at the Autonomous Exadata Infrastructure level, and it applies to all Autonomous Container Databases and Autonomous Databases created in the Autonomous Exadata Infrastructure resource as well as to the resource itself.

Additionally, you can set a maintenance window for each individual Autonomous Container Database, thereby setting the window for all of its contained Autonomous Databases.



Oracle recommends that you set a maintenance window for at least Exadata Infrastructure resources. Doing so will prevent maintenance operations from occurring at times that would be disruptive to regular database operations.

You can set the maintenance window for an Exadata Infrastructure and Autonomous Container Database resources when you create them or you can set or change it later. Once a maintenance activity is scheduled based on the maintenance window you set, you can manage the actual timing of the activity, even to the point of changing the patch version, applying the patch immediately, or skipping the activity.

Related Topics

- Using the Console to Create Exadata Cloud@Customer Infrastructure To create your Exadata Cloud@Customer infrastructure, be prepared to provide values for the fields required for configuring the infrastructure.
- 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.
- Patch and Update an Exadata Database Service on Cloud@Customer System Learn to update and patch the Exadata Database Service on Cloud@Customer System
- Create an Autonomous Container Database
- Change the Maintenance Schedule of an Autonomous Container Database

Specifying What Kind of Patches to Apply

One standard maintenance operation is to apply database software patches to your Autonomous Container Databases and, by extension, the Autonomous Databases created in them. By default, Oracle applies quarterly Release Updates (RUs). Currently, the Release Update Revision (RUR) maintenance type is not supported.

To help you decide whether to have Oracle apply RUs or RURs to a given Autonomous Container Database, see My Oracle Support Note 2285040.1.

Related Topics

https://support.oracle.com/epmos/faces/DocContentDisplay?id=2285040.1

Migrating On-Premises Databases to Oracle Autonomous Database on Exadata Cloud@Customer

Zero Downtime Migration supports migrations to Oracle Autonomous Database on Exadata Cloud@Customer from any on-premises Oracle Database, including existing Exadata Cloud@Customer systems, using logical migration method and NFS as a data transfer medium.

Supported usecases:



- Logical Online Migration
 For more information, see Logical Online Migration.
- Logical Offline Migration
 For more information, see Logical Offline Migration.

Prerequisites: For more information, see *Preparing for a Logical Database Migration* and *Migrating to Oracle Autonomous Database on Exadata Cloud@Customer.*

Related Topics

- Logical Online Migration
- Logical Offline Migration
- Preparing for a Logical Database Migration
- Migrating to Oracle Autonomous Database on Exadata Cloud@Customer

Using Autonomous Data Guard with Autonomous Database on Exadata Cloud@Customer

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

- Enabling Autonomous Data Guard on an Autonomous Container Database When you enable Data Guard, a separate Data Guard association is created for the primary and the standby database.
- Enabling Autonomous Data Guard on an Autonomous Database Autonomous Databases inherit Data Guard settings from the parent container database.
- Maintenance Scheduling and Patching Data Guard Enabled Autonomous
 Container Database
 Follow these steps to change the maintenance schedule of a Data Guard enabled
 Autonomous Container Database.

Enabling Autonomous Data Guard on an Autonomous Container Database

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

Note:

Replication of data happens only over the client network.

• Create an Autonomous Data Guard Enabled Autonomous Container Database Follow these steps to create an Autonomous Data Guard Enabled Autonomous Container Database on an Oracle Exadata Cloud@Customer system.



View Details of a Data Guard Enabled Primary or Standby Autonomous Container
 Database

Follow these steps to view detailed information about a primary or standby Autonomous Container Database on an Oracle Exadata Database Service on Cloud@Customer system.

Edit Autonomous Container Database Backup Settings

If automatic backups were disabled while provisioning an Autonomous Container Database (ACD), you can enable them later from the Oracle Cloud Infrastructure (OCI) console.

- Convert a Physical Standby ACD to Snapshot Standby ACD
 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 ACD to Physical Standby ACD A snapshot standby database will automatically convert back to a physical standby database after 7 days.
- Rotate CDB Encryption Key
- Managing a Standby Autonomous Container Database
 Enabling Autonomous Data Guard on an Autonomous Container Database creates a standby (peer) Autonomous Container Database that provides data protection, high availability, and facilitates disaster recovery for the primary database.
- Perform a Failover to Standby Autonomous Container Database Initiate a failover operation by using the Data Guard association of the standby database.
- Perform a Switchover to Standby or Primary Autonomous Container Database Initiate a switchover operation by using the Data Guard association of the primary database.
- Reinstate Data Guard Enabled Standby Autonomous Container 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.
- Terminate a Data Guard Enabled Primary Autonomous Container Database Follow these steps to terminate an autonomous container database on an Oracle Exadata Cloud@Customer system.
- Terminate a Data Guard Enabled Standby Autonomous Container Database Follow these steps to terminate an autonomous container database on an Oracle Exadata Cloud@Customer system.
- Operations Performed Using the APIs Learn how to use the API to manage Autonomous Data Guard Enabled Autonomous Container Database.

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.



Create an Autonomous Data Guard Enabled Autonomous Container Database

Follow these steps to create an Autonomous Data Guard Enabled Autonomous Container Database on an Oracle Exadata Cloud@Customer system.

Note:

For better management and sharing of the underlying SGA/memory resources, Oracle recommends that all Autonomous Databases configured for In-Memory be in the same Autonomous Container Database.

Minimum Resource Requirements

To create one Autonomous Container Database, you need at least:

- 2 OCPUs or 8 ECPUs
- 50 GB local storage
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. Click Create Autonomous Container Database.

The Create Autonomous Container Database page is displayed.

- 4. Provide the following basic information:
 - **Compartment**: Choose the compartment in which your autonomous container database will be created.
 - **Display Name**: Enter a user-friendly description or other information that helps you easily identify the autonomous container database. The display name does not have to be unique. Avoid entering confidential information.
- 5. Select the Autonomous Exadata VM Cluster you wish to use to create your autonomous container database.

Note:

If the selected Autonomous Exadata VM Cluster does not have 2 available OCPUs or 8 available ECPUs per node, which is the minimum requirement for creating an Autonomous Container Database, then this field is greyed out. Select an Autonomous Exadata VM Cluster that has enough resources to create an autonomous container database.

- 6. Choose Container DB software version.
 - Select version from base images: Create a database with the Oracle Database version selected.
 - Select base image: The latest version is selected by default. Select a database version (N, N-1) if needed.



- 7. Under Configure Autonomous Data Guard, select the Enable Autonomous Data Guard checkbox and provide the following details.
 - **Peer Autonomous Container Database Compartment**: Choose the compartment in which your standby autonomous container database will be created.
 - **Display Name**: Enter a user-friendly description or other information that helps you easily identify the autonomous container database. The display name does not have to be unique. Avoid entering confidential information.
 - Select peer Autonomous Exadata VM Cluster: Specify the following values for the standby:
 - Peer Region: Select a peer region.
 - Peer Exadata: Select the Exadata Cloud@Customer infrastructure where the standby database will be created. Click the CHANGE COMPARTMENT hyperlink to choose a compartment.
 - Peer Autonomous Exadata VM Cluster: Select the Autonomous VM Cluster in which the standby ACD must be created. Click the CHANGE COMPARTMENT hyperlink to choose a compartment.

The primary and standby Autonomous Container Databases must be on two different Autonomous VM Clusters on the same Exadata infrastructure or different Exadata infrastructures.

Note:

If the selected Autonomous Exadata VM Cluster does not have 2 available OCPUs per node, which is the minimum requirement for creating an Autonomous Container Database, then this field is greyed out. Select an Autonomous Exadata VM Cluster that has enough resources to create an Autonomous Container Database.

- **Data Protection Mode**: Specify the protection mode used for this Data Guard association.
 - Maximum Performance: Provides the highest level of data protection that is possible without compromising the availability of a primary database.
 - Maximum Availability: Provides the highest level of data protection that is
 possible without affecting the performance of a primary database. This is the
 default protection mode.

See Oracle Data Guard Concepts and Administration for more information about Oracle Data Guard Protection Modes.

• **Enable automatic failover**: Select this checkbox to enable automatic failover and set the FSFO lag limit.

Fast-Start Failover (FSFO) lag limit: Set Fast-Start Failover (FSFO) lag limit in increments of 1. **Minimum:** 5 and **Maximum:** 3600 seconds. **Default:** 30 seconds.

Note:

FSFO Lag Limit is applicable only to Maximum Performance protection mode.

8. Optionally, you can configure an automatic maintenance schedule.



- a. Click Edit Maintenance Preferences.
- b. Configure Container Database maintenance version
 - Next Release Update (RU): Update to the next release update in the next maintenance cycle.
 - Latest Release Update (RU): Update to the latest release update in the next maintenance cycle.

For more information, see Management Operations.

c. To configure the maintenance schedule, select **Specify a schedule**.

Choose your preferred month, week, weekday, and start time for autonomous container database maintenance.

- 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.
- Under **Day of the week**, specify the day of the week on which the maintenance will occur.
- Under **Start hour**, specify the hour during which the maintenance run will begin.
- d. Click Save Changes.
- 9. Enable automatic backups.

By default, automatic backups are enabled for an ACD. Optionally, you can choose to disable them by deselecting the Enable automatic backups check box.

While provisioning an ACD with Autonomous Data Guard, you can not disable the automatic backups.

Note:

If disabled for an ACD, automatic backups can be enabled anytime later from the Oracle Cloud Infrastructure (OCI) console by following the steps outlined in Edit Autonomous Container Database Backup Settings. However, once enabled you can not disable automatic backups for the ACD.

If enabling automatic backups fail for some reason, the ACD provisioning also fails with an error message. As a workaround, you can provision the ACD with automatic backups disabled, and enable them from the ACD's Details page later.

10. Select a Backup Destination Type:

Note:

The backup destination type can only be set while enabling automatic backups on an ACD and cannot be changed later.



The possible options are:

 Object Storage: Stores backups in an Oracle-managed object storage container on Oracle Cloud Infrastructure.

If you choose Object Storage as the type, you can optionally specify an internet HTTP proxy to use when connecting to the storage container. Oracle recommends using a proxy when possible for enhanced security.

 Network File System (NFS): Stores backups in a Network File System (NFS) storage location.
 If you choose Network File System (NES) as the type, select a previously define

If you choose Network File System (NFS) as the type, select a previously defined Backup Destination that uses Network File System (NFS) storage.

Recovery Appliance: Stores backups in one of your previously defined backup destinations that use Oracle Zero Data Loss Recovery Appliance.
 If you choose Recovery Appliance as the type, select a previously defined Backup Destination that uses Oracle Zero Data Loss Recovery Appliance, the DB_UNIQUE_NAME of the Autonomous Container Database, and the VPC user name password.

Provide the connection string that connects to the recovery appliance in an Oracle "easy connect" string format, that is, <host>:<port>/<service name>, where <host> is the SCAN hostname of the Zero Data Loss Recovery Appliance.

- **11**. The following Advanced Options are available:
 - Backup retention period: Specify a Backup retention period value to meet your needs. You can choose any value between 7 to 95 days.
 For the Object Storage and Network File System (NFS) backup destination types, the backup retention policy value defaults to 30 days.

For the Recovery Appliance backup destination type, this value is controlled by the Recovery Appliance protection policy.

All the backups are automatically deleted after the backup retention period.

 Encryption Key: Choose an encryption option, Encrypt using Oracle-managed keys or Encrypt using customer-managed keys. The default option is Oraclemanaged keys.

To use customer-managed keys, select the **Encrypt using customer-managed keys** option, 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.



- Autonomous Container Databases and Autonomous Databases only support 256-bit Hardware Security Module (HSM) Vault keys.
- Validate OKV Key encryption post restart: OKV TDE Master Key is validated every time you start or restart your ACD. 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 remains in OKV and will not be deleted.
- Management: Choose a Character Set and National Character from the drop-down list.
- Database In-memory:
 - Enable database In-memory: It requires at least four OCPUs and a
 percentage of the System Global Area (SGA) to enable in-memory. If you
 enable In-memory, select the percentage of SGA to allocate to the IM
 Column Store. In-memory may have an impact on the autonomous
 database's performance if a large amount of memory is allocated or if it is
 disabled.

Note:

If you enable In-memory in a primary database with Data Guard enabled, the configuration will be replicated to the standby database as read-only.

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. 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. 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. Setting the appropriate application tag key in the

Oracle-ApplicationName tag namespace while deploying an Autonomous Database, can ensure that the database is configured and ready for the particular application, for example, PeopleSoft out of the box.

For more information, see *Tagging Oracle Exadata Database Service on Cloud@Customer Resources*.

12. Click Create Autonomous Container Database.

Related Topics

- Oracle Data Guard Protection Modes
- Management Operations
- Resource Tags
- 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.

View Details of a Data Guard Enabled Primary or Standby Autonomous Container Database

Follow these steps to view detailed information about a primary or standby Autonomous Container Database on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- **3.** In the list of Autonomous Container Databases, click the display name of the database you wish to view details.
- 4. In the Autonomous Container Database Details page, check the Autonomous Data Guard association status and peer database state.
- To change the protection mode and Fast-Start Failover (FSFO) lag limit of the primary database, select Update Autonomous Data Guard from the More Actions drop-down list.
 - In the resulting Update Autonomous Data Guard dialog, make changes and click **Save Changes**.
- 6. Under Resources, click Autonomous Data Guard to view association details.

Edit Autonomous Container Database Backup Settings

If automatic backups were disabled while provisioning an Autonomous Container Database (ACD), you can enable them later from the Oracle Cloud Infrastructure (OCI) console.

- **1.** Go to the Details page of the Autonomous Container Database whose backup settings you want to change.
- 2. Under More actions, click Edit Backup Settings.



You can also edit the backup settings by clicking the **Edit** link under the **Backup** section on the Autonomous Container Database Information tab.

The Edit Backup Settings dialog opens.

- If automatic backups are disabled for this ACD, enable them by selecting the Enable automatic backups checkbox, and choose appropriate values for the following settings:
 - **Backup destination type:** Select a Backup destination type and then specify options based on the selected type.

Note:

The backup destination type can only be set while enabling automatic backups on an ACD and cannot be changed later.

The possible options are:

- Object Storage: Stores backups in an Oracle-managed object storage container on Oracle Cloud Infrastructure.
 If you choose Object Storage as the type, you can optionally specify an internet HTTP proxy to use when connecting to the storage container.
 Oracle recommends using a proxy when possible for enhanced security.
- Network File System (NFS): Stores backups in a Network File System (NFS) storage location.

If you choose Network File System (NFS) as the type, select a previously defined Backup Destination that uses Network File System (NFS) storage.

 Recovery Appliance: Stores backups in one of your previously defined backup destinations that use Oracle Zero Data Loss Recovery Appliance. If you choose Recovery Appliance as the type, select a previously defined Backup Destination that uses Oracle Zero Data Loss Recovery Appliance, the DB_UNIQUE_NAME of the Autonomous Container Database, and the VPC user name password.

Provide the connection string that connects to the recovery appliance in an Oracle "easy connect" string format, that is, <host>:<port>/<service name>, where <host> is the SCAN hostname of the Zero Data Loss Recovery Appliance.

 Backup retention period (in days): Specify a Backup retention period value to meet your needs. You can choose any value between 7 to 95 days.
 For the Object Storage and Network File System (NFS) backup destination types, the backup retention policy value defaults to 30 days.

For the Recovery Appliance backup destination type, this value is controlled by the Recovery Appliance protection policy.

All the backups are automatically deleted after the backup retention period.

4. Click Save Changes.

Convert a Physical Standby ACD to Snapshot Standby ACD

A snapshot standby database is a fully updateable standby database created by converting a physical standby database into a snapshot standby database.

A snapshot standby database receives and archives redo data received from a primary database. Redo data received from a primary database, however, is not applied. The redo data is applied when a snapshot standby database is converted back into a physical standby database after discarding all local updates to the snapshot standby database. Once the conversion is complete, the role of the standby ACD will change to SNAPSHOT_STANDBY. You can perform DDL and DML operations on all ADBs in the SNAPSHOT_STANDBY ACD.

Note:

- A snapshot standby will automatically convert back to physical standby after 7 days.
- Automatic Backups will not be taken on a snapshot standby.
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the physical standby database you are interested in.

Autonomous Container Database Details page is displayed.

4. Select Convert to snapshot standby from the More Actions drop-down list.

Note:

You cannot convert a physical standby to a snapshot standby with Automatic Failover enabled. Disable Automatic Failover to convert your standby database to snapshot standby mode.

- 5. To disable automatic failover, do the following:
 - a. Under Resources, click Autonomous Data Guard.
 - b. Click the name of the peer (primary) database.

Primary ACD details page is displayed.

- c. Select Update Autonomous Data Guard from the More Actions drop-down list.
- d. In the resulting Update Autonomous Data Guard dialog, clear the **Enable automatic** failover checkbox.
- e. Click Save Changes.
- 6. After disabling automatic failover, continue with converting your physical standby to snapshot standby.

Review the warning message on the Convert to Snapshot Standby window.



Convert to Snapshot Standby supports two options:

- Use New Database Services: Connect to snapshot standby database using new services that are active only in snapshot standby mode.
- Use Primary Database Services: Connect to the snapshot standby database using the same services in the primary database.

Selecting the **Use Primary Database Services** option will display an additional warning message about configuring proper connection strings to connect to primary and snapshot standby databases respectively to avoid incorrect database connections.

7. Click Convert.

After the conversion, the role of the standby database changes to **Snapshot Standby**, and the **Convert to Physical Standby** option becomes available in the **More Actions** drop-down list.

Note:

Convert to physical standby ACD will only be enabled when ACD is in SNAPSHOT STANDBY mode.

- Converting to physical standby ACD will discard all local updates from all ADBs and apply redo data from primary ACD.
- Converting to physical standby will revert the standby ACD role and all its ADBs roles to STANDBY.

Convert a Snapshot Standby ACD to Physical Standby ACD

A snapshot standby database will automatically convert back to a physical standby database after 7 days.

The system displays a banner with the actual date when the snapshot standby database will convert back to a physical standby database. Database role on all the ADBs in that ACD will also change accordingly. Banner message about automatic conversion will be available only on ACDs.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the snapshot standby database you are interested in.

Autonomous Container Database Details page is displayed.

4. Select Convert to physical standby from the More Actions drop-down list.



If you convert your snapshot standby Autonomous Container Database to physical standby, all local updates to your snapshot standby will be discarded and data from your primary Autonomous Container Database will be applied.

5. Click Convert.

Rotate CDB Encryption Key

Follow these steps to rotate the TDE Master key. On key rotation, the ACD life cycle goes through the regular updating state and returns to available.

You can rotate the TDE Master key as many times as you want. The new TDE Master Key is stored in the same wallet in which the previous key was stored. Rotating the TDE Master Key leads to the new key being generated in OKV and assigned to this database. You can view all of the keys in OKV.

Note:

You can rotate both Oracle-managed and customer-managed encryption keys.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the primary or standby database you wish to view details.
- 4. On the Autonomous Container Database Details page, click Rotate Encryption Key.
- 5. On the Rotate Encryption Key dialog, click **Rotate Encryption Key**.

Managing a Standby Autonomous Container Database

Enabling Autonomous Data Guard on an Autonomous Container Database creates a standby (peer) Autonomous Container Database that provides data protection, high availability, and facilitates disaster recovery for the primary database.

If the primary Autonomous Container Database becomes unavailable because of a region failure, a failure of the Autonomous Exadata Infrastructure, or the failure of the Autonomous Container Database, itself, then it automatically fails over to the standby Autonomous Container Database if Automatic Failover is enabled. The role of the failed primary database becomes *Disabled Standby* and, after a brief period of time, the standby database assumes the role of the primary database.

Besides hardware failures and regional outages, the following table lists database health conditions that also trigger automatic failover:



Table 6-8 Database Health Condition	Table 6-8	Database	Health	Condition
-------------------------------------	-----------	----------	--------	-----------

Database Health Condition	Description
Datafile Write Errors	The system initiates a fast-start failover if write errors occur in any data files, including temp files, system data files, and undo files.
Corrupted Dictionary	Dictionary corruption of a critical database. Currently, this state can be detected only when the database is open.
Corrupted Controlfile	Controlfile is permanently damaged because of a disk failure.

- After automatic failover concludes, a message displays on the details page of the disabled standby database advising you that failover has occurred.
- Automatic failover is optional while configuring Autonomous Data Guard. You can enable or disable automatic failover after configuring Autonomous Data Guard.
- The FastStartFailoverLagLimit configuration attribute establishes an acceptable limit, in seconds, up to which the standby database can fall behind the primary database, with respect to the redo applied. If the limit is reached, then a fast-start failover does not occur. This attribute is used when fast-start failover is enabled and the configuration is operating in maximum performance mode.

The FastStartFailOverLagLimit attribute:

- Has a default value of 30 seconds
- Cannot be configured
- Is only applicable when in maximum performance protection mode

After the service resolves the issues with the former primary Autonomous Container Database, you can perform a manual switchover to return both databases to their initial roles.

Once you provision the standby database, you can perform various management tasks related to the standby database, including:

- Manually switching over a primary database to a standby database
- Manually failing over a primary database to a standby database
- Reinstating a primary database to standby role after failover
- Terminating a standby database



Perform a Failover to Standby Autonomous Container Database

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

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- **3.** In the list of Autonomous Container Databases, click the display name of the infrastructure resource you are interested in.
- 4. Click the name of the standby database or snapshot standby associated with the primary Autonomous Container Database that you want to failover.

The system displays a warning when you perform a failover operation when the standby ACD is in snapshot standby mode:

WARNING:

Your standby database is in a snapshot standby role. Failover will convert your snapshot standby database to physical standby by discarding all local updates to your snapshot standby and applying data from your primary database.

- 5. Click Failover.
- 6. In the Confirm Manual Failover to Standby dialog box, enter the name of the Autonomous Container Database you want to failover, and then click **Failover**.

Alternatively,

- a. Under **Resources**, click Autonomous Data Guard to display a list of peer databases for the primary database you are managing.
- **b.** For the Data Guard association on which you want to perform a failover, click the Actions icon (three dots), and then click **Failover**.
- c. In the Confirm Manual Failover to Standby dialog box, enter the name of the Autonomous Container Database you want to failover, and then click **Failover**.

Note:

After successful completion of failover, the standby ACDs role will change to primary and the primay's role will become disabled-standby.

Perform a Switchover to Standby or Primary Autonomous Container Database

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

- You can perform Switchover only when both primary and standby are in available state.
- You cannot perform switchover if patching or maintenance is in progress on primary or standby.
- You cannot perform switchover operation If the standby ACD is in snapshot standby mode.



- After the switchover, the maintenance preference for the new standby and primary will remain the same as the old standby and primary.
- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the infrastructure resource you are interested in.
- 4. Click the name of the primary or secondary database.
- 5. Click Switchover.
- 6. In the confirmation dialog box, click Switchover.

Alternatively,

- a. Under Resources, click Data Guard Associations.
- **b.** For the Data Guard association on which you want to perform a switchover, click the Actions icon (three dots), and then click **Switchover**.
- c. In the Confirm Switchover to Standby dialog box, click Swichover.

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

Reinstate Data Guard Enabled Standby Autonomous Container 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 the operations team correct the cause of failure, you can reinstate the failed database as a functioning standby for the current primary by using its Data Guard association. you can reinstate the failed database as a functioning standby for the current primary by using its Data Guard association.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the infrastructure resource you are interested in.
- 4. Click the name of the **Disabled Standby** database.
- 5. Under Resources, click Autonomous Data Guard..
- 6. For the Data Guard association on which you want to reinstate this database, click the Actions icon (three dots), and then click **Reinstate**.
- 7. In the Reinstate Database dialog box, click Reinstate.

This database should now be reinstated as the standby in the Data Guard association. You can now perform a switchover operation to revert the respective databases to their original roles.



Terminate a Data Guard Enabled Primary Autonomous Container Database

Follow these steps to terminate an autonomous container database on an Oracle Exadata Cloud@Customer system.

You must terminate all Autonomous Databases within a container database before you can terminate the container database itself. Terminating Autonomous Container Database will disable Autonomous Data Guard, which affects high availability and disaster recovery of your Autonomous Databases.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the infrastructure resource you are interested in.
- 4. In the Autonomous Container Database Details page, select **Terminate** from the More Actions drop-down list.
- 5. Click Terminate.
- 6. In the confirmation dialog, type the name of the Autonomous Container Database, and then click **Terminate Autonomous Container Database**.

Terminate a Data Guard Enabled Standby Autonomous Container Database

Follow these steps to terminate an autonomous container database on an Oracle Exadata Cloud@Customer system.

You can terminate a standby Autonomous Container Database even if there are standby Autonomous Databases inside it. However, you cannot terminate standby Autonomous Databases inside a standby Autonomoous Container Database. To terminate a standby Autonomoous Database, you must first terminate the primary Autonomous Database. Terminating Autonomous Container Database will disable Autonomous Data Guard, which affects high availability and disaster recovery of your Autonomous Databases.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the infrastructure resource you are interested in.
- 4. In the Autonomous Container Database Details page, select **Terminate** from the More Actions drop-down list.
- 5. Click Terminate.
- 6. In the confirmation dialog, type the name of the Autonomous Container Database, and then click **Terminate Autonomous Container Database**.

Operations Performed Using the APIs

Learn how to use the API to manage Autonomous Data Guard Enabled Autonomous Container Database.



For information about using the API and signing requests, see "REST APIs" and "Security Credentials". For information about SDKs, see "Software Development Kits and Command Line Interface".

The following table lists the REST API endpoints to manage Autonomous Data Guard Enabled Autonomous Container Database.

Operation	REST API Endpoint
Creates Autonomous Container Databases (update to the existing API)	CreateAutonomousContainerDatabase
View details of the specified Autonomous Container Database	GetAutonomousContainerDatabase
View a list of Autonomous Container Databases with Autonomous Data Guard associations	ListAutonomousContainerDatabaseDatag uardAssociations
Fetch details of an Autonomous Container Database Autonomous Data Guard associations	GetAutonomousContainerDatabaseDatagu ardAssociation
Fail over the standby Autonomous Container Database identified by the autonomousContainerDatabaseId parameter to the primary Autonomous Container Database after the existing primary Autonomous Container Database fails or becomes unreachable.	FailoverAutonomousContainerDatabaseD ataguardAssociation
Switch over the primary Autonomous Container Database of an Autonomous Data Guard peer association to standby role. The standby Autonomous Container Database associated with autonomousContainerDatabaseDataguard AssociationId assumes the primary Autonomous Container Database role.	SwitchoverAutonomousContainerDatabas eDataguardAssociation
Reinstate a disabled standby Autonomous Container Database identified by the autonomousContainerDatabaseId parameter to an active standby Autonomous Container Database.	ReinstateAutonomousContainerDatabase DataguardAssociation
View a list of the Autonomous Data Guard- enabled databases associated with the specified Autonomous Database.	ListAutonomousDatabaseDataguardAssoc iations
Fetch details of an Autonomous Database Autonomous Data Guard associations	GetAutonomousDatabaseDataguardAssoci ation
Fetches details such as peer database role, lag time, transport lag, and state	GetAutonomousDatabase

Related Topics

- REST APIs
- Security Credentials
- Software Development Kits and Command Line Interface



Enabling Autonomous Data Guard on an Autonomous Database

Autonomous Databases inherit Data Guard settings from the parent container database.

- View Autonomous Data Guard Enablement Autonomous Data Guard settings are configured on the Autonomous Container Databases in which the databases are running.
- Create an Autonomous Data Guard Enabled Autonomous Database
- View Details of a Data Guard Enabled Primary or Standby Autonomous Database Follow these steps to view detailed information about a primary or standby Autonomous Database on an Oracle Exadata Database Service on Cloud@Customer system.
- Rotate ADB Encryption Key

View Autonomous Data Guard Enablement

Autonomous Data Guard settings are configured on the Autonomous Container Databases in which the databases are running.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Container Databases.

This page displays if an autonomous database is Data Guard enabled or not, and if enabled, then the role of the database in the Data Guard association.

Create an Autonomous Data Guard Enabled Autonomous Database

Follow these steps to create an autonomous database on an Oracle Exadata Cloud@Customer system.

Note:

You cannot create an Autonomous Database for Developers instance on an Autonomous Data Guard-enabled container database.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. Click Create Autonomous Database.
- 4. In the Create Autonomous Database dialog, enter the following:

Basic Database Information

- **Compartment**: Select the compartment of the Autonomous Database.
- **Display Name**: A user-friendly description or other information that helps you easily identify the resource. The display name does not have to be unique. Avoid entering confidential information.
- **Database Name**: The database name must consist of letters and numbers only, starting with a letter. Avoid entering confidential information.



Workload Type

Select the desired workload type. See *Autonomous Data Warehouse* and *Autonomous Transaction Processing* for information about each workload type.

Autonomous Container Database: Select the Autonomous Data Guardenabled Autonomous Container Databases checkbox, and then select an Autonomous Container Database.

Compartment: Specify the compartment containing the Autonomous Container Database you wish to use.

Database CPU Core Count and Storage Configuration

 CPU Count: The total number of cores available to all databases within the Autonomous Exadata Infrastructure depends on the infrastructure shape and what is already allocated to other Autonomous Databases.
 The CPU type, that is, OCPU or ECPU is determined by the parent Autonomous Exadata VM Cluster resource's compute type.

The selected CPU count is validated against a list of provisionable CPUs, and if the database can not be scaled up to the chosen CPU count, you will be provided with the two nearest provisionable CPU values.

Based on the resource utilization on each node; not all the values of the available CPUs can be used to provision or scale Autonomous Databases. For example, suppose you have 20 CPUs available at the AVMC level, not all the values from 1 to 20 CPUs can be used to provision or scale Autonomous Databases depending on the resource availability at the node level. The list of CPU values that can be used to provision or scale an Autonomous Database is called **Provisionable CPUs**.

On the console, when you try to provision or scale an Autonomous Database, the CPU count will be validated against the list of provisionable CPUs, and if the value is not provisionable, you will be provided with the two nearest provisionable CPU values. Alternatively, if you want to see the complete list of provisionable CPU values for an Autonomous Exadata VM Cluster, you can use the following API:

GetAutonomousContainerDatabase returns a list of provisionable OCPU values that can be used to create a new Autonomous Database in the given Autonomous Container Database. See *GetAutonomousContainerDatabase* for more details.

For ECPUs, this value defaults to 2 ECPUs. For databases that need 2 or more ECPUs, you must specify the number of assigned ECPUs in increment of 1.

Note:

CPU Overprovisioning is not allowed with ECPUs.

For OCPUs, the default value is 1 OCPU. However, you can assign a fractional OCPU value from 0.1 to 0.9 (in increments of 0.1 OCPU) to databases that do not need a full OCPU. This allows you to over-provision CPU and run more databases on each infrastructure instance. For databases that need 1 or more OCPUs, you must specify the number of assigned OCPUs



as an integer. For example, you cannot assign 3.5 OCPUs to a database. The next available number of OCPUs above 3 is 4.

Databases with CPU over-provisioning can only connect using tp and low services.

Deselect Auto Scaling to disable auto-scaling. By default, auto-scaling is enabled to allow the system to automatically use up to three times more CPU and IO resources to meet workload demand.

• **Storage (TB)**: Specify the storage you wish to make available to your Autonomous Database, in terabytes. The available storage depends on the infrastructure shape and what is already consumed by other Autonomous Databases.

Administrator Credentials

Set the password for the Autonomous Database Admin user by entering a password that meets the following criteria. You use this password when accessing the Autonomous Database service console and when using an SQL client tool.

- Contains from 12 to 30 characters
- Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

Configure network access

You can optionally create an ACL during database provisioning, or at any time thereafter.

- a. Click Modify Access Control.
- **b.** In the Edit Access Control List dialog, select the **Enable database level access control** checkbox.
- c. Under the Primary database access control list, specify the following types of addresses in your list by using the IP notation type drop-down selector:
 IP Address allows you to specify one or more individual public IP addresses. Use commas to separate your addresses in the input field.

CIDR Block allows you to specify one or more ranges of public IP addresses using CIDR notation. Use commas to separate your CIDR block entries in the input field.

 Under Standby database access control, do the following: (Default) Same as primary database: Leave as is if you want the same access control list for the secondary database.

Define standby database access control: Initialized with the same details as primary. Add or modify entries, as applicable.

Advanced Options

Tags: Optionally, you can 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, skip this option (you can apply tags later) or ask your administrator. Avoid entering confidential information.

Encryption Key: ADB inherits encryption settings from the parent ACD. If the parent ACD is configured for customer-managed OKV-based encryption, then the child ADB will

ORACLE

also have TDE Master Key generated and managed in the same OKV wallet used to store ACD master keys. Additionally, any backups taken on the Autonomous Database will have the OKV-based key associated with it.

5. Click Create Autonomous Database.

Note:

The following naming restrictions apply to Autonomous Transaction Processing and Autonomous Data Warehouse databases:

- Names associated with databases terminated within the last 60 days cannot be used when creating a new database.
- A database name cannot be used concurrently for both an Autonomous Data Warehouse and an Autonomous Transaction Processing database.

Related Topics

- About Autonomous Database
- GetAutonomousContainerDatabase
- Resource Tags

View Details of a Data Guard Enabled Primary or Standby Autonomous Database

Follow these steps to view detailed information about a primary or standby Autonomous Database on an Oracle Exadata Database Service on Cloud@Customer system.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Databases, click the display name of the database you wish to view details.
- 4. In the Autonomous Database Details page, check the Autonomous Data Guard association status and peer database state.
- 5. Under **Resources**, click **Autonomous Data Guard** to view association details.

Rotate ADB Encryption Key

Follow these steps to rotate the TDE Master key. On key rotation, the ADB life cycle goes through the regular updating state and returns to available.

You can rotate the TDE Master key as many times as you want. The new TDE Master Key is stored in the same wallet in which the previous key was stored. Rotating the TDE Master Key leads to the new key being generated in OKV and assigned to this database. You can view all of the keys in OKV.



You can rotate both Oracle-managed and customer-managed encryption keys.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Cloud@Customer.
- 2. Click Autonomous Databases.
- **3.** In the list of Autonomous Databases, click the display name of the database you wish to view details.
- 4. On the Autonomous Database Details page, from the **More Actions** drop-down list, select **Rotate Encryption Key**.
- 5. On the Rotate Encryption Key dialog, click Rotate Encryption Key.

Maintenance Scheduling and Patching Data Guard Enabled Autonomous Container Database

Follow these steps to change the maintenance schedule of a Data Guard enabled Autonomous Container Database.

- Configure Automatic Maintenance Schedule for a Data Guard Enabled Autonomous Container Database
- View the Next Scheduled Maintenance Run of a Data Guard Enabled Autonomous Container Database
- View the Maintenance History of a Data Guard Enabled Autonomous Container Database
- Immediately Patch a Data Guard Enabled Autonomous Container Database
- Reschedule or Skip scheduled Maintenance for Data Guard Enabled Autonomous Container Database

Configure Automatic Maintenance Schedule for a Data Guard Enabled Autonomous Container Database

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the container database you are interested in.
- 4. On the Autonomous Container Database details page, click **Edit Maintenance Preferences**.

In the Edit Automatic Maintenance dialog that opens, you can configure both the maintenance schedule and the patch type.



The standby database will have **No preference** by default. Standby Maintenance depends on the primary maintenance schedule.

5. Optionally, you can change the maintenance patch type. To edit this setting, select either **Release Update (RU)** or **Release Update Revision (RUR)**.

Release Update (RU): Autonomous Database installs only the most current release update.

Release Update Revision (RUR): Autonomous Database installs the release update plus additional fixes.

Note:

Standby will be always patched before primary and the default gap between standby and primary is 7 days. You have have an option to change the default gap to anytime between 1 - 7 days.

Configure Container Database maintenance version

- Next Release Update (RU): Update to the next release update in the next maintenance cycle.
- Latest Release Update (RU): Update to the latest release update in the next maintenance cycle.
- 6. To configure the maintenance schedule, select Specify a schedule in the Configure the automatic maintenance schedule section. Choose your preferred month, week, weekday, and start time for container database maintenance.
 - Under Maintenance months, specify at least one month for each maintenance quarter during which you want Autonomous Exadata Infrastructure maintenance to occur.

Note:

Maintenance quarters begin in February, May, August, and November, with the first maintenance quarter of the year beginning in February.

- 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.
- Under Day of the week, specify the day of the week on which the maintenance will occur.
- Under Start hour, specify the hour during which the maintenance run will begin.



- Choose the buffer period between primary and standby maintenance execution. Buffer period is the number of days before which the standby Autonomous Container Database Maintenance will be scheduled before primary Autonomous Container Database Maintenance
- 7. Click Save Changes.

View the Next Scheduled Maintenance Run of a Data Guard Enabled Autonomous Container Database

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the container database you are interested in.
- 4. On the Autonomous Container Database details page, under **Maintenance**, click the **View** link in the **Next Maintenance** field.
- 5. On the Maintenance page, under Autonomous Database Maintenance, click Maintenance.

In the list of maintenance events, you can the details of scheduled maintenance runs. Maintenance event details include the following:

- The status of the scheduled maintenance run
- The type of maintenance run (quarterly software maintenance or a critical patch)
- The OCID of the maintenance event
- The start time and date of the maintenance

View the Maintenance History of a Data Guard Enabled Autonomous Container Database

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the container database you are interested in.
- 4. On the Autonomous Container Database details page, under **Maintenance**, click the **View** link in the **Next Maintenance** field.
- 5. On the Maintenance page, under Autonomous Database Maintenance, click Maintenance History.

In the list of past maintenance events, you can click on an individual event title to read the details of the maintenance that took place. Maintenance event details include the following:

- The category of maintenance (quarterly software maintenance or a critical patch)
- Whether the maintenance was scheduled or unplanned
- The OCID of the maintenance event
- The start time and date of the maintenance



Immediately Patch a Data Guard Enabled Autonomous Container Database

Note:

Patching primary immediately will result in standby being patched first, if standby is not already patched.

- 1. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer.
- 2. Click Autonomous Databases.
- 3. In the list of Autonomous Container Databases, click the display name of the Autonomous Container Database that you want to patch.
- 4. On the Autonomous Container Database Details page, in the **Maintenance** section, click the **View** link in the **Next Maintenance** field to display the Maintenance page for the Autonomous Container Database that you want to patch.
- 5. In the Autonomous Container Database section, click Patch Now in the Scheduled Start Time field to display the Run Maintenance dialog.
- 6. Click **Patch Now** to start the patching operation.

Reschedule or Skip scheduled Maintenance for Data Guard Enabled Autonomous Container Database

Note: Skipping primary will skip standby also. If standby is patched, then skipping on primary is not allowed. Open the navigation menu. Under Oracle Database, click Exadata Database Service on Cloud@Customer. Click Autonomous Databases. In the list of Autonomous Container Databases, click the display name of the container database that you want to manage. On the Autonomous Container Database details page, in the Maintenance section, click the View link in the Next Maintenance field. On the Maintenance page, any container database maintenance events planned for the next 15 days will appear in the list of maintenance events.

To skip scheduled maintenance for a container database, click Skip.



You cannot skip scheduled maintenance more than twice, consecutively.

To reschedule maintenance, click **Edit** and enter a start time for the update in the Edit Maintenance dialog. Ensure that your specified container database maintenance window is later in the quarter than your scheduled Exadata infrastructure maintenance.

Using Performance Hub

Use Performance Hub to monitor database activity, diagnose issues, and tune queries to improve the performance of Oracle databases.

For more information, see:

- What's New in Performance Hub
- About Performance Hub
- Performance Hub Features
- Using Performance Hub to Analyze Database Performance



7

Reference Guides for Exadata Database Service on Cloud@Customer

- Using the dbaascli Utility with Exadata Database Service on Cloud@Customer Learn to use the dbaascli utility on Exadata 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 Exadata Database Service on Cloud@Customer in the Monitoring Service Review the metrics emitted by the 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 Exadata Cloud@Customer resources emit events, which are structured messages that indicate changes in resources.
- Policy Details for Exadata Database Service on Cloud@Customer Learn to write policies to control access to 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.
- Security Guide for Exadata Database Service on Cloud@Customer Systems This guide describes security for an Oracle Exadata Cloud@Customer System. It includes information about the best practices for securing the Oracle Exadata Cloud@Customer System.
- Troubleshooting 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 Exadata Database Service on Cloud@Customer

Learn to use the dbaascli utility on Exadata Cloud@Customer.

• About Using the dbaascli Utility on Exadata Database Service on Cloud@Customer You can use the dbaascli utility to perform various database lifecycle and administration operations on Exadata 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 Exadata Database Service on Cloud@Customer Software Images Using the Dbaascli Utility You can list and download the Oracle database software images on an 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 Exadata Database Service on Cloud@Customer

You can use the dbaascli utility to perform various database lifecycle and administration operations on Exadata 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.



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.

- 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



```
2.IMAGE_TAG=19.8.0.0.0-NC
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 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.



- 1. Connect to the virtual machine as the opc user. For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
- 2. Start a root user command shell:

sudo -s



3. Run the following command:

dbaascli dbhome create --version *Oracle Home Version* --imageTag *image Tag Value*

Where:

- --version specifies the Oracle Database version
- --imageTag specifies the Image Tag of the image to be used

For example:

dbaascli dbhome create --version 19.9.0.0.0

Note:

Specifying imageTag is optional. To view the Image Tags, refer to command dbaascli cswlib showImages. Image Tags are typically same as the version of the database. However, it is kept as a provision for cases where multiple images may need to be released for the same version - each catering to a specific customer requirement.

4. Exit the root user command shell:

exit

For more details on advanced supported options, see dbaascli dbhome create.

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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



Databases created on a subset of nodes will not be displayed in the OCI console.

- Create Oracle Database version 12.1.0.2 or higher with the release update JAN 2021 or higher. For databases with lower versions, it is recommended to use the OCI Console based API.
- 1. Connect to the virtual machine as the opc user. For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
- 2. Start a root user command shell:

sudo -s

3. Run the following command:

dbaascli database create --dbName *database name* --oracleHome *Oracle Home Path*

Where:

- --dbName specifies the name of the database
- --oracleHome specifies Oracle home location

To create a CDB, run the following command:

dbaascli database create --dbName *database name* --oracleHome *Oracle Home Path*

To create a non-CDB, run the following command:

```
dbaascli database create --dbName database name --oracleHome Oracle
Home Path --createAsCDB false
```

When prompted, enter the sys and tde passwords.

4. Exit the root user command shell:

exit

For more details on advanced supported options, see dbaascli database create.

- Running Prerequisite Checks Prior to Creating Oracle Database To run 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 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. For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
- 2. Start a root user command shell:

sudo -s

3. Run the following command:

```
dbaascli database create --dbName database name --oracleHome Oracle Home Path --executePrereqs
```

Where:

- --dbName specifies the name of the database
- --oracleHome specifies the Oracle home location
- 4. Exit the root user command shell:

exit

For more details on advanced supported options, see dbaascli database create.

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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.

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an 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.



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

Managing Exadata Database Service on Cloud@Customer Software Images Using the Dbaascli Utility

You can list and download the Oracle database software images on an Exadata Database Service on Cloud@Customer instance, which can then be used for provisioning a database home.



You can create custom database software images for your 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 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 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

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@dg1llrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
    showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file
location:
    /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
AM_46941.log
```

```
###########
```

17.IMAGE_TAG=18.17.0.0.0 VERSION=18.17.0.0.0 DESCRIPTION=18c JAN 2022 DB Image

18.IMAGE_TAG=19.10.0.0.0
VERSION=19.10.0.0.0
DESCRIPTION=19c JAN 2021 DB Image

- 19.IMAGE_TAG=19.11.0.0.0 VERSION=19.11.0.0.0 DESCRIPTION=19c APR 2021 DB Image
- 20.IMAGE_TAG=19.12.0.0.0 VERSION=19.12.0.0.0 DESCRIPTION=19c JUL 2021 DB Image
- 21.IMAGE_TAG=19.13.0.0.0 VERSION=19.13.0.0.0 DESCRIPTION=19c OCT 2021 DB Image

Images can be downloaded using their image tags. For details, see help using
'dbaascli cswlib download --help'.
dbaascli execution completed

Related Topics

- Connecting to a Virtual Machine with SSH
 You can connect to the virtual machines in an Exadata 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 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 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 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.
- 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):

- Connect to the virtual machine as the opc user. For detailed instructions, see Connecting to a Virtual Machine with SSH.
- 2. Start a root user command shell:

sudo -s

3. Run the following command:

```
dbaascli dbhome patch --oracleHome dbhome_path --targetVersion
Oracle Database version
```

Where:

- --oracleHome identifies the path of the Oracle home to be patched.
- --targetVersion specifies the target Oracle Database version to use for patching, specified as five numeric segments separated by periods (e.g. 19.12.0.0.).



For example:

```
dbaascli dbhome patch --oracleHome /u02/app/oracle/product/19.0.0.0/
dbhome 2 --targetVersion 19.9.0.0.0
```

4. Exit the root user command shell:

exit

For more details on advanced supported options, see dbaascli dbHome patch.

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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:

- Connect to the virtual machine as the opc user. For detailed instructions, see Connecting to a Virtual Machine with SSH.
- 2. Start a root user command shell:

sudo -s

3. Run the following command:

```
dbaascli database move --oracleHome path_to_target_oracle_home --
dbname database name
```

Where:

- --oracleHome identifies the path of the target Oracle home that uses the desired Oracle Database software version. Note that the target Oracle home must exist in your system prior to using the database move command.
- --dbname specifies the name of the database that is being moved.

For example:

```
dbaascli database move --oracleHome /u02/app/oracle/product/
19.0.0.0/dbhome_2 --dbname xyz
```



4. Exit the root user command shell:

exit

For more details on advanced supported options, see dbaascli database move.

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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 Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli grid patch To patch Oracle Grid Infrastructure to the specified minor version, use the dbaascli grid patch command.



Patching Oracle Grid Infrastructure (GI) Using GI Software Image

To patch Oracle Grid Infrastructure (GI) using GI software image, use this procedure.

Oracle Grid Infrastructure can also be patched by first creating a patched software image, and then using that image to perform the patching operation. This provides the advantage that an image can be created ahead of time outside of the patching window. It also helps in conflict resolution as any conflicts among the patches are highlighted during the image creation process without impacting the patching window.

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.

- Connect to the virtual machine as the opc user. For detailed instructions, see Connecting to a Virtual Machine with SSH.
- 2. Start a root user command shell:

sudo -s

3. Run the following command:

dbaascli cswlib showImages --product database

The command output lists the available database software images.

dbaascli cswlib showImages --product grid

The command output lists the available grid software images.

4. Exit the root user command shell:

exit



For more details on advanced supported options, see dbaascli cswlib showImages.

Example 7-3 dbaascli cswlib showImages

```
[root@dg1llrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
    showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file
location:
    /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
AM_46941.log
```

```
###########
```

17.IMAGE_TAG=18.17.0.0.0 VERSION=18.17.0.0.0 DESCRIPTION=18c JAN 2022 DB Image

18.IMAGE_TAG=19.10.0.0.0
VERSION=19.10.0.0.0
DESCRIPTION=19c JAN 2021 DB Image

- 19.IMAGE_TAG=19.11.0.0.0 VERSION=19.11.0.0.0 DESCRIPTION=19c APR 2021 DB Image
- 20.IMAGE_TAG=19.12.0.0.0 VERSION=19.12.0.0.0 DESCRIPTION=19c JUL 2021 DB Image
- 21.IMAGE_TAG=19.13.0.0.0 VERSION=19.13.0.0.0 DESCRIPTION=19c OCT 2021 DB Image

Images can be downloaded using their image tags. For details, see help using
'dbaascli cswlib download --help'.
dbaascli execution completed

Related Topics

- Connecting to a Virtual Machine with SSH
 You can connect to the virtual machines in an Exadata 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.

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 --executePreregs 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.

- 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).
- 4. Exit the root user command shell:

exit

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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

3. Run the following command:

```
dbaascli database move --oracleHome path_to_target_oracle_home --dbname
database_name --executePrereqs
```

Where:

- --oracleHome identifies the path of the target Oracle Home that uses the desired
 Oracle Database software version. Note that the target Oracle Home must exist in
 your system prior to using the database move command.
- --dbname specifies the name of the database that is being moved
- 4. Exit the root user command shell:

exit

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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.

4. Exit the root user command shell:

exit



Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.
- dbaascli grid patch To patch Oracle Grid Infrastructure to the specified minor version, use the dbaascli grid patch command.

Resuming or Rolling Back a Patching Operation

You can resume or revert a failed patching operation. Reverting a patch is known as a rollback.

- Resuming a Patch Operation
 To resume a patching operation, use the --resume flag with the original patching command.
- Rolling Back a Patch Operation
 Use the --rollback flag with the original patching command to roll back (revert) a
 patching operation.

Resuming a Patch Operation

To resume a patching operation, use the --resume flag with the original patching command.

- Connect to the virtual machine as the opc user.
 For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
- 2. Start a root user command shell:

sudo -s

3. Run the original patching command to resume a patching operation: For example:

```
dbaascli dbhome patch --oracleHome /u02/app/oracle/product/19.0.0.0/
dbhome 2 --targetVersion 19.9.0.0.0 --resume
```

4. Exit the root user command shell:

exit

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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*.
- 2. Start a root user command shell:

sudo -s

3. Run the original patching command to roll back (revert) a patching operation: For example:

dbaascli grid patch --targetVersion 19.11.0.0.0 --rollback

Note:

- Resume and Rollback operations are supported for Oracle Home patching, Oracle Grid Infrastructure patching, and database move operations.
- When resuming or rolling back a patching operation, you must run the resume or rollback command from the same node that was used to run the original patching command, and you must run the original command with the addition of the --resume or --rollback flag.
- 4. Exit the root user command shell:

exit

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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.

- Connect to a virtual machine as the opc user. For detailed instructions, see Connecting to a Virtual Machine with SSH.
- 2. Start a root user command shell:

sudo -s

3. To update to the latest available cloud tooling release, run the following command:

dbaascli admin updateStack

The command takes care of updating the cloud tooling release on all the nodes of the cluster.

For more details and other available options, refer to dbaascli admin updateStack --help.

Related Topics

- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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.

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
```

• 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> --executePreregs

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
- 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 24.1.2.0.0 (240327)
- 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
- Release 21.2.1.x.x

Release 24.1.2.0.0 (240327)

- 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

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.
- 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 rolechange.
- 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
 - 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
 - Added option to not open the PDB at the end of relocate. For details, see skipOpenPDB argument in dbaascli pdb relocate --help. After using this option, the pdb relocate can be completed by running the command using completePDBRelocate argument.
 - Added option to clean up the relocated PDB metadata/services at the source location. For details, see cleanupRelocatedPDB argument in dbaascli pdb delete --help
- New dbaascli commands
 - dbaascli database createTemplate This command can be used to create database templates (DBCA templates) that can subsequently be used to create databases. DBCA templates are widely used for creating a clone database with DBCA - a tool that is shipped with Oracle Database server software. For details, see dbaascli database createTemplate --help



- Introduced dbaascli tde rotateMasterKey to rotate the master key for database encryption. For details, see dbaascli tde rotateMasterKey --help. The command dbaascli tde rotate masterkey is now deprecated.
- Database Lifecycle related improvements
 - Added support to use dbca templates in database creation workflows. For details, see dbcaTemplateFilePath argument in dbaascli database create --help
 - Improved performance for duplicate database creation. For details on how to create duplicate database, see dbaascli database duplicate --help
 - Added support to create a duplicate database from a source database which is not TDE-encrypted.
- TDE management
 - Introduced dbaascli tde rotateMasterKey to rotate the master key for database encryption. For details, see dbaascli tde rotateMasterKey --help. The command dbaascli tde rotate masterkey is now deprecated.
 - Revamped workflow for all TDE operations. For details, see dbaascli tde --help
- Grid Infrastructure (GI) Patching
 - Added support to allow parallel execution of patching operation on nodes. This option needs to be carefully exercised as it results into reduced database availability.
- Database Backup and Recovery
 - Revamped workflow for creating database from standalone backups
- Includes AHF version 22.2.4
- Various bug fixes and stability improvements

Release 22.3.1.1.0 (221003)

- New dbaascli commands
 - dbaascli database getDetails This command shows the detailed information of a given database, for example, dbname, node information, pluggable databases information, and so on. For details, see dbaascli database getDetails --help.
- Pluggable Database Operations
 - Added support for creating pluggable databases as refreshable clone using refreshablePDB argument. For details, see dbaascli pdb remoteClone --help
- · Various bug fixes and stability improvements

Release 22.3.1.0.1 (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.
 - 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
- 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:
 - 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.
- Various bug fixes and stability improvements.

Release 21.3.1.2.0

- Improved the timing of dbaascli operations with enhanced Control Plane metadata synchronization logic.
- Enhanced dbaascli logs to have millisecond-level information along with the associated thread.
- Introduced more prerequisite checks in database home patching and database move operations to catch potential failures scenarios with suggestions to corrective action.
- Database patching operations now retain the state of the databases to be same as it was prior to patching. For pluggable databases, pdb saved state is honored.
- Various bug fixes and stability improvements.

Release 21.3.1.1.0

- Added support to unlock PDB Admin user account as part of PDB creation, localClone, or remoteClone operation. For details, see option -lockPDBAdminAccount in dbaascli pdb create --help.
- Fixed an issue that updates the database resource registered with Oracle Grid Infrastructure in existing environments with the correct value of database name.
- Enhanced PDB lifecycle operations.
- Various bug fixes and stability improvements.

Release 21.3.1.0.1

- Support for the following dbaascli commands to be run as oracle user.
 - dbaascli pdb bounce
 - dbaascli pdb close
 - dbaascli pdb connectString
 - dbaascli pdb create
 - dbaascli pdb delete
 - dbaascli pdb getDetails
 - dbaascli pdb list
 - dbaascli pdb localClone



- dbaascli pdb open
- dbaascli pdb remoteClone
- Revamped out-of-place patching of database. For details, see dbaascli database move -help.
- Timing related enhancements in Oracle Grid Infrastructure patching workflow. For details, see dbaascli grid patch -help.
- Deprecated the support for exadbcpatchmulti / dbaascli patch for patching operations. The dbaascli dbhome patch and dbaascli grid patch commands are provided for patching operation for database homes and Oracle Grid Infrastructure. Refer to the *Patching Oracle Grid Infrastructure* and *Oracle Database Using dbaascli* section for details. Also see, *dbaascli Command Reference* section.
- Deprecated the support for dbaascli tools patch command to bring consistency in the dbaascli command conventions. The new command is dbaascli admin updateStack. For details, see section Updating Cloud Tooling using dbaascli.
- Ability to run dbaascli in disconnected mode for long running operations. Executing dbaascli command with --waitForCompletion false gets you a job ID that can be queried later to get the status of the operation, using dbaascli job getStatus -jobid *job_id*. This is useful for long running operations where users may want to get the control back immediately after command execution. In this release, this option is available only for dbaascli database create command. More commands will be added in subsequent releases to have this support. The help output for those commands will reflect the support of --waitForCompletion option.
- Deprecated the support for dbaascli shell. It is recommended that users run the complete dbaascli commands on the command prompt as suggested in all the document examples. Execution of just dbaascli will show the output of its usage help instead of entering into a dbaascli shell.
- Various bug fixes and stability improvements.

Release 21.2.1.x.x

- Redesigned Oracle Grid Infrastructure patching operation and added ability to resume from failed point, patch on subset of nodes, instance draining, and other enhancements. For details, see dbaascli grid patch --help. Also refer to the Patching Oracle Grid Infrastructure and Oracle Database Using dbaascli section.
- Deprecated the support for exadbcpatchmulti / dbaascli patch for patching operations. dbaascli dbhome patch and dbaascli grid patch commands are provided for patching operation for database homes and Oracle Grid Infrastructure. Refer to the *Patching Oracle Grid Infrastructure and Oracle Database Using dbaascli* section for details. Also see, *dbaascli Command Reference* section.
- Deprecated the support for dbaascli tools patch command to bring consistency in the command conventions. The new command is dbaascli admin updateStack.
- Redesigned PDB management APIs for create, local clone, and remote clone operations. For details, see dbaascli pdb --help.
- Redesigned database delete API. For details, see dbaascli database delete --help.
- Revamped dbhome creation (support for custom software image, scale-out operation). For details, see dbaascli dbhome create --help.



- Support for database creation on subset of cluster nodes. For details, see dbaascli database create --help.
- Ability to run dbaascli in disconnected mode for long running operations. Executing dbaascli command with --waitForCompletion false gets you a job ID that can be queried later to get the status of the operation, using dbaascli job getStatus -jobid job_id. This is useful for long running operations where users may want to get the control back immediately after command execution. In this release, this option is available only for dbaascli database create command. More commands will be added in subsequent releases to have this support. The help output for those commands will reflect the support of --waitForCompletion option.
- Enhanced dbhome patching experience with introduction of multiple options like skipPDBs, continueWithDowntime, and so on. For details, see dbaascli dbhome patch --help.
- Support for better diagnostic collection. For details, see dbaascli diag collect --help.
- Minor improvements in the area of database upgrade automation.
- Various bug fixes and stability improvements.

dbaascli Command Reference

You must use dbaascli to create databases and integrate them with the cloud automation framework.

dbaascli is a cloud native interface that can take DBCA templates as inputs, calls the functionality of DBCA to create databases, and then calls OCI APIs to integrate the database into the cloud automation framework. Customers using DBCA in scripts today can update their existing scripts to call dbaascli instead of DBCA. If dbaascli cannot be used due to a particular feature of DBCA being unavailable in dbaascl, then customers should open a My Oracle Support (MOS) request to add that functionality to dbaascli.

- dbaascli admin updateStack
 To install or update a dbaastools RPM, use the dbaascli admin updateStack 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.

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 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 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.
- 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 getPDBs To view the list of all pluggable databases in a container database, use the dbaascli database getPDBs command.
- dbaascli database modifyParameters
 To modify or reset initialization parameters for an Oracle Database, use the dbaascli database modifyParameters command.
- dbaascli database move To move the database from one home to another, use the dbaascli database move command.



- dbaascli database recover
 To recover a database, use the dbaascli database recover command.
- dbaascli database runDatapatch
 To patch an Oracle Database, use the dbaascli database runDatapatch command.
- 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 dataguard prepareStandbyBlob
 To generate a blob file containing various files that are required on the standby site in case of a dataguard environment, use the dbaascli dataguard prepareStandbyBlob command.
- dbaascli dataguard updateDGConfigAttributes To update Data Guard automation attributes across all the cluster nodes, use the dbaascli dataguard updateDGConfigAttributes command.
- dbaascli dbhome create

To create an Oracle Database home of desired version, use the dbaascli dbhome create command.

- dbaascli dbHome delete To delete a given Oracle Database home, use the dbaascli dbHome delete command.
- dbaascli dbhome getDatabases

To view information about all Oracle Databases running from a given database Oracle home, use the dbaascli dbHome getDatabases command. Specify either the Oracle home location or Oracle home name.

dbaascli dbHome getDetails

To view information about a specific Oracle home, use the dbaascli dbHome getDetails command. Specify either the Oracle home location or Oracle home name.

• dbaascli dbHome patch

To patch Oracle home from one patch level to another, use the dbaascli dbHome patch command.

- dbaascli diag collect To collect diagnostics, use the dbaascli diag collect command.
- dbaascli diag healthCheck To run diagnostic health checks, use the dbaascli diag healthCheck command.
- dbaascli gridHome create To configure Grid Infrastructure home, use the dbaascli gridHome create command.
- dbaascli grid configureTCPS
 To configure TCPS for the existing cluster, use the dbaascli grid configureTCPS command.



dbaascli grid patch

To patch Oracle Grid Infrastructure to the specified minor version, use the dbaascli grid patch command.

- dbaascli grid removeTCPSCert To remove existing TCPS certificates from Grid Infrastructure wallet, use the dbaascli grid removeTCPSCert command.
- dbaascli grid rotateTCPSCert
 To rotate TCPS certificates, use the dbaascli grid rotateTCPSCert command.
- dbaascli grid upgrade

To upgrade Oracle Grid Infrastrucure from one major version to another, use the dbaascli grid upgrade command.

- dbaascli job getStatus
 To view the status of a specified job, use the dbaascli job getStatus command.
- dbaascli patch db apply
- dbaascli patch db prereq
- dbaascli pdb backup

To backup a pluggable database (PDB), query PDB backups, and delete a PDB backup, use the dbaascli pdb backup command.

- dbaascli pdb bounce
 To bounce a pluggable database (PDB), use the dbaascli pdb bounce command.
- dbaascli pdb close
 To close a pluggable database (PDB), use the dbaascli pdb close command.
- dbaascli pdb getConnectString To display Oracle Net connect string information for a pluggable database (PDB) run the dbaascli pdb getConnectString command.
- dbaascli pdb create
 To create a new pluggable database (PDB), use the dbaascli pdb create command.
- dbaascli pdb delete To delete a pluggable database (PDB) run the dbaascli pdb delete command.
- dbaascli pdb getDetails
 To view details of a pluggable database (PDB), use the dbaascli pdb getDetails command.
 - dbaascli pdb list To view the list of pluggable databases (PDB) in a container database, use the dbaascli pdb list command.
- dbaascli pdb localClone
 To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the dbaascli pdb localClone command.
- dbaascli pdb open
 To open a pluggable database (PDB), use the dbaascli pdb open command.
- dbaascli pdb recover
 To recover a pluggable database (PDB), use the dbaascli pdb recover command.
- dbaascli pdb refresh
 To refresh a specified pluggable database (PDB), use the dbaascli pdb refresh
 command.



• dbaascli pdb relocate

To relocate the specified PDB from the remote database into local database, use the dbaascli pdb relocate command.

dbaascli pdb remoteClone
 To create a new pluggable database (PDB) as a clone of an existing PDB in
 another container database (CDB), use the dbaascli pdb remoteClone
 command.

dbaascli system getDBHomes

To view information about all the Oracle homes, use the dbaascli system getDBHomes command.

dbaascli system getGridHomes

To list the details of all Grid homes, use the dbaascli system getGridHomes command.

• dbaascli tde 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.

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 <code>dbaasclitde status</code> 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

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.



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 ${\tt root}$ 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

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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 number of CPUs that you want to scale up or down per VM in a cluster
- --message optionally, you can include a message for your reference

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli cswlib deleteLocal

To delete the local image, use the dbaascli cswlib deleteLocal command.

Run the command as the root user.

Syntax

dbaascli cswLib deleteLocal --imageTag <value>

Where:

--imageTag specifies Oracle home image tag

Example 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 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.

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

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 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 ${\tt root}$ user.

Syntax

dbaascli cswLib listLocal [--product <value>]

Where:



--product identifies Oracle home product type. Valid values: database or grid.

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
1.IMAGE TAG=12.2.0.1.220419
 IMAGE SIZE=5GB
 VERSION=12.2.0.1.220419
 DESCRIPTION=12.2 APR 2022 DB Image
2.IMAGE TAG=18.16.0.0.0
 IMAGE SIZE=6GB
 VERSION=18.16.0.0.0
 DESCRIPTION=18c OCT 2021 DB Image
3.IMAGE TAG=19.14.0.0.0
 IMAGE SIZE=5GB
 VERSION=19.14.0.0.0
 DESCRIPTION=19c JAN 2022 DB Image
dbaascli execution completed
```

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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.

Run the command as the root user.

Syntax

```
dbaascli cswlib showImages
[--product]
```

Where:

• --product identifies Oracle home product type. Valid values: database or grid.

Example 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 Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli database addInstance

To add the database instance on the specified node, use the dbaascli database addInstance command.

Prerequisite

• Run the command as the root user.

Syntax

```
dbaascli database addInstance --dbname <value> --node <value> [--
newNodeSID <value>]
```

Where:

- --dbname specifies Oracle Database name
- --node specifies the node name for the database instance
 - --newNodeSID specifies SID for the instance to add in the new node

dbaascli database backup

To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the dbaascli database backup command.

Prerequisite

• Run the command as the root user.

Syntax

```
dbaascli database backup --dbname <value>
{
        --list
        {
            [--backupType <value>]
            [[--json <value>]
        }
        ! [--json <value>]
        }
        ! --start [--level0] [--level1]
        {
            [[--archival --tag <value>]
            [[--archivelog]
        }
        ! [--archivelog]
        }
        ! --delete --backupTag <value>
        ! --status --uuid <value>
        ! --getBackupReport
        {
        }
    }
}
```



```
--tag <value>
   | --latest
   |
   --json <value>
   | --configure
   {
        --configFile <value>
        | --enableRTRT
        | --disableRTRT
    }
   | --getConfig [--configFile <value>]
   | --validate [--untilTime <value>]
   | --showHistory [--all]
```

Where:

}

```
--dbname: Oracle Database name.
--list | --start | --delete | --status | --getBackupReport | --configure | --
getConfig
--list: Returns database backup information.
    [--json: Specify the file name for JSON output.]
--start: Begins database backup.
        [--level0 | --level1 | --archival]
        [--level0: Creates a Level-0 (full) backup. ]
        [--level1: Creates a Level-1 (incremental) backup. ]
        [--archival: Creates an Archival full backup. ]
            --tag: Specify backup tag.
--delete: Deletes Archival backup.
           --backupTag <value>
--status
           --uuid <value>
--getBackupReport: Returns backup report.
           --tag: Specify backup tag.
           --latest: Returns latest backup report (all types of database backup).
           --json: Specify the file name for JSON output.
--configure: Configures database for backup.
           --configFile | --enableRTRT | --disableRTRT
           --configFile: Specify database backup configuration file.
           --enableRTRT: Enables Real Time Redo Transport.
           --disableRTRT: Disables Real Time Redo Transport.
--getConfig: Returns database backup configuration.
            [--configFile: Specify the database backup configuration file.]
--validate: Validates that backups are complete and corruption-free.
            [--untilTime: Validates from closest Level-0 (full) backup until time
provided. Input format: DD-MON-YYYY HH24:MI:SS.]
--showHistory: Displays the history of backup operations.
            [--all: Displays all backup operations.]
```

Note:

enableRTRT and disableRTRT are applicable only for ZDLRA backup destination on Exadata Database Service on Cloud@Customer.



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
```

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.

Example 7-10 dbaascli database bounce

dbaascli database bounce --dbname dbname

dbaascli database changepassword

To change the password of a specified Oracle Database user, use the dbaascli database changePassword command. When prompted enter the user name for which you want to change the password and then enter the password.

Prerequisites

Run the command as the root or oracle user.

Syntax

```
dbaascli database changePassword [--dbname <value>] [--user <value>]
{
    [--prepareStandbyBlob <value> [--blobLocation <value>]] | [--
standbyBlobFromPrimary <value>]
}
[--resume [--sessionID <value>]]
```

Where:

- --dbname specifies the name of the Oracle Database that you want to act on
- --user specifies the user name whose password change is required
- --prepareStandbyBlob specifies true to generate a blob file containing the artifacts needed to change the password in a Data Guard environment. Valid values: true|false
- --blobLocation specifies the custom path where blob file will be generated
- --standbyBlobFromPrimary specifies the standby blob file, which is prepared from the primary database
- --resume specifies to resume the previous execution
 - --sessionID specifies to resume a specific session ID

Example 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

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:

- --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

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>]
```

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

dbaascli database delete

To delete an Oracle Database, use the dbaascli database delete 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 database delete --dbname
[--deleteArchiveLogs]
[--deleteBackups]
[--precheckOnly]
[--waitForCompletion]
```

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.

Example 7-14 dbaascli database delete

dbaascli database delete --dbname db19

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli database deleteInstance

To delete the database instance on the specified node, use the dbaascli database deleteInstance command.

Prerequisite

• Run the command as the root user.

Syntax

```
dbaascli database deleteInstance --dbname <value> --node <value> [--
continueOnUnreachableNode]
```

Where:

- --dbname specifies Oracle Database name
- --node specifies the node name for database instance
- --continueOnUnreachableNode specifies to perform the operation even if the node is unreachable

Example 7-15 database deleteinstance



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>]
```

- --dbName specifies Oracle Database name
- --sourceDBConnectionString specifies source database connection string in the format of <scan name>:<scan port>/<database service name>
- -- oracleHome specifies Oracle home location
- --oracleHomeName specifies Oracle home name
- --dbSID specifies database SID
- --dbUniqueName specifies database unique name
- --sgaSizeInMB specifies sga target value in mega byte unit
- --pgaSizeInMB specifies pga aggregate target value in mega byte unit
- --datafileDestination specifies ASM disk group name to use for database datafiles
- --fraDestination specifies ASM disk group name to use for database fast recovery area
- --fraSizeInMB specifies fast recovery area size value in mega byte unit
- --sourceDBWalletLocation specifies source database TDE wallet file location. This is required to duplicate database from active database
- --nodeList specifies a comma-delimited list of nodes for the database
- --resume specifies to resume the previous execution
 - -- sessionID specifies to resume a specific session ID
- --revert specifies to rollback the previous execution
 - --sessionID specifies to rollback a specific session ID
- --rmanParallelism specifies parallelsim value
- --rmanSectionSizeInGB specifies RMAN section size in GB
- --tdeConfigMethod specifies TDE configuration method. Allowed values are FILE and KMS.
- --kmsKeyOCID specifies KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE.
- --sourceDBTdeConfigMethod specifies source database TDE configuration method. Allowed values are FILE and KMS.
- --sourceDBKmsKeyOCID specifies source database KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE.
- --executePrereqs specifies yes to run only the prereqs for this operation. Valid values: yes|no
- --waitForCompletion specifies false to run the operation in background. Valid values: true false



 --skipPDBs specifies a comma-delimited list of source database PDB names, which needs to be excluded for the duplicate database operation. Example: pdb1,pdb2...

```
Example 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 getPDBs

To view the list of all pluggable databases in a container database, use the dbaascli database getPDBs command.

Run the command as the root or oracleuser.

Syntax

dbaascli database getPDBs --dbname <value>

Where:

• --dbname specifies the name of the container database

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.

Syntax

```
dbaascli database modifyParameters --dbname --setParameters | --
resetParameters | --responseFile
[--backupPrepared]
[--instance]
[--allowBounce]
```

- --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



- --resetParameters specifies a comma-delimited list of parameters to be reset to their corresponding default values. For example, parameter1=valueA,parameter2=valueB, and so on
- --responseFile specifies the absolute location of the response JSON file to modify the database parameters
- --backupPrepared acknowledges that a proper database backup is in place prior to modifying critical or sensitive parameters.
- --instance specifies the name of the instance on which the parameters will be processed. If not specified, then the operation will be performed at the database level.
- --allowBounce grants permission to bounce the database in order to reflect the changes on applicable static parameters.

```
Example 7-18 dbaascli database modifyParameters
```

```
dbaascli database modifyParameters --dbname dbname --setParameters "log archive dest state 17=ENABLE"
```

dbaascli database move

To move the database from one home to another, use the dbaascli database move command.

Prerequisites

- Before performing a move operation, ensure that all of the database instances associated with the database are up and running.
- Run the command as the root user.

Syntax

```
dbaascli database move
{
    --oracleHome <value> | --oracleHomeName <value>
}
--dbname <value>
[--executePrereqs]
[--resume [--sessionID <value>]]
[--rollback [--sessionID <value>]]
[--skipDatapatch]
[--skipPDBs <value>]
[--skipClosedPDBs]
[--continueWithDbDowntime]
[--allowParallelDBMove]
[--waitForCompletion <value>]
[--nodeList <value>]
```

- -- 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

```
Example 7-19 dbaascli database move
```

```
dbaascli database move --dbname testdb1 --oracleHome /u02/app/oracle/product/
12.1.0/dbhome 2
```

dbaascli database recover

To recover a database, use the dbaascli database recover command.

Prerequisite

- Run the command as the root user.
- Database must have been configured with backup storage destination details where backups are stored.

Syntax



Where:

```
--dbname: Oracle Database name.
    --start | --status
--start: Begins database recovery.
    --untilTime | --untilSCN | --latest | --tag
    --untilTime: Recovers database until time. Input format: DD-MON-YYYY
HH24:MI:SS.
    --untilSCN: Recovers database until SCN.
    --latest: Recovers database to last known state.
    --tag: Recovers database to archival tag.
--status
    --uuid <value>
```

Example 7-20 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]
```

- --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

```
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.

Example 7-21 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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, Connecting to a Virtual Machine with SSH.



Syntax

```
dbaascli database stop
[--dbname]
[--mode]
```

Where:

- --dbname specifies the name of the database that you want to stop
- --mode specifies the mode of the database. Valid values: abort, immediate, normal, transactional

The command performs a database shutdown in immediate mode. No new connections or new transactions are permitted. Active transactions are rolled back, and all connected users are disconnected.

```
Example 7-22 dbaascli database stop
```

dbaascli database stop --dbname db19

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli database upgrade

To upgrade an Oracle Database, use the dbaascli database upgrade 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 database upgrade --dbname --targetHome | --targetHomeName
[--standBy | --allStandbyPrepared]
[--removeGRP]
[--increaseCompatibleParameter]
[--executePrereqs | --postUpgrade | --revert]
[--upgradeOptions]
```

- --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
- --revert 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.

Example 7-23 dbaascli database upgrade pre-upgrade requisite checks

dbaascli database upgrade --dbbname *dbname* --targetHome *Target Oracle home location* --executePrereqs

Related Topics

Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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>

- --dbname specifies the Oracle Database name
- --blobLocation specifies the custom directory location where the standby blob file will be generated in a Data Guard environment



dbaascli dataguard updateDGConfigAttributes

To update Data Guard automation attributes across all the cluster nodes, use the dbaascli dataguard updateDGConfigAttributes command.

Run the command as the root or oracleuser.

Syntax

dbaascli dataguard updateDGConfigAttributes --attributes <value>

Where:

 --attributes contains the Data Guard automation attributes that are to be modified. Accepts comma-delimited values in the format <attribute=value>. Attributes must be predefined in the Data Guard configuration file.

dbaascli dbhome create

To create an Oracle Database home of desired version, use the <code>dbaascli</code> dbhome create command.

Prerequisite

Run the command as the root user.

Syntax

```
dbaascli dbhome create --version
[--oracleHome]
[--oracleHomeName]
[--enableUnifiedAuditing]
[--imageTag]
```

Where:

- --version specifies the version of Oracle Home specified as five numeric segments separated by periods, for example, 19.12.0.00
- --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

Example 7-24 dbaascli dbhome create

dbaascli dbhome create --version 19.11.0.0.0

Alternatively, dbaascli dbhome create --version 19.8.0.0.0.0 --imageTag 19.8.0.0.0 for cases where image tags are different from version.



dbaascli dbHome delete

To delete a given Oracle Database home, use the dbaascli dbHome delete command.

Prerequisite

Run the command as the root user.

Syntax

```
dbaascli dbHome delete
{ --oracleHome <value>
| --oracleHomeName <value> } [--resume [--sessionID <value>]]
```

Where:

- --oracleHome specifies the location of the Oracle home
- --oracleHomeName specifies the name of the Oracle home
- --resume resumes the previous execution
 - -- sessionID specifies to resume a specific session ID

dbaascli dbhome getDatabases

To view information about all Oracle Databases running from a given database Oracle home, use the dbaascli dbHome getDatabases command. Specify either the Oracle home location or Oracle home name.

Run the command as the root user.

Syntax

```
dbaascli dbHome getDatabases
{ --oracleHomeName value | --oracleHome value }
```

Where:

- --oracleHomeName specifies user-defined Oracle home name
- --oracleHome specifies the location (path) of Oracle home

Example 7-25 dbaascli dbHome getDatabases --oracleHome

dbaascli dbHome getDatabases --oracleHome /u02/app/mar home/

dbaascli dbHome getDetails

To view information about a specific Oracle home, use the dbaascli dbHome getDetails command. Specify either the Oracle home location or Oracle home name.

Prerequisite

Run the command as the root user.



Syntax

```
dbaascli dbHome getDetails
{ --oracleHomeName value | --oracleHome value }
```

Where:

- --oracleHomeName specifies user-defined Oracle home name
- --oracleHome specifies the location of Oracle home

Example 7-26 dbaascli dbHome getDetails - using Oracle home location

dbaascli dbHome getDetails --oracleHome /u02/app/home db19c/

Example 7-27 dbaascli dbHome getDetails - using Oracle home name

dbaascli dbHome getDetails --oracleHomeName home db19c

dbaascli dbHome patch

To patch Oracle home from one patch level to another, use the dbaascli dbHome patch command.

Prerequisite

Run the command as the root user.

Syntax

```
dbaascli dbHome patch --oracleHome | --oracleHomeName
--targetVersion
[--resume]
    [--sessionID]
[--continueWithDbDowntime]
[--skipUnreachableNodes]
[--nodes]
[--nodes]
[--executePrereqs]
[--skipDatapatch]
[--skipDatapatch]
[--skipPDBs]
[--skipClosedPDBs]
[--rollback]
```

- --oracleHome specifies the path of Oracle home
- --oracleHomeName specifies the name of Oracle home
- --targetVersion specifies the target version of Oracle Home specified as five numeric segments separated by periods, for example, 19.12.0.0.0.
- --resume resumes the previous run



- --sessionID specifies to resume a specific session ID
- --continueWithDbDowntime continues patching with database downtime. This option can be used in environments wherein there is only one active instance up and the patching operation can be continued even with a downtime.
- --skipUnreachableNodes skips operation on unreachable nodes
- --nodes specifies a comma-delimited list of nodes if patching has to be performed on a subset of nodes
- --executePrereqs runs prereqs
- --skipDatapatch skips running datapatch on the databases
- --imageFilePath specifies the absolute path of the image file to be used
- --skipPDBs skips running the datapatch on a specified comma-delimited list of PDBs. For example: cdb1:pdb1,cdb2:pdb2, and so on
- --skipClosedPdbs skips running datapatch on closed PDBs
- --rollback rolls back patched Oracle home.

Example 7-28 dbaascli dbhome patch

```
dbaascli dbhome patch --targetVersion 19.10.0.0.0 --oracleHome /u02/app/
oracle/product/19.0.0.0/dbhome 2
```

dbaascli 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>]
```

- --components specifies a list of components for log collection. Valid values:
 - db
 - gi
 - os
 - dbaastools



– all

- --startTime specifies the start time for log collection. Valid date and time format: YYYY-MM-DDTHH24:MM:SS
- --endTime specifies the end time for log collection. Valid date and time format: YYYY-MM-DDTHH24:MM:SS
- --nodes specifies a comma-delimited list of nodes to collect logs
- --dbNames specifies the database name for which to collect logs. You can specify only one database name.
- --objectStoreBucketURI specifies an Object Storage service pre-authenticated request (PAR) URL used to upload collected logs. Logs are collected from Guest VM. For more information, see Using Pre-Authenticated Requests.
- --destLocation specifies the location on Guest VM to collect logs.
 Default: /var/opt/oracle/dbaas_acfs
- --waitForCompletion Values: true|false. Default true. Specify false to run in the background.

Related Topics

Using Pre-Authenticated Requests

dbaascli diag healthCheck

To run diagnostic health checks, use the dbaascli diag healthCheck 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 diag healthCheck
[--destLocation]
[--nodes]
[--objectStoreBucketURI]
```

Where:

- --destLocation specifies the directory location on domU to collect logs.
 Default: /var/opt/oracle/dbaas acfs
- --nodes specifies a comma-delimited list of nodes to collect logs
- --objectStoreBucketURI specifies OSS PAR URL to upload log collection

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.



dbaascli gridHome create

To configure Grid Infrastructure home, use the dbaascli gridHome create command.

Prerequisite

Run the command as the root user.

Syntax

```
dbaascli gridHome create --version value [--resume [--sessionID value]] [--
waitForCompletion value]
```

Where:

- --version specifies the Grid home version
- --resume resumes the previous run
 - -- sessionID specifies to resume a specific session ID
- --waitForCompletion specifies false to run the operation in the background. Valid values: true|false

dbaascli grid configureTCPS

To configure TCPS for the existing cluster, use the dbaascli grid configureTCPS command.

Prerequisite

Run the command as the ${\tt root}$ user.

Syntax

Note:

By default, TCPS is enabled for databases on Oracle Exadata Database Service on Dedicated Infrastructure systems.



Note:

TCPS is not enabled for databases on Exadata Database Service on Cloud@Customer systems. To enable TCPS for a given database, update the database specific sqlnet.ora file with WALLET_LOCATION = (SOURCE=(METHOD=FILE)(METHOD_DATA=(DIRECTORY=/var/opt/oracle/ dbaas_acfs/grid/tcps_wallets))) on all database nodes and then bounce the database. This will enable TCPS usage for the database. However, enabling TCPS will cause ZDLRA connection to fail. On Exadata Database Service on Cloud@Customer systems, you can enable either ZDLRA or TCPS configuration. Enabling both ZDLRA and TCPS simultaneously will not work.

```
dbaascli grid configureTCPS
[--pkcs12WalletFilePath]
[--caCertChain]
[--precheckOnly]
[--serverCert]
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

Where:

- --pkcs12WalletFilePath specifies the absolute path of the certificate file, which is in the pkcs12 wallet format
- --caCertChain concatenated list of certs, containing intermediate CA's and root CA certs
- --precheckOnly specifies yes to run only the prechecks for this operation. Valid values: yes or no.
- --serverCert specifies the path of PEM certificate to use or rotate for TCPS configuration.
- --privateKey specifies the path of the private key file of the certificate.
- --certType type of the cert to be added to the Grid Infrastructure wallet. Accepted values are: SELF_SIGNED_CERT, CA_SIGNED_CERT, or PKCS12_CERT. Default: SELF_SIGNED_CERT
- --privateKeyPasswordProtected specifies if the private key is password protected or not. Valid values: true or false. Default: true.

Example 7-29 dbaascli grid configureTCPS

To configure grid using self-signed certificate:

dbaascli grid configureTCPS

To configure grid using CA-signed certificate:

```
dbaascli grid configureTCPS --cert_type CA_SIGNED_CERT --
server_cert /tmp/certs/server_cert.pem --ca_cert_chain /tmp/certs/
```



```
ca.pem --private_key /tmp/certs/encrypted_private.key --
private key password protected false
```

dbaascli grid patch

To patch Oracle Grid Infrastructure to the specified minor version, use the <code>dbaascli grid</code> patch command.

Prerequisites

Run the command as the ${\tt root}$ user.

Syntax

- --targetVersion specifies the target version of Oracle Home specified as five numeric segments separated by periods (e.g. 19.12.0.0.0)
- --containerURL specifies custom URL for fetching Grid Infrastructure image
- --executePrereqs runs prereqs
- --nodeList specifies a comma-delimited list of nodes if patching has to be performed on a subset of nodes
- --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.
- --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
- --waitForCompletion specifies false to run the operation in background. Valid values: true|false



Example 7-30 dbaascli grid patch

dbaascli grid patch --targetVersion 19.12.0.0.0

dbaascli grid removeTCPSCert

To remove existing TCPS certificates from Grid Infrastructure wallet, use the dbaascli grid removeTCPSCert command.

Prerequisite

Run the command as the root user.

Syntax

```
dbaascli grid removeTCPSCert --subject <value>
{
    --userCert | --trustedCert | --requestedCert
}
[--serialNumber <value>] [--executePrereqs] [--resume [--sessionID
<value>]] [--bounceListeners]
```

Where:

- --subject specifies subject of the certificate
- --userCert flag to indicate user certificate
- --trustedCert flag to indicate trusted certificate
- --requestedCert flag to indicate requested certificate
- --serialNumber specifies the serial number of the certificate
- --executePrereqs runs the prerequisite checks and reports the results
- --resume resumes the previous run
 - --sessionID specifies to resume a specific session ID
- --bounceListeners flag to bounce the Grid Infrastructure listener and scan listener

dbaascli grid rotateTCPSCert

To rotate TCPS certificates, use the dbaascli grid rotateTCPSCert command.

Prerequisite

Run the command as the root user.

Syntax

```
dbaascli grid rotateTCPSCert
[--pkcs12WalletFilePath]
[--caCertChain]
[--precheckOnly]
```



```
[--serverCert]
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

Where:

- --pkcs12WalletFilePath specifies the absolute path of the certificate file, which is in the pkcs12 wallet format
- --caCertChain concatenated list of certs, containing intermediate CA's and root CA certs
- --precheckOnly specifies yes to run only the prechecks for this operation. Valid values: yes Or no.
- --serverCert specifies the path of PEM certificate to use or rotate for TCPS configuration.
- --privateKey specifies the path of the private key file of the certificate.
- --certType type of the cert to be added to the Grid Infrastructure wallet. Accepted values are: SELF SIGNED CERT, CA SIGNED CERT, or PKCS12 CERT. Default: SELF SIGNED CERT
- --privateKeyPasswordProtected specifies if the private key is password protected or not. Valid values: true or false. Default: true.

Example 7-31 dbaascli grid rotateTCPSCert

To rotate cert using self-signed certificate (default option):

dbaascli grid rotateTCPSCert

To rotate cert using CA-signed certificate:

```
dbaascli grid rotateTCPSCert --cert_type CA_SIGNED_CERT --server_cert /tmp/
certs/server_cert.pem --ca_cert_chain /tmp/certs/ca.pem --private_key /tmp/
certs/encrypted private.key --privateKeyPasswordProtected true
```

dbaascli grid upgrade

To upgrade Oracle Grid Infrastrucure from one major version to another, use the dbaascli grid upgrade 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 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

Example 7-32 dbaascli grid upgrade

```
daascli grid upgrade --version 19.11.0.0.0 --executePrereqs
DBAAS CLI version MAIN
Executing command grid upgrade --version 19.11.0.0.0 --executePrereqs
```

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an 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 ${\tt root}$ 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

Example 7-33 dbaascli job getStatus

dbaascli job getStatus --jobID 13c82031-f202-41b7-9aef-f4a71df0f551



Related Topics

Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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

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 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.
- To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.
 Soc. Connecting to a Virtual Machine with SSU.

See, Connecting to a Virtual Machine with SSH.

Syntax

```
dbaascli pdb backup --pdbName <value> --dbname <value> {
```



```
--start
{
    [--level1]
    | [--archival --tag <value>]
  }
! --delete --backupTag <value>
! --status --uuid <value>
! --getBackupReport --json <value> --tag <value>
! --list [--json <value>]
```

Where:

}

```
--pdbName: PDB name.
--dbname: Oracle Database name.
--start | --delete | --status | --getBackupReport | --list
--start: Begins PDB backup.
      [--level1 | --archival]
     [--level1: Creates a Level-1 (incremental) backup.]
     [--archival: Creates an archival full backup.]
         --tag: Specify backup tag.
--delete: Deletes archival backup.
         --backupTag: Specify backup tag to delete.
--status
          --uuid <value>
--getBackupReport: Returns backup report.
        --json: Specify the file name for JSON output.
        --taq: Specify backup taq.
--list: Returns PDB backup information.
         [--json: Specify the file name for JSON output.]
```

Example 7-34 Examples

To take level1 backup for a PDB pdb1 in a CDB myTestDb:

dbaascli pdb backup --dbname myTestDb --pdbName pdb1 --start --level1

 To query the status of PDB backup request submitted with uuid eef16b26361411ecb13800163e8e4fac:

dbaascli pdb backup --dbname *myTestDb* --pdbName *pdb1* --status --uuid *eef16b26361411ecb13800163e8e4fac*

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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 $\ensuremath{\mathsf{oracle}}$ 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 pdb bounce --dbname --pdbName | --pdbUID
[-openMode]
```

Where:

- --dbname specifies the name of the container database that hosts the PDB
- --pdbName specifies the name of the PDB
- --pdbUID specifies the identifier of the PDB
- -- openMode specifies the target OPEN MODE of PDB

Example 7-35 dbaascli pdb bounce

dbaascli pdb bounce --dbname
 cdb_name --pdbName pdb name associated with the CDB

dbaascli pdb bounce --dbname cdb name --pdbUID con uid of that pdb

Optional:

- --openMode READ WRITE
- --openMode READ ONLY

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli pdb close

To close a pluggable database (PDB), use the dbaascli pdb close command.

Prerequisite

Run the command as the oracle 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 pdb close --dbname --pdbName | --pdbUID

Where:

• --dbname specifies the name of the container database that hosts the PDB.



- --pdbname specifies the name of the PDB that you want to close.
- --pdbUID specifies the identifier of the PDB

Upon successful completion of running this command, the PDB is closed on all of the container database instances.

Example 7-36 dbaascli pdb close

dbaascli pdb close --dbname *cdb name* --pdbName *pdb name associated with the CDB*

dbaascli pdb close --dbname cdb name --pdbUID con uid of that pdb

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, Connecting to a Virtual Machine with SSH.

Syntax

```
dbaascli pdb getConnectString --dbname --pdbName | --pdbUID
```

Where:

- --dbname specifies the name of the container database that hosts the PDB
- --pdbname specifies the name of the PDB for which you want to display connect string information
- --pdbUID specifies the identifier of the PDB

Example 7-37 dbaascli pdb getConnectString

dbaascli pdb getConnectString --dbname dbname --pdbName pdbName

Related Topics

 Connecting to a Virtual Machine with SSH
 You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.



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 <value>]
[--waitForCompletion <value>]
[--standbyBlobFromPrimary <value>]
```

Where:

- --pdbName specifies the name of the new PDB that you want to create
- --dbName specifies the name of the container database that hosts the new PDB
- --maxCPU optionally specifies the maximum number of CPUs that are available to the PDB. Setting this option is effectively the same as setting the CPU_COUNT parameter in the PDB
- --maxSize optionally specifies the maximum total size of data files and temporary files for tablespaces belonging to the PDB. Setting this option is effectively the same as setting the MAXSIZE PDB storage clause in the CREATE PLUGGABLE DATABASE SQL command. You can impose a limit by specifying an integer followed by a size unit (K, M, G, or T), or you can specify UNLIMITED to explicitly enforce no limit
- --pdbAdminUserName specifies the new PDB admin user name
- --lockPDBAdminAccount specifies true or false to lock the PDB admin user account. Default value is true.
- --resume resumes the previous run
 - -- sessionID specifies to resume a specific session ID
- --executePrereqs specifies yes to run only the prereqs for this operation. Valid values: yes or no
- --waitForCompletion specifies false to run the operation in the background. Valid values: true or false
- --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.

During the PDB creation process, you are prompted to specify the administration password for the new PDB.



Example 7-38 dbaascli pdb create

To create a PDB from seed in a standard database in a non-Data Guard environment:

dbaascli pdb create --dbName db721 --pdbName new_pdb1 --maxsize 5G --maxcpu 2

To create PDB in Data Guard environment:

dbaascli pdb create --dbName db721 --pdbName new pdb1

```
dbaascli pdb create --dbName db721 --pdbName new_pdb1 --
standbyBlobFromPrimary /tmp/send db721.tar
```

dbaascli pdb 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 value]
[--waitForCompletion value]
[--resume [--sessionID value]]
[--allStandbyPrepared]
[--cleanupRelocatedPDB]
```

Where:

- --dbName specifies the name of the container database that hosts the PDB
- --pdbName specifies the name of the PDB that you want to delete
- --pdbUID specifies the UID of the PDB that you want to delete
- --executePrereqs specifies yes to run only the prereqs for this operation. Valid values: yes or no
- --waitForCompletion specifies false to run the operation in the background. Valid values: true or false
- --resume specifies to resume the previous execution
 - --sessionID specifies to resume a specific session ID
- --allStandbyPrepared specifies to confirm that the operation has been successfully run on all the standby databases
- --cleanupRelocatedPDB option to cleanup source database after a PDB has been relocated.

Example: dbaascli pdb delete



To delete a PDB from a standard database in a non-Data Guard environment or from Standby database in Data Guard environment.

dbaascli pdb delete --dbName db721 --pdbName pdb1

To create PDB from Primary database in Data Guard environment:

dbaascli pdb create --dbName db721 --pdbName pdb1 --allStandbyPrepared

dbaascli pdb getDetails

To view details of a pluggable database (PDB), use the dbaascli pdb getDetails command.

Prerequisite

Run the command as the oracle user.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, Connecting to a Virtual Machine with SSH.

Syntax

dbaascli pdb getDetails --dbname --pdbName | --pdbUID

Where:

- --dbname specifies the name of the container database that hosts the PDB
- --pdbname specifies the name of the PDB that you want to delete
- --pdbUID specifies the identifier of the PDB

Example 7-39 dbaascli pdb getDetails

dbaascli pdb getDetails--dbname *cdb name* --pdbName *pdb name associated* with the CDB

dbaascli pdb getDetails--dbname cdb name --pdbUID con_uid of that pdb

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.



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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, Connecting to a Virtual Machine with SSH.

Syntax

dbaascli pdb list --dbname

Where:

--dbname specifies the name of the container database that hosts the PDB

```
Example 7-40 dbaascli pdb list
```

dbaascli pdb list --dbname cdb name

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, Connecting to a Virtual Machine with SSH.

Syntax

```
dbaascli pdb localClone --dbName --pdbName
[--resume]
[--maxCPU]
[--waitForCompletion]
[--primaryDBWalletTar]
[--maxSize]
[--targetPDBName]
[--executePrereqs]
[--powerLimit]
```



Where:

- --dbName specifies the name of the database
- --pdbName specifies the name of the new PDB that you want to clone
- --resume resumes the previous run
- --maxCPU specifies the maximum number of CPUs to be allocated for the PDB
- --waitForCompletion specifies false to run the operation in the background. Valid values: true or false
- --primaryDBWalletTar specifies the primary database wallet tar file. This is required only for standby database PDB operations
- --maxSize specifies the maximum storage size in GB for the new PDB
- --targetPDBName specifies the name for the target PDB (new cloned PDB)
- --executePrereqs specifies yes to run only the prereqs for this operation. Valid values: yes or no
- --powerLimit specifies the degree of parallelism to be used for the clone operation. Valid value is between 1 and 128

The newly cloned PDB inherits administration passwords from the source PDB.

Example 7-41 dbaascli pdb localClone

```
dbaascli pdb localClone --dbName db35 --pdbName PDB35 --targetPDBName local clone1 --maxCPU 2 --maxSize 15
```

Related Topics

Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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>]]
```

- --pdbName specifies the name of the PDB that you want to open
- --pdbUID specifies the identifier of the PDB



- --dbname specifies the name of the container database that hosts the PDB.
- -- openMode specifies the target OPEN MODE of PDB
- --startServices: specifies to start all or list all services corresponding to a PDB. Accepted values are all or a comma-delimited list of PDB services.
- --waitForCompletion: specify false to run the operation in the background. Valid values: true|false
- --setPDBRefreshModeNone: specifies to convert a refreshable PDB to non-refreshable PDB
 - -- skipPDBRefresh: specifies to skip refreshable PDB refresh
 - --pdbAdminUserName: specifies new PDB admin user name

Upon successful completion, the PDB is opened on all of the container database instances.

Example 7-42 dbaascli pdb open

dbaascli pdb open --dbname *cdb name* --pdbName *pdb name associated with the CDB*

dbaascli pdb open --dbname cdb name --pdbUID con uid of that pdb

Optional: -- openMode READ WRITE/READ ONLY

dbaascli pdb recover

To recover a pluggable database (PDB), use the dbaascli pdb recover command.

Prerequisite

- Run the command as the root user.
- Database must be configured with backup storage destination details where backups are stored.

Syntax

```
dbaascli pdb recover --pdbName <value> --dbname <value>
{
    --start
    {
        --untilTime <value>
        | --untilSCN <value>
        | --latest
        | --tag <value>
    }
    | --status --uuid <value>
}
```

Where:

--pdbName: PDB name. --dbname: Oracle Database name.



```
--start | --status
--start
        --untilTime | --untilSCN | --latest | --tag
        --untilTime: Recovers PDB until time. Input format: DD-MON-YYYY
HH24:MI:SS.
        --untilSCN: Recovers PDB until SCN.
        --latest: Recovers PDB to last known state.
        --tag: Recovers PDB to last known state.
        --tag: Recovers PDB to archival tag.
--status
        --uuid <value>
```

Example 7-43 Examples

• To recover a PDB *pdb1* in a CDB *myTestDb* to latest:

```
dbaascli pdb recover --dbname myTestDb --pdbName pdb1 --start -- latest
```

• To query the status of PDB recovery request submitted with uuid 81a17352362011ecbc3000163e8e4fac:

```
dbaascli pdb recover --dbname myTestDb --pdbName pdb1 --status --
uuid 81a17352362011ecbc3000163e8e4fac
```

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata 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>]
```

- --dbname: specifies the name of the Oracle Database
- --pdbName: specifies the name of the pluggable database
- --pdbUID: specifies the identifier of the pluggable database
- --waitForCompletion: specify false to run the operation in the background. Valid values: true|false



Related Topics

Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli pdb relocate

To relocate the specified PDB from the remote database into local database, use the dbaascli pdb relocate command.

Prerequisite

Run the command as the oracle user. When prompted, you must supply the SYS user password for the source database.

Syntax

```
dbaascli pdb relocate --pdbName <value> --dbName <value> --
sourceDBConnectionString <value>
[--targetPDBName <value>]
[--powerLimit <value>]
[--maxCpu <value>]
[--maxSize <value>]
[--resume [--sessionID <value>]]
[--executePreregs <value>]
[--sourcePDBServices <value>]
[--sourcePDBReadOnlyServices <value>]
[--waitForCompletion <value>]
{
    [--blobLocation <value>] | [--standbyBlobFromPrimary <value>]
}
[--upgradePDB <value>]
[--updateDBBlockCacheSize]
{
    [skipOpenPDB] | [--completePDBRelocate]
}
```

- --pdbName specifies the source PDB name to relocate
- --dbName specifies the target database name
- --sourceDBConnectionString specifies the source database connection string in the format <scan_name>:<scan_port>/<database_service_name>
- --targetPDBName specifies a name for the target PDB (new relocated PDB)
- --powerLimit specifies the degree of parallelism to be used for the relocate operation
- --maxCpu specifies the maximum number of CPUs to be allocated for the PDB
- --maxSize specifies the maximum storage size in GB for the new PDB
- --resume specifies to resume the previous execution
 - --sessionID specifies to resume a specific session ID



- --executePrereqs specifies yes to run only the prereqs for this operation. Valid values: yes|no
- --sourcePDBServices specifies a list of comma-delimited source PDB services
- --sourcePDBReadOnlyServices specifies a comma-delimited list of source PDB read-only services
- --waitForCompletion specifies false to run the operation in the background. Valid values: true|false
- --blobLocation specifies the location of a custom directory where the standby BLOB file will be generated in a Data Guard environment.
- --standbyBlobFromPrimary specifies the location of the standby BLOB file, which is prepared from the primary database. This is required only for standby operations.

Note:

The parameters --blobLocation and mutually exclusive.

- --upgradePDB specifies true to upgrade the PDB as part of this operation. Valid values : true | false.
- --updateDBBlockCachesize option to enable application to set db block cache size initialization parameters to support data copy with different block size.
- --skipOpenPDB indicates that the PDB should not be opened at the end of the current operation.
- --completePDBRelocate completes the PDB relocation if done as a two-step operation.

Example 7-44 dbaascli pdb relocate

```
dbaascli pdb relocate --sourceDBConnectionString test-
scan.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521/
source cdb service name --pdbName source pdb --dbName target db
```

dbaascli pdb remoteClone

To create a new pluggable database (PDB) as a clone of an existing PDB in another container database (CDB), use the dbaascli pdb remoteClone command.

Run the command as the root or oracle user.

Syntax

```
dbaascli pdb remoteClone --pdbName <value> --dbName <value> --
sourceDBConnectionString <value> [--targetPDBName <value>] [--
powerLimit <value>] [--maxCPU <value>] [--maxSize <value>] [--resume
[--sessionID <value>]] [--executePrereqs] [--waitForCompletion
<value>] [--sourcePDBExportedTDEKeyFile <value>]
{
        [--blobLocation <value>]
```



Where:

- --pdbName specifies the name of the source PDB that you want to clone
- --dbname specifies the name (DB NAME) of the CDB that hosts the newly cloned PDB
- --sourceDBConnectionString specifies the source database connection string in the format scan name:scan port/database service name
- --targetPDBName specifies the name for the target PDB (new cloned PDB)
- --powerLimit specifies the degree of parallelism to be used for the clone operation. Valid value is between 1 and 128
- --maxCPU specifies the maximum number of CPUs to be allocated for the PDB
- --maxSize specifies the maximum storage size in GB for the new PDB
- --resume resumes the previous run
 - --sessionID specifies to resume a specific session ID
- --executePrereqs specifies yes to run only the prereqs for this operation. Valid values: yes or no
- --waitForCompletion specifies false to run the operation in the background. Valid values: true or false
- --sourcePDBExportedTDEKeyFile specifies the source PDB exported key file. This variable is applicable to only 12.1 database.
- --blobLocation specifies the custom path where the standby blob file will be generated in a Data Guard environment
- --standbyBlobFromPrimary specify the location of the standby blob file, which is prepared from the primary database. This is required only for standby database PDB operations

Note:

The parameters --blobLocation and --standbyBlobFromPrimary are mutually exclusive.

- --excludeUserTablespaces option to skip user table spaces, example *t1,t2,t3*.
- --excludePDBData specify true/yes to skip user data from source PDB.
- --pdbAdminUserName specifies new PDB admin user name



- --lockPDBAdminAccount specify true or false to lock the PDB admin user account. Default value is true.
- --sourcePDBServiceConvertList specify a comma-delimited list of source to target service names, which need to be converted. Syntax is source_srv1:new_srv1, source_srv2:new_srv2.
- --refreshablePDB specifies to create refreshable PDB
 - --refreshMode specifies refresh mode for refreshable PDB. Valid values: AUTO[MANUAL
 - * --refreshIntervalInMinutes specifies refresh interval for refreshablePDB in minutes
 - --dblinkUsername specifies common user of a remote database used for database link to connect to the remote database
 - * --honorCaseSensitiveUserName indicates specified username is case sensitive
- --updateDBBlockCacheSize: specifies to enable application to set db block cache size initialization parameters to support data copy with a different block size

When promoted, you must supply the SYS user password for the source PDB. The newly cloned PDB inherits administration passwords from the source PDB. The cloned PDB is named using the following format: dbname_sourcepdbname. This command is supported only for databases that are not in a Data Guard configuration and use Oracle Database version 12.2.0.1, or later.

Example 7-45 dbaascli pdb remoteClone

```
dbaascli pdb remoteClone --sourceDBConnectionString test-
can.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521 --pdbName
source_pdb1 --dbName db9944 --targetPDBName new_pdb1 --maxsize 5 --
maxcpu 2
```

```
dbaascli pdb remoteClone --sourceDBConnectionString

orcla.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com --pdbName

source_pdb1 --dbName db9944 --targetPDBName new_pdb1 --maxsize 5 --

maxcpu 2
```

dbaascli system getDBHomes

To view information about all the Oracle homes, use the dbaascli system getDBHomes command.

Prerequisite

Run the command as the root or oracle user.

Syntax

dbaascli system getDBHomes



Example 7-46 dbaascli system getDBHomes

dbaascli system getDBHomes

dbaascli system getGridHomes

To list the details of all Grid homes, use the dbaascli system getGridHomes command.

Prerequisite

Run the command as the root or oracle user.

Syntax

dbaascli system getGridHomes

dbaascli tde 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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, Connecting to a Virtual Machine with SSH.

Syntax

```
dbaascli tde addSecondaryHsmKey --secondaryKmsKeyOCID
[--dbname]
[--precheckOnly]
```

Where:

- --secondaryKmsKeyOCID specifies the secondary KMS key to add to the existing HSM (KMS) configuration
- --dbname specifies the name of the database
- --precheckOnly specifies yes to run only the prechecks for this operation. Valid values: yes or no

Example 7-47 dbaascli tde addSecondaryHsmKey

```
dbaascli tde addSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocidl.key.ocl.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5u
xainvvxza
```

```
dbaascli tde addSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
```



frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krx
wllk5uxainvvxza --precheckOnly yes

Related Topics

Connecting to a Virtual Machine with SSH

You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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>] [--pdbName <value>]

Where:

- --dbname specifies the name of the database
- --pdbName specifies the name of the PDB for which the TDE Keystore needs to be changed

dbaascli tde enableWalletRoot

To enable wallet_root spfile parameter for the existing database, use the dbaasclitde enableWalletRoot 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 tde enableWalletRoot
[--dbRestart]
[--dbname]
[--precheckOnly]
```

Where:

 --dbrestart specifies the database restart option. Valid values are: rolling or full. Default value: rolling
 If you do not pass the dbreatart argument, then the database restarts in a

If you do not pass the ${\tt dbrestart}$ argument, then the database restarts in a rolling manner.



- --dbname specifies the name of the Oracle Database.
- --precheckOnly runs only the precheck for this operation. Valid values are: yes or no

```
Example 7-48 dbaascli tde enableWalletRoot
```

```
dbaascli tde enableWalletRoot --dbname db name --dbrestart rolling|full
dbaascli tde enableWalletRoot --dbname orcl
dbaascli tde enableWalletRoot --dbname orcl--dbrestart full
```

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, Connecting to a Virtual Machine with SSH.

Syntax

```
dbaascli tde encryptTablespacesInPDB --pdbName
[--dbname]
[--precheckOnly]
[--useSysdbaCredential]
```

Where:

- --pdbName specifies the name of the PDB to encrypt all the tablespaces.
- --dbname specifies the name of the Oracle Database.
- --precheckOnly runs only the precheck for this operation. Valid values: yes or no
- --useSysdbaCredential uses SYSDBA credentials for this operation if passed value is true. Valid values: true or false



Example 7-49 dbaascli tde encryptTablespacesInPDB

dbaascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb

```
dbaascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb -- precheckOnly yes
```

dbaascli tde encryptTablespacesInPDB --dbname *dbname* --pdbName *pdb* -- useSysdbaCredential *true*

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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 <value>]
[--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 specifies yes to run only the prechecks for this operation. Valid values: yes or no
- --primarySuc specify this property in the standby database of the Data Guard environment once the command is successfully run on the primary database
- --resume specifies to resume the previous run



- -- sessionID specifies to resume a specific session ID
- --revert specifies to rollback the previous run
 - -- sessionID specifies to rollback a specific session ID
- --waitForCompletion specify false to run the operation in background. Valid values : true|false.

Example 7-50 dbaascli tde fileToHsm --kmsKeyOCID

```
dbaascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza
```

```
dbaascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza --precheckOnly yes
```

```
dbaascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux
ainvvxza --resume
```

dbaascli tde getHsmKeys

To get TDE active key details, use the dbaascli tde getHsmKeys 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 tde getHsmKeys
[--dbname]
[--infoFile]
```

Where:

- --dbname specifies the name of the database
- --infoFile specifies the file path where the list of OCIDs will be saved. The output is in JSON format

Example 7-51 dbaascli tde getHsmKeys

```
dbaascli tde getHsmkeys --dbname dbname
```

dbaascli tde getHsmkeys --dbname dbname --infoFile infoFilePath



Related Topics

 Connecting to a Virtual Machine with SSH
 You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

dbaascli tde getMkidForKeyVersionOCID

To get Master Key ID associated with the KMS key version OCID, use the dbaasclitde getMkidForKeyVersionOCID 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 tde getMkidForKeyVersionOCID --kmsKeyVersionOCID
[--dbname]
```

Where:

- --kmsKeyVersionOCID specifies the KMS key version OCID to set
- --dbname specifies the name of the database

Example 7-52 dbaascli tde getMkidForKeyVersionOCID

```
dbaascli tde getMkidForKeyVersionOCID --dbname dbname --
kmsKeyVersionOCID ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcw
vylkd21x56lu2s6iwnxwgigu23nha
```

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, Connecting to a Virtual Machine with SSH.



Syntax

```
dbaascli tde getPrimaryHsmKey
[--dbname]
```

Where:

• --dbname specifies the name of the database

Example 7-53 dbaascli tde getPrimaryHsmKey

dbaascli tde getPrimaryHsmKey --dbname dbname

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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>]
[--primaryDBWalletTar <value>]
[--skipPatchCheck <value>]
[--executePrereqs <value>]
[--primarySuc <value>]
[--resume [--sessionID <value>]] [--revert]
[--waitForCompletion <value>]
```

Where:

- --dbname specifies the name of the database
- --primaryDBWalletTar specifies the primary database wallet tar file. This is required only for standby database hsmToFile conversion
- --skipPatchCheck skips validation check for required patches if the value passed for this argument is true. Valid values: true or false
- --executePrereqs specifies yes to run only the prechecks for this operation. Valid values: yes or no
- --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
- --waitForCompletion specifies false to run the operation in background. Valid values: true|false

```
Example 7-54 dbaascli tde hsmToFile
```

```
dbaascli tde hsmToFile --dbname dbname
dbaascli tde hsmToFile --dbname dbname --executePrereqs yes
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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine. See, *Connecting to a Virtual Machine with SSH*.

Syntax

```
dbaascli tde listKeys
[--file]
[--dbname]
```

Where:

- --file specifies the file path to save the results
- --dbname specifies the name of the database

Example 7-55 dbaascli tde listKeys

dbaascli tde listKeys --dbname dbname

dbaascli tde listKeys --dbname dbname --file infoFilePath

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.



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.

To use the utility, you must connect to an Exadata Cloud@Customer virtual machine.

See, Connecting to a Virtual Machine with SSH.

Syntax

```
dbaascli tde removeSecondaryHsmKey
[--force]
[--secondaryKmsKeyOCID]
[--dbname]
[--precheckOnly]
```

Where:

- --force if not specified, you will be prompted while deleting all of the existing HSM (KMS) keys
- --secondaryKmsKeyOCID specifies the secondary KMS key to add to the existing HSM (KMS) configuration
- --dbname specifies the name of the database
- --precheckOnly specifies yes to run only the prechecks for this operation. Valid values: yes or no

Example 7-56 dbaascli tde removeSecondaryHsmKey

dbaascli tde removeSecondaryHsmKey --dbname dbname

```
dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5u
xainvvxza
```

```
dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocidl.key.ocl.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5u
xainvvxza --precheckOnly yes
```

Related Topics

 Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.



dbaascli tde rotateMasterKey

To rotate the master key for database encryption, use the dbaascli tde rotateMasterKey command.

Prerequisites:

Run the command as the root user.

Syntax

```
dbaascli tde rotateMasterKey --dbname <value>
[--rotateMasterKeyOnAllPDBs]
[--pdbName <value>]
[--executePrereqs]
[--resume [--sessionID <value>]]
{
    [--prepareStandbyBlob <value> [--blobLocation <value>]]
    [ [--standbyBlobFromPrimary <value>]
}
```

Where:

- --dbname specifies the name of the Oracle Database
- --rotateMasterKeyOnAllPDBs specifies true to rotate master key of all PDBs in CDB. Valid values: true|false
- --pdbName specifies the name of the PDB
- --executePrereqs runs the prerequisites checks and report the results
- --resume specifes to resume the previous execution
- --sessionID specifies to resume a specific session ID
- --prepareStandbyBlob specifies true to generate a BLOB file containing the artifacts needed to perform the operation in a Data Guard environment
- --blobLocation specifies the location of the custom directory where the standby BLOB file will be generated in a Data Guard environment
- --standbyBlobFromPrimary specifies the location of the standby BLOB file, which is prepared from the primary database. This is required only for standby operations.

dbaascli tde setKeyVersion

To set the version of the primary key to be used in DB/CDB or PDB, use the dbaasclitde setKeyVersion command.

Run the command as the root user.

Syntax

```
dbaascli tde setKeyVersion --kmsKeyVersionOCID <value>
[--dbname <value>]
```



```
[--pdbName <value>]
[--masterKeyID <value>]
[--standbySuc <value>]
[--executePrereqs <value>]
[--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 specifies yes to run only the prechecks for this operation. Valid values: yes or no
- --waitForCompletion specify false to run the operation in background. Valid values: true|false

Example 7-57 dbaascli tde setKeyVersion

```
dbaascli tde setKeyVersion --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2
lx56lu2s6iwnxwgigu23nha
```

```
dbaascli tde setKeyVersion --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2
lx56lu2s6iwnxwgigu23nha --precheckOnly yes
```

```
dbaascli tde setKeyVersion --dbname dbname --pdbName pdb --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2
lx56lu2s6iwnxwgigu23nha
```

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>]
[--standbySuc <value>]
```



```
[--precheckOnly]
[--bounceDatabase]
```

Where:

- --primaryKmsKeyOCID specifies the primary KMS key to set
- --dbname specifies the name of the database
- --standbySuc specify this property in the primary database of the Data Guard environment once the command is successfully run on the standby database
- --precheckOnly specifies yes to run only the prechecks for this operation. Valid values: yes or no
- --bounceDatabase specify this flag to do rolling database bounce for this operation

Example 7-58 dbaascli tde setPrimaryHsmKey

```
dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocidl.key.ocl.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krx
wllk5uxainvvxza
```

```
dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxtygoiqpm2pu2afgta54krx
wllk5uxainvvxza --precheckOnly yes
```

dbaascli tde status

To display information about the keystore for the specified database, use the <code>dbaasclitde status</code> command.

Prerequisite

Run the command as the oracle 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 tde status --dbname dbname
```

Where:

• --dbname specifies the name of the database that you want to check.

Output from the command includes the type of keystore, and the status of the keystore.

Example 7-59 dbaascli tde status

dbaascli tde status --dbname dbname



Related Topics

 Connecting to a Virtual Machine with SSH
 You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

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 Exadata Database Service on
 Cloud@Customer services 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.

You can determine the name of the VM cluster by running the following $\tt crsctl$ command as the $\tt grid$ 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 Exadata Database Service on Cloud@Customer services and objects.

Purpose

ExaCLI (exacli) enables you configure your 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]
[-1 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 filename of the cookie jar to use. If you do not specify a filename, then the cookie is stored in a default cookie jar located at <i>HOME</i> /.exacli/cookiejar, where <i>HOME</i> 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.



Option	Description
<pre>-e command or -e 'command[; command]' or -e @batchFile</pre>	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 orno-prompt	Suppresses prompting for user input.

Command Parameters

To obtain information about objects and services on your system, use these ExaCLI command parameters.

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.



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 <code>object_id</code> value of an object				
	<pre>select object_name, data_object_id from user_objects where object_name = 'BIG_CENSUS'; OBJECT_NAME DATA_OBJECT_ID</pre>				
	 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.				

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.				

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.

Table 7-1 (Cont.) Command

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.



- 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 7-60 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 -1 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 7-61 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.

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=edlcl03clu01-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.
- 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	
	••
	Value
<pre>> Send CEF metrics to OCI Monitoring (defaultOciMonitoring) '</pre>	ON



If the defaultocimonitoring flag is set to OFF, then run the tfact1 set defaultocimonitoring=on or tfact1 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
 1_____
+----'
```

• 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.

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 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*.

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 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** dropdown 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*.

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

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.

N	oto	
IN	ote	

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.

• To check the run status of Oracle Trace File Analyzer, run the tfact1 status command as root or a non-root user:

<pre># tfactl status</pre>						
	·					
Host	Status of TFA	PID		Port	Version	I
Build ID	Inventory	Status				
+			-+-		+	-
+	+	+				
node1	RUNNING	41312		5000	22.1.0.0.0	
22100020220310214	615 COMPLETE					
node2	RUNNING	272300		5000	22.1.0.0.0	
22100020220310214	615 COMPLETE					
!	+	-+	-+-		+	-
+	+	'				



• To start the Oracle Trace File Analyzer daemon on the local node, run the tfact1 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 tfact1 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 Exadata Database Service on Cloud@Customer in the Monitoring Service

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

Dimensions

All the metrics discussed in this topic include the following dimensions.

- RESOURCEID: The OCID of the VM Cluster.
- **RESOURCENAME**: The name of the VM Cluster.

NOT_SUPPORTED

The metrics listed in the following table are automatically available for the VM cluster.



Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collectio n Frequenc y	Dimensions	
ASMDiskgrou pUtilizatio n	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	hostName deploymentTy pe diskgroupNam e	
CpuUtilizat ion	CPU Utilization	percentage	Percent CPU utilization	1 minute	hostName deploymentTy pe	
FilesystemU tilization	Filesystem Utilization	percentage	Percent utilization of provisioned filesystem	1 minute	hostName deploymentTy pe filesystemNa me	
LoadAverage	Load Average	integer	System load average over 5 minutes	1 minute	hostName deploymentTy pe	
MemoryUtili zation	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 deploymentTy pe	
NodeStatus	Node Status	integer	Indicates whether the host is reachable.	1 minute	hostName deploymentTy pe	
OcpusAlloca ted	OCPU Allocated	integer	The number of OCPUs allocated	1 minute	deploymentTy pe	
SwapUtiliza tion	Swap Utilization	percentage	Percent utilization of total swap space	1 minute	hostName deploymentTy pe	

NOT_SUPPORTED

The metrics listed in the following table are automatically available for the database.



Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions
CpuUtilizati on	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 r instanceName hostName deploymentTyp e resourceId_{da abase pdb} resourceName _{database pdb}
StorageUtili zation	Storage Utilization	percentage	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.	1 hour	deploymentTyp e resourceId_{dat abase pdb} resourceName _{database pdb}
BlockChanges	DB Block Changes	Changes per second	The Average number of blocks changed per second.	5 minutes	instanceNumbe r instanceName hostName deploymentTyp e resourceId_{dat abase pdb} resourceName _{database pdb}
ExecuteCount	Execute Count	Count	The number of user and recursive calls that executed SQL statements during the selected interval.	5 minutes	instanceNumbe r instanceName hostName deploymentTyp e



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	instanceNumbe r instanceName hostName deploymentTyp e
CurrentLogon s	Current Logons	Count	The number of successful logons during the selected interval.	5 minutes	instanceNumbe r instanceName hostName deploymentTyp e resourceId_{dat abase pdb} resourceName _{database pdb}
TransactionC ount	Transaction Count	Count	The combined number of user commits and user rollbacks during the selected interval.	5 minutes	instanceNumbe r instanceName hostName deploymentTyp e resourceId_{dat abase pdb} resourceName _{database pdb}
UserCalls	User Calls	Count	The combined number of logons, parses, and execute calls during the selected interval.	5 minutes	instanceNumbe r instanceName hostName deploymentTyp e resourceId_{dat abase pdb} resourceName _{database pdb}

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions
ParseCount	Parse Count	Count	The number of hard and soft parses during the selected interval.	5 minutes	instanceNumbe r instanceName hostName deploymentTyp e resourceId_{dat abase pdb} resourceName _{database pdb}
StorageUsed	Storage Space Used	GB	Total amount of storage space used by the database at the collection time.	1 hour	deploymentTyp e resourceId_{dat abase pdb} resourceName _{database pdb}
StorageAlloc ated	Storage Space Allocated	GB	Total amount of storage space allocated to the database at the collection time	1 hour	deploymentTyp e resourceId_{dat abase pdb} resourceName _{database pdb}
StorageUsedB yTablespace	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	tablespaceNam e, tablespaceType deploymentTyp e resourceId_{dat abase pdb} resourceName _{database pdb}
StorageAlloc atedByTables pace	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	TablespaceNa me, tablespaceType , deploymentTyp e, resourceId_{dat abase pdb} resourceName _{database pdb}

Metric Name	Metric Display Name	Unit	Dsicription and Metric Chart Defaults	Collection Frequency	Dimensions
StorageUtili zationByTabl espace	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	tablespaceNam e, tablespaceType deploymentTyp e

Oracle Exadata Database Service on Cloud@Customer Events

Exadata Cloud@Customer resources emit events, which are structured messages that indicate changes in resources.

- About Event Types on Exadata Cloud@Customer Learn about the event types available for Exadata Cloud@Customer resources.
- Exadata Infrastructure Event Types Review the list of event types that Exadata Infrastructure instances 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 (Cloud@Customer)
 Review the list of event types that database nodes emit.
- Database Home Event Types (Cloud@Customer) Review the list of event types that Database Homes emit.
- Database Event Types (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.
- Autonomous VM Cluster Event Types
 Review the list of event types that Autonomous VM clusters emit.
- Autonomous Container Database Event Types Review the list of event types that Autonomous Container Database emit.
- Autonomous Database Event Types Review the list of event types that Autonomous Database emit.



- Data Guard Event Types Review the list of event types that Data Guard associations emit.
- Autonomous Data Guard Association Event Types Review the list of event types that Autonomous Data Guard associations emit.
- Key Store Event Types Review the list of event types that Key Store emits.
- Exadata Cloud@Customer Infrastructure Maintenance Event Types Review the list of event types that Exadata Cloud@Customer 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.
- 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 Exadata Cloud@Customer

Learn about the event types available for Exadata Cloud@Customer resources.

Exadata 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.



Friendly Name	Event Type
Activate Begin	com.oraclecloud.databaseservice.acti vateexadatainfrastructure.begin
Activate End	com.oraclecloud.databaseservice.acti vateexadatainfrastructure.end
Change Compartment	com.oraclecloud.databaseservice.chan geexadatainfrastructurecompartment
Configuration File Download	com.oraclecloud.databaseservice.down loadexadatainfrastructureconfigfile
Create Begin	com.oraclecloud.databaseservice.crea teexadatainfrastructure.begin
Create End	com.oraclecloud.databaseservice.crea teexadatainfrastructure.end
Delete Begin	com.oraclecloud.databaseservice.dele teexadatainfrastructure.begin
Delete End	com.oraclecloud.databaseservice.dele teexadatainfrastructure.end
Update Begin	com.oraclecloud.databaseservice.upda teexadatainfrastructure.begin
Update End	com.oraclecloud.databaseservice.upda teexadatainfrastructure.end
Exadata Infrastructure - Connectivity Status	com.oraclecloud.databaseservice.exad atainfrastructureconnectstatus

Table 7-2 Exadata Infrastructure Event Types	Table 7-2	Exadata	Infrastructure	Event	Types
--	-----------	---------	----------------	-------	-------

Example 7-62 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",
     "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..aaaaaaaayryg14guhplo5rtiumx3eh4mk7grrkrqspzaltmvbxecn
bvhkrga",
    "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"
      }
    }
  },
  "eventID" : "cde92d45-a06b-4b69-a125-6005dd8c2f0c",
  "extensions" : {
    "compartmentId" :
```



```
"ocid1.compartment.oc1..aaaaaaaayrygl4guhplo5rtiumx3eh4mk7grrkrqspzaltm
vbxecnbvhkrga"
}
```

VM Cluster Network Event Types

Review the list of event types that VM cluster networks emit.

Friendly Name	Event Type
Create Begin	com.oraclecloud.databaseservice.crea tevmclusternetwork.begin
Create End	com.oraclecloud.databaseservice.crea tevmclusternetwork.end
Network Validation File Download	com.oraclecloud.databaseservice.down loadvmclusternetworkconfigfile
Terminate Begin	com.oraclecloud.databaseservice.dele tevmclusternetwork.begin
Terminate End	com.oraclecloud.databaseservice.dele tevmclusternetwork.end
Update Begin	com.oraclecloud.databaseservice.crea tevmclusternetwork.begin
Update End	com.oraclecloud.databaseservice.crea tevmclusternetwork.end
Validate Begin	com.oraclecloud.databaseservice.vali datevmclusternetwork.begin
Validate End	com.oraclecloud.databaseservice.vali datevmclusternetwork.end

Table 7-3 VM Cluster Network Event Types

Example 7-63 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",
            "compartmentId": "example name",
            "compartmentName": "example name",
            "exempted
            "example name",
            "exempted
            "exempted
            "compartmentName": "example name",
            "compartmentName": "example name",
            "exempted
            "compartmentName": "exampted
            "compar
```



```
"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"
    }
 }
}
```

VM Cluster Event Types

Review the list of event types that VM clusters emit.

Friendly Name	Event Type
Change Compartment	com.oraclecloud.databaseservice.changev mclustercompartment
Create Begin	<pre>com.oraclecloud.databaseservice.createv mcluster.begin</pre>
Create End	com.oraclecloud.databaseservice.createv mcluster.end
Terminate Begin	com.oraclecloud.databaseservice.deletev mcluster.begin
Terminate End	<pre>com.oraclecloud.databaseservice.deletev mcluster.end</pre>
Update Begin	com.oraclecloud.databaseservice.updatev mcluster.begin
Update End	<pre>com.oraclecloud.databaseservice.updatev mcluster.end</pre>

Table 7-4 VM Cluster Event Types

Example 7-64 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",
      "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.

 Table 7-5
 Backup Destination Event Types

Friendly Name	Event Type
Change Compartment	com.oraclecloud.databaseservice.changeb ackupdestinationcompartment
Create	com.oraclecloud.databaseservice.createb ackupdestination
Terminate	com.oraclecloud.databaseservice.deleteb ackupdestination
Update	com.oraclecloud.databaseservice.updateb ackupdestination

Example 7-65 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 (Cloud@Customer)

Review the list of event types that database nodes emit.

Table 7-6 Database Node Event Types

Friendly Name	Event Type
Update Begin	com.oraclecloud.databaseservice.dbno deaction.begin
Update End	<pre>com.oraclecloud.databaseservice.dbno deaction.end</pre>

Example 7-66 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 (Cloud@Customer)

Review the list of event types that Database Homes emit.

Table 7-7 Database Home Event Types

Friendly Name	Event Type
Create Begin	<pre>com.oraclecloud.databaseservice.created bhome.begin</pre>
Create End	<pre>com.oraclecloud.databaseservice.created bhome.end</pre>
Terminate Begin	<pre>com.oraclecloud.databaseservice.deleted bhome.begin</pre>
Terminate End	<pre>com.oraclecloud.databaseservice.deleted bhome.end</pre>
Update Begin	<pre>com.oraclecloud.databaseservice.updated bhome.begin</pre>
Update End	<pre>com.oraclecloud.databaseservice.updated bhome.end</pre>

Example 7-67 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 (Cloud@Customer)

}

Review the list of event types that databases emit.

Friendly Name	Event Type	
Create Begin	com.oraclecloud.databaseservice.crea tedatabase.begin	
Create End	com.oraclecloud.databaseservice.crea tedatabase.end	
Delete Backup Begin	com.oraclecloud.databaseservice.dele tebackup.begin	
Delete Backup End	com.oraclecloud.databaseservice.dele tebackup.end	
Restore Begin	com.oraclecloud.databaseservice.rest oredatabase.begin	
Restore End	com.oraclecloud.databaseservice.rest oredatabase.end	
Terminate Begin	com.oraclecloud.databaseservice.dele tedatabase.begin	
Terminate End	com.oraclecloud.databaseservice.dele tedatabase.end	
Update Begin	com.oraclecloud.databaseservice.upda tedatabase.begin	
Update End	com.oraclecloud.databaseservice.upda tedatabase.end	

Table 7-8 Database Event Types

Example 7-68 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",
      "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.

Friendly Name	Event Type
VM Cluster - Patch Begin	com.oraclecloud.databaseservice.patchvm cluster.begin
VM Cluster - Patch End	com.oraclecloud.databaseservice.patchvm cluster.end
DB Home - Patch Begin	<pre>com.oraclecloud.databaseservice.patchdb home.begin</pre>
DB Home - Patch End	com.oraclecloud.databaseservice.patchdb home.end
Database - Move Begin	com.oraclecloud.databaseservice.movedat abase.end

Table 7-9 Database and Grid Infrastructure Patching Events



Friendly Name	Event Type
Database - Move End	com.oraclecloud.databaseservice.movedat abase.end

Table 7-9 (Cont.) Database and Grid Infrastructure Patching Events

Example 7-69 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": {
      "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",
        "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"
    }
 }
}
```

Autonomous VM Cluster Event Types

Review the list of event types that Autonomous VM clusters emit.

Friendly Name	Event Type	
Autonomous VM Cluster - Change Compartment	com.oraclecloud.databaseservice.chan geautonomousvmclustercompartment	
Autonomous VM Cluster - Create Begin	com.oraclecloud.databaseservice.crea teautonomousvmcluster.begin	
Autonomous VM Cluster - Create End	com.oraclecloud.databaseservice.crea teautonomousvmcluster.end	
Autonomous VM Cluster - Terminate Begin	<pre>com.oraclecloud.databaseservice.dele teautonomousvmcluster.begin</pre>	
Autonomous VM Cluster - Terminate End	<pre>com.oraclecloud.databaseservice.dele teautonomousvmcluster.end</pre>	
Autonomous VM Cluster - Update Begin	com.oraclecloud.databaseservice.upda teautonomousvmcluster.begin	
Autonomous VM Cluster - Update End	com.oraclecloud.databaseservice.upda teautonomousvmcluster.end	

Table 7-10 Autonomous VM Cluster Events



Example 7-70 Autonomous VM Cluster Examples

This is a reference event for Autonomous VM Cluster - Change Compartment:

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.changeautonomousvmclustercompartment",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "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": "AutonomousVmCluster-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": "VmClusterNetwork-unique ID",
        "cpuCoreCount": "2",
        "availableCpus": "1",
        "dataStorageSizeInTBs": "4",
        "availableDataStorageSizeInTBs": "1",
        "licenseType": "BRING YOUR OWN LICENSE",
        "timeZone": "US/Pacific"
      }
   }
  }
```

This is a reference event for Autonomous VM Cluster - Create Begin

```
{
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.createautonomousvmcluster.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTypeVersion": "1.0",
    "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": "AutonomousVmCluster-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": "VmClusterNetwork-unique ID",
      "cpuCoreCount": "2",
      "availableCpus": "1",
      "dataStorageSizeInTBs": "4",
      "availableDataStorageSizeInTBs": "1",
      "licenseType": "BRING YOUR OWN LICENSE",
      "timeZone": "US/Pacific"
   }
 }
}
```

This is a reference event for Autonomous VM Cluster - Terminate Begin

```
{
   "cloudEventsVersion": "0.1",
   "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
   "eventType":
"com.oraclecloud.databaseservice.deleteautonomousvmcluster.begin",
   "source": "databaseservice",
   "eventTypeVersion": "1.0",
   "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": "AutonomousVmCluster-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": "VmClusterNetwork-unique_ID",
"cpuCoreCount": "2",
"availableCpus": "1",
"dataStorageSizeInTBs": "4",
"availableDataStorageSizeInTBs": "1",
"licenseType": "BRING_YOUR_OWN_LICENSE",
"timeZone": "US/Pacific"
}
```

This is a reference event for Autonomous VM Cluster - Update Begin

```
{
   "cloudEventsVersion": "0.1",
   "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
   "eventType":
"com.oraclecloud.databaseservice.updateautonomousvmcluster.begin",
   "source": "databaseservice",
   "eventTypeVersion": "1.0",
   "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": "AutonomousVmCluster-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": "VmClusterNetwork-unique ID",
        "cpuCoreCount": "2",
       "availableCpus": "1",
       "dataStorageSizeInTBs": "4",
       "availableDataStorageSizeInTBs": "1",
       "licenseType": "BRING YOUR OWN LICENSE",
       "timeZone": "US/Pacific"
      }
```

}

Autonomous Container Database Event Types

Review the list of event types that Autonomous Container Database emit.

 Table 7-11
 Autonomous Container Database Events

Friendly Name	Event Type
Change Compartment	com.oraclecloud.databaseservice.chan geautonomouscontainerdatabasecompart ment
Create Backup Begin	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.backup.beg in</pre>
Create Backup End	com.oraclecloud.databaseservice.auto nomous.container.database.backup.end
Create Begin	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.instance.c reate.begin</pre>
Create End	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.instance.c reate.end</pre>
Maintenance Begin	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.maintenanc e.begin</pre>
Maintenance End	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.maintenanc e.end</pre>
Maintenance Reminder	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.maintenanc e.reminder</pre>
Maintenance Scheduled	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.maintenanc e.scheduled</pre>
Restart Begin	com.oraclecloud.databaseservice.rest artautonomouscontainerdatabase.begin
Restart End	com.oraclecloud.databaseservice.rest artautonomouscontainerdatabase.end
Restore Begin	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.restore.be gin</pre>
Restore End	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.restore.en d</pre>



Friendly Name	Event Type	
Terminate Begin com.oraclecloud.databaseservi inateautonomouscontainerdatab in		
Terminate End	com.oraclecloud.databaseservice.term inateautonomouscontainerdatabase.end	
Update Begin	com.oraclecloud.databaseservice.auto nomous.container.database.instance.u pdate.begin	
Update End	<pre>com.oraclecloud.databaseservice.auto nomous.container.database.instance.u pdate.begin</pre>	

 Table 7-11 (Cont.) Autonomous Container Database Events

Example 7-71 Autonomous Container Database Examples

This is a reference event for create Autonomous Container Database:

```
{lifecycleState=AVAILABLE,
autonomousVmClusterId=ocid1.autonomousvmcluster.oc1.sea.abzwkljrevsjcfuskw7mg
i6ulzfg2epjjxhbnhwj63q7q3kzvuwg3djqnd2a,
displayName=CDB2-NFS,
dbName=wqxtzzfn,
dbUniqueName=wqxtzzfn sealtd,
freeTags={},
autonomousContainerDatabaseId=ocid1.autonomouscontainerdatabase.oc1.sea.abzwk
ljrxkuruqe3qzgw432adr5ug7ridmwi4ifvjlsahcdqdhbhzbf543xa,
compartmentId=ocid1.compartment.region1..aaaaaaaaass5x4witduxzgrs7qmzqk3m5kmog
uve7urwploqef6763w3o42ta,
timeUpdated=2020-06-15 21:52:24.085,
tenantId=ocid1.tenancy.region1..aaaaaaaagyw5okosjg54csr3u5qgaxvtjufz55537h44m
jy2umiqur4z5w3a,
timeCreated=2020-06-15 19:08:03.797,
id=ocid1.autonomouscontainerdatabase.oc1.sea.abzwkljrxkuruqe3qzgw432adr5ug7ri
dmwi4ifvjlsahcdqdhbhzbf543xa,
definedTags={}}
```

Autonomous Database Event Types

Review the list of event types that Autonomous Database emit.

Friendly Name	Event Type
Change Compartment Begin	com.oraclecloud.databaseservice.changea utonomousdatabasecompartment.begin
Change Compartment End	com.oraclecloud.databaseservice.changea utonomousdatabasecompartment.end

Table 7-12 Autonomous Database Events



Friendly Name	Event Type	
Create Backup Begin	com.oraclecloud.databaseservice.autonom ous.database.backup.begin	
Create Backup End	com.oraclecloud.databaseservice.autonom ous.database.backup.end	
Create Begin	com.oraclecloud.databaseservice.autonom ous.database.instance.create.begin	
Create End	com.oraclecloud.databaseservice.autonom ous.database.instance.create.end	
Restore Begin	com.oraclecloud.databaseservice.autonom ous.database.restore.begin	
Restore End	com.oraclecloud.databaseservice.autonom ous.database.restore.end	
Start Begin	com.oraclecloud.databaseservice.startau tonomousdatabase.begin	
Start End	com.oraclecloud.databaseservice.startau tonomousdatabase.end	
Stop Begin	com.oraclecloud.databaseservice.stopaut onomousdatabase.begin	
Stop End	com.oraclecloud.databaseservice.stopaut onomousdatabase.end	
Terminate Begin	com.oraclecloud.databaseservice.deletea utonomousdatabase.begin	
Terminate End	com.oraclecloud.databaseservice.deletea utonomousdatabase.end	
Update Begin	com.oraclecloud.databaseservice.updatea utonomousdatabase.begin	
Update End	com.oraclecloud.databaseservice.updatea utonomousdatabase.end	

Table 7-12	(Cont.) Autonomou	s Database Events
------------	-------------------	-------------------

Data Guard Event Types

Review the list of event types that Data Guard associations emit.

Friendly Name	Event Type
Data Guard Association - Create Begin	com.oraclecloud.databaseservice.crea tedataguardassociation.begin
Data Guard Association - Create End	com.oraclecloud.databaseservice.crea tedataguardassociation.end
Switchover Begin	com.oraclecloud.databaseservice.swit choverdataguardassociation.begin



Friendly Name	Event Type
Switchover End	com.oraclecloud.databaseservice.swit choverdataguardassociation.end
Failover Begin	<pre>com.oraclecloud.databaseservice.fail overdataguardassociation.begin</pre>
Failover End	com.oraclecloud.databaseservice.fail overdataguardassociation.end
Reinstate Begin	com.oraclecloud.databaseservice.rein statedataguardassociation.begin
Reinstate End	com.oraclecloud.databaseservice.rein statedataguardassociation.end

Table 7-13 (Cont.) Data Guard Association Events

Example 7-72 Data Guard Associations Examples

This is a reference event for Data Guard Association - Create Begin:

```
{
   "eventId": "022a63a4-ff77-11e9-a0af-f45c89b1cb17",
    "eventTime": "2019-11-05T02:50:21.000Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique id"
   },
    "eventType":
"com.oraclecloud.databaseservice.createdataguardassociation.begin",
   "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique id",
      "compartmentName": "example name",
      "resourceName": "my container database",
      "resourceId": "ocid1.dataguard.oc1.phx.unique id",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dataguard.oc1.phx.unique id",
        "timeCreated": "2019-06-27T21:15:59.000Z",
        "timeUpdated": "2019-06-27T21:16:04.389Z",
        "lifecycleState": "ACTIVE",
        "lifecycleMessage": "",
        "dbSystemId": "ocid1.vmcluster.oc1.phx.unique id",
        "DatabaseId": "ocid1.database.oc1.phx.unique id",
        "DbHomeId": "ocid1.dbhome.oc1.phx.unique id",
        "DGConfigId": "022a67c8-ff77-11e9-af6e-f45c89b1cb17",
        "DGConfigState": "SUCCESS",
        "LastSyncedTime": "2019-06-27T21:16:04.389Z",
        "ApplyLag": "2 hours",
```

```
"SyncState": "SYNCED",
    "dcsDgUpdateTimestamp": "2019-06-27T21:16:04.389Z",
    "lastUpdatedIdentifier": "022a6912-ff77-11e9-9e77-
f45c89b1cb17",
    "displayName": "Data Guard Association - Create Begin",
    "licenseType": "BRING_YOUR_OWN_LICENSE",
    "workloadType": "Transaction Processing"
    }
  }
}
```

This is a reference event for Data Guard Association - Create End:

```
{
    "eventId": "022aa7cc-ff77-11e9-90cd-f45c89b1cb17",
    "eventTime": "2019-11-05T02:50:21.000Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique id"
    },
    "eventType":
"com.oraclecloud.databaseservice.createdataguardassociation.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique id",
      "compartmentName": "example name",
      "resourceName": "my container database",
      "resourceId": "ocid1.dataguard.oc1.phx.unique id",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.dataguard.oc1.phx.unique id",
        "timeCreated": "2019-06-27T21:15:59.000Z",
        "timeUpdated": "2019-06-27T21:16:04.389Z",
        "lifecycleState": "ACTIVE",
        "lifecycleMessage": "",
        "dbSystemId": "ocid1.vmcluster.oc1.phx.unique id",
        "DatabaseId": "ocid1.database.oc1.phx.unique id",
        "DbHomeId": "ocid1.dbhome.oc1.phx.unique id",
        "DGConfigId": "022aab34-ff77-11e9-b89c-f45c89b1cb17",
        "DGConfigState": "SUCCESS",
        "LastSyncedTime": "2019-06-27T21:16:04.389Z",
        "ApplyLag": "2 hours",
        "SyncState": "SYNCED",
        "dcsDgUpdateTimestamp": "2019-06-27T21:16:04.389Z",
        "lastUpdatedIdentifier": "022aac10-ff77-11e9-8041-
f45c89b1cb17",
        "displayName": "Data Guard Association - Create End",
        "licenseType": "BRING YOUR OWN LICENSE",
        "workloadType": "Transaction Processing"
```



```
}
}
}
```

Autonomous Data Guard Association Event Types

Review the list of event types that Autnomous Data Guard associations emit.

Table 7-14 Autonomous Data Guard Association Events

Friendly Name	Event Type
Autonomous Data Guard Association -	com.oraclecloud.DatabaseService.CreateA
Create Autonomous Data Guard Begin	utonomousDataGuardAssociation.begin
Autonomous Data Guard Association –	com.oraclecloud.DatabaseService.CreateA
Create Autonomous Data Guard End	utonomousDataGuardAssociation.end
Autonomous Data Guard Association -	com.oraclecloud.DatabaseService.Switcho
Switchover Begin	verAutonomousDataguardAssociation.begin
Autonomous Data Guard Association -	com.oraclecloud.DatabaseService.Switcho
Switchover End	verAutonomousDataguardAssociation.end
Autonomous Data Guard Association -	com.oraclecloud.DatabaseService.Failove
Failover Begin	rAutonomousDataguardAssociation.begin
Autonomous Data Guard Association -	com.oraclecloud.DatabaseService.Failove
Failover End	rAutonomousDataguardAssociation.end
Autonomous Data Guard Association -	com.oraclecloud.DatabaseService.Reinsta
Reinstate Begin	teAutonomousDataGuardAssociation.begin
Autonomous Data Guard Association -	com.oraclecloud.DatabaseService.Reinsta
Reinstate End	teAutonomousDataGuardAssociation.end

Example 7-73 Autonomous Data Guard Associations Examples

This is a reference event for Autonomous Data Guard Association - Create Autonomous Data Guard Begin:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique ID>",
 "eventType":
"com.oraclecloud.DatabaseService.CreateAutonomousDataGuardAssociation.begin",
 "source": "databaseservice",
 "eventTypeVersion": "2.0",
 "eventTime": "2019-06-27T21:16:04Z",
 "contentType": "application/json",
 "data": {
   "eventGroupingId": "<unique ID>",
   "eventName"="CreateAutonomousDataGuardAssociation"
   "compartmentId": "ocid1.compartment.oc1..<unique ID>",
   "compartmentName": "example name",
   "resourceVersion":null,
   "resourceName": "my container database",
   "resourceId": "<unique_ID>",
   "availabilityDomain": "all",
```

```
"tagSlug": "<slug ID>",
    "definedTags": {},
    "additionalDetails": {
      "lifecycleState": "PROVISIONING",
      "DGConfigId"="91f298da-b890-42ce-935b-9b841e908756",
      "ApplyLag"=null,
      "LastRoleChangeTime"=null,
      "TransportLag"=null,
"autonomousContainerDatabaseId"="ocid1.autonomouscontainerdatabase.oc1.
sea.<unique ID>",
      "DGConfigState"=null,
      "lifeCycleMessage"=null,
      "lastUpdatedIdentifier"=null,
      "SyncState"=null,
"autonomousExadataInfrastructureId"="ocid1.autonomousvmcluster.oc1.sea.
<unique ID>",
      "timeUpdated"="2020-10-18 23:02:34.864",
      "timeCreated"="2020-10-18 23:02:34.864",
      "dbHomeId"="ocid1.autonomouspodhome.oc1.sea.<unique ID>",
      "LastSyncedTime"=null,
      "dcsDgUpdateTimestamp"=null,
   }
  }
}
```

This is a reference event for Autonomous Data Guard Association - Switchover Begin:

```
{
  "cloudEventsVersion": "0.1",
 "eventID": "<unique ID>",
  "eventType":
"com.oraclecloud.DatabaseService.CreateAutonomousDataGuardAssociation.b
egin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "<unique ID>",
    "eventName"="SwitchoverAutonomousDataguardAssociation"
    "compartmentId": "ocid1.compartment.oc1..<unique ID>",
    "compartmentName": "example name",
    "resourceVersion":null,
    "resourceName": "my container database",
    "resourceId": "ocid1.autonomousdgassociation.oc1.sea.<unique ID>",
    "availabilityDomain": "all",
    "tagSlug": "<slug ID>",
    "definedTags": {},
    "stateChange": {
        "previous"=null,
        "current: {
            "lifecycleState"="ROLE CHANGE IN PROGRESS
```



```
}
    }
    "additionalDetails": {
      "lifecycleState": "ROLE CHANGE IN PROGRESS",
      "DGConfigId"="<unique ID>",
      "ApplyLag"="0 seconds computed 2 seconds ago",
      "LastRoleChangeTime"=null,
      "TransportLag"="0 seconds computed 2 seconds ago",
"autonomousContainerDatabaseId"="ocid1.autonomouscontainerdatabase.oc1.sea.<u
nique ID>",
      "DGConfigState"="SUCCESS",
      "lifeCycleMessage"=null,
      "lastUpdatedIdentifier"=null,
      "SyncState"="SYNCED",
"autonomousExadataInfrastructureId"="ocid1.autonomousvmcluster.oc1.sea.<uniqu
e ID>",
      "timeUpdated"="2020-10-18 23:02:34.864",
      "timeCreated"="2020-10-18 23:02:34.864",
      "dbHomeId"="ocid1.autonomouspodhome.oc1.sea.<unique ID>",
      "LastSyncedTime"=null,
      "dcsDgUpdateTimestamp"=null,
    }
  }
}
```

This is a reference event for Autonomous Data Guard Association - Failover Begin:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique ID>",
 "eventType":
"com.oraclecloud.DatabaseService.CreateAutonomousDataGuardAssociation.begin",
 "source": "databaseservice",
 "eventTypeVersion": "2.0",
 "eventTime": "2019-06-27T21:16:04Z",
 "contentType": "application/json",
 "data": {
   "eventGroupingId": "<unique ID>",
   "eventName"="FailoverAutonomousDataguardAssociation"
   "compartmentId": "ocid1.compartment.oc1..<unique ID>",
   "compartmentName": "example name",
   "resourceVersion":null,
   "resourceName": "my container database",
   "resourceId": "ocid1.autonomousdgassociation.oc1.sea.<unique ID>",
   "availabilityDomain": "all",
   "tagSlug": "<slug ID>",
   "definedTags": {},
   "stateChange": {
       "previous"=null,
       "current: {
            "lifecycleState"="ROLE CHANGE IN PROGRESS
       }
```



```
}
    "additionalDetails": {
     "lifecycleState": "ROLE CHANGE IN PROGRESS",
      "DGConfigId"="<unique ID>",
      "ApplyLag"="0 seconds computed 2 seconds ago",
      "LastRoleChangeTime"=null,
      "TransportLag"="0 seconds computed 2 seconds ago",
"autonomousContainerDatabaseId"="ocid1.autonomouscontainerdatabase.oc1.
sea.<unique ID>",
      "DGConfigState"="SUCCESS",
      "lifeCycleMessage"=null,
      "lastUpdatedIdentifier"=null,
      "SyncState"="SYNCED",
"autonomousExadataInfrastructureId"="ocid1.autonomousvmcluster.oc1.sea.
<unique ID>",
      "timeUpdated"="2020-10-18 23:02:34.864",
      "timeCreated"="2020-10-18 23:02:34.864",
      "dbHomeId"="ocid1.autonomouspodhome.oc1.sea.<unique ID>",
      "LastSyncedTime"=null,
      "dcsDgUpdateTimestamp"=null,
   }
  }
}
```

This is a reference event for Autonomous Data Guard Association - Reinstate Begin:

```
{
  "cloudEventsVersion": "0.1",
  "eventID": "<unique ID>",
  "eventType":
"com.oraclecloud.DatabaseService.CreateAutonomousDataGuardAssociation.b
egin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "<unique ID>",
    "eventName"="ReinstateAutonomousDataGuardAssociation"
    "compartmentId": "ocid1.compartment.oc1..<unique ID>",
    "compartmentName": "example name",
    "resourceVersion":null,
    "resourceName": "my container database",
    "resourceId": "ocid1.autonomousdgassociation.oc1.sea.<unique ID>",
    "availabilityDomain": "all",
    "tagSlug": "<slug ID>",
    "definedTags": {},
    "stateChange": {
        "previous"=null,
        "current: {
            "lifecycleState"="ROLE CHANGE IN PROGRESS
        }
```

```
}
    "additionalDetails": {
     "lifecycleState": "ROLE_CHANGE_IN_PROGRESS",
      "DGConfigId"="<unique ID>",
      "ApplyLag"="0 seconds computed 2 seconds ago",
      "LastRoleChangeTime"=null,
      "TransportLag"="0 seconds computed 2 seconds ago",
"autonomousContainerDatabaseId"="ocid1.autonomouscontainerdatabase.oc1.sea.<u
nique ID>",
      "DGConfigState"="SUCCESS",
      "lifeCycleMessage"=null,
      "lastUpdatedIdentifier"=null,
      "SyncState"="SYNCED",
"autonomousExadataInfrastructureId"="ocid1.autonomousvmcluster.oc1.sea.<uniqu
e ID>",
      "timeUpdated"="2020-10-18 23:02:34.864",
      "timeCreated"="2020-10-18 23:02:34.864",
      "dbHomeId"="ocid1.autonomouspodhome.oc1.sea.<unique ID>",
      "LastSyncedTime"=null,
      "dcsDgUpdateTimestamp"=null,
    }
  }
}
```

Key Store Event Types

Review the list of event types that Key Store emits.

Friendly Name	Event Type
Key Store - Create	<pre>com.oraclecloud.databaseservice.createk eystore</pre>
Key Store - Update	<pre>com.oraclecloud.databaseservice.updatek eystore</pre>
Key Store - Terminate	<pre>com.oraclecloud.databaseservice.deletek eystore</pre>
Key Store - Change Compartment	<pre>com.oraclecloud.databaseservice.changek eystorecompartment</pre>
Autonomous Container Database - Rotate Key Begin	com.oraclecloud.databaseservice.rotatek eyautonomouscontainerdatabase.begin
Autonomous Container Database - Rotate Key End	com.oraclecloud.databaseservice.rotatek eyautonomouscontainerdatabase.end
Autonomous Database - Rotate Key Begin	com.oraclecloud.databaseservice.rotatek eyautonomousdatabase.begin
Autonomous Database - Rotate Key End	com.oraclecloud.databaseservice.rotatek eyautonomousdatabase.end

Table 7-15 Key Store Events



Example 7-74 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 Cloud@Customer Infrastructure Maintenance Event Types

Review the list of event types that Exadata Cloud@Customer Infrastructure Maintenance emits.

Friendly Name	Event Type	Event Message
Exadata Infrastructure - Virtual Machines Maintenance Begin	<pre>com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancevm.begi n</pre>	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of the Database Server component of your ExaCC Infrastructure instance <i><infra- name></infra- </i> , ocid <i><infra-ocid></infra-ocid></i> has started. A follow-up notice will be sent when the Database Server maintenance operation has completed.

 Table 7-16
 Exadata Cloud@Customer Infrastructure Maintenance Events



Friendly Name	Event Type	Event Message
Exadata Infrastructure - Virtual Machines Maintenance End	com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancevm.end	This is an Oracle Cloud Operations notice that quarterly maintenance update of the Database Server component of your ExaCC Infrastructure instance <infra-name>, ocid <infra- ocid>; Database Server <dbserver name=""> ocid <dbserver ocid=""> has completed.</dbserver></dbserver></infra- </infra-name>
Exadata Infrastructure - Maintenance Scheduled	com.oraclecloud.databas eservice.exaccinfrastru cturemaintenanceschedul ed	Rolling Oracle Cloud Operations is announcing the availability of a new quarterly maintenance update for ExaCC Infrastructure.
		Oracle has scheduled the installation of this new update on your service instance <infra-name>, ocid <infra- ocid> on <time- scheduled>.</time- </infra- </infra-name>
		The maintenance method for this maintenance is Rolling as selected per the maintenance preferences.
		Non Rolling
		Oracle Cloud Operations is announcing the availability of a new quarterly maintenance update for ExaCC Infrastructure.
		Oracle has scheduled the installation of this new update on your service instance <infra-name>, ocid <infra- ocid> on <time- scheduled>.</time- </infra- </infra-name>
		The maintenance method for this maintenance is Non Rolling as selected per the maintenance preferences.
		Non-rolling maintenance minimizes maintenance time but will result in full system downtime.

 Table 7-16 (Cont.) Exadata Cloud@Customer Infrastructure Maintenance

 Events



Friendly Name	Event Type	Event Message
Exadata Infrastructure - Maintenance Reminder	com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancereminde r	Rolling: This is an Oracle Cloud Operations reminder notice. Oracle has scheduled a quarterly maintenance update installation for ExaCC Infrastructure instance <infra-name>, ocid <ocid> in approximately <no-of- days> days on <time- scheduled>. The maintenance method for this maintenance is Rolling as selected per the maintenance preferences. Non Rolling: This is an Oracle Cloud Operations reminder notice.</time- </no-of- </ocid></infra-name>
		Oracle has scheduled a quarterly maintenance updat installation for ExaCC Infrastructure instance <infra-name>, ocid <ocid> in approximately <no-of- days> days on <time- scheduled>. The maintenance method for this maintenance is Non Rolling a selected per the maintenance preferences. Non-rolling maintenance time but will result in full system downtime</time- </no-of- </ocid></infra-name>
Exadata Infrastructure - Maintenance Begin	<pre>com.oraclecloud.databas eservice.exaccinfrastru cturemaintenance.begin</pre>	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update installation for your ExaCC Infrastructure instance <i><infra-name></infra-name></i> , ocid <i><infra< i=""> <i>ocid></i>. The update installation for the service started at <i><time scheduled=""></time></i>. A follow-up notice will be sen when the maintenance update operation has completed.</infra<></i>

Table 7-16 (Cont.) Exadata Cloud@Customer Infrastructure Maintenance Events



Friendly Name	Event Type	Event Message
- Maintenance End eservice.exaccinfi	com.oraclecloud.databas eservice.exaccinfrastru cturemaintenance.end	stru Cloud Operations notice that
		Failed: This is an Oracle Cloud Operations notice that your ExaCC Infrastructure quarterly maintenance update installation for service instance <infra-name>, ocid <infra- ocid> which started at <maintenance-start- time> has failed to complete due to technical reasons and the operations team are currently looking into the issue. You will receive regular notifications to track the progress of this maintenance.</maintenance-start- </infra- </infra-name>
Exadata Infrastructure - Maintenance Custom action time Begin	<pre>com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancecustoma ctiontime.begin</pre>	This is an Oracle Cloud Operations notice that the custom action timeout for your ExaCC Infrastructure instance <infra-name>, ocid <infra- ocid>; Database Server <dbserver name="">, ocid <dbserver ocid=""> has started. A follow-up notice will be sent when the custom action timeout has ended.</dbserver></dbserver></infra- </infra-name>
Exadata Infrastructure - Maintenance Custom action time End	com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancecustoma ctiontime.end	This is an Oracle Cloud Operations notice that the custom action timeout for your ExaCC Infrastructure instance <infra-name>, ocid <infra- ocid>; Database Server <dbserver name="">, ocid <dbserver ocid=""> has ended.</dbserver></dbserver></infra- </infra-name>

 Table 7-16 (Cont.) Exadata Cloud@Customer Infrastructure Maintenance

 Events



Friendly Name	Event Type	Event Message
Exadata Infrastructure - IB Switch maintenance Begin	<pre>com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancenetwork switches.begin</pre>	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of the network fabric switches component of your ExaCC Infrastructure instance <infra-name>, ocid <infra- ocid> has started. A follow-up notice will be sent when the network fabric switches maintenance operation has completed.</infra- </infra-name>
Exadata Infrastructure - IB Switch maintenance End	com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancenetwork switches.end	This is an Oracle Cloud Operations notice that quarterly maintenance update of the network fabric switches component of your ExaCC Infrastructure instance <infra-name>, ocid <infra- ocid> has completed.</infra- </infra-name>
Exadata Infrastructure - maintenance Storage servers Start	<pre>com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancestorage servers.start</pre>	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of Storage servers component of your ExaCC Infrastructure instance <infra-name>, ocid <infra- ocid> has started. A follow-up notice will be sent when storage servers maintenance operation has completed.</infra- </infra-name>
Exadata Infrastructure - maintenance Storage servers End	com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancestorage servers.end	This is an Oracle Cloud Operations notice that quarterly maintenance update of Storage servers component of your ExaCC Infrastructure instance <infra-name>, ocid <infra-ocid> has completed.</infra-ocid></infra-name>
Exadata Infrastructure - Maintenance Rescheduled	<pre>com.oraclecloud.databas eservice.exaccinfrastru cturemaintenanceresched uled</pre>	Oracle Cloud Operations is announcing reschedule of a quarterly maintenance update for ExaCC Infrastructure. A maintenance run has been rescheduled on your service instance <infra-name>, ocid <infra-ocid> to <new- schedule-time>.</new- </infra-ocid></infra-name>

Table 7-16	(Cont.) Exadata Cloud@Customer Infrastructure Maintenance
Events	



Friendly Name	Event Type	Event Message
Exadata Infrastructure - Maintenance Method Change	com.oraclecloud.databas eservice.exaccinfrastru cturemaintenancemethodc hange	Oracle Cloud Operations is announcing a change related to quarterly maintenance update for ExaCC Infrastructure. There's a change in maintenance method on your service instance <infra-name>, ocid <infra-ocid> to <new- maintenance-method>.</new- </infra-ocid></infra-name>
Exadata Infrastructure - Maintenance ReScheduled With Reason	<pre>com.oraclecloud.databas eservice.exaccinfrastru cturemaintenanceresched uledwithreason</pre>	Event message provided is a custom message sent from Oracle Cloud Operations.

 Table 7-16 (Cont.) Exadata Cloud@Customer Infrastructure Maintenance

 Events

Example 7-75 Exadata Cloud@Customer Infrastructure Quarterly Maintenance Events Examples

This is a reference event for Exadata Infrastructure - Virtual Machines Maintenance Begin:

```
"exampleEvent": {
 "cloudEventsVersion": "0.1",
 "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
 "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancevm.begin",
 "source": "databaseservice",
 "eventTypeVersion": "2.0",
 "eventTime": "2019-06-27T21:16:04.000Z",
 "contentType": "application/json",
 "data": {
   "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
   "eventName": "ExaccInfrastructureMaintenanceVM",
   "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": "detailed message",
            "shape": "ExadataCC.Base3.48",
            "displayName": "testDisplayName"
```

```
}
}
```

This is a reference event for Exadata Infrastructure - Virtual Machines Maintenance End:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
  "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancevm.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
    "eventName": "ExaccInfrastructureMaintenanceVM",
    "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",
          "timeEnded": "2019-08-29T12:30:00.000Z",
          "description": "ExaCC Infrastructure Maintenance
Notifications",
          "message": "detailed message",
          "shape": "ExadataCC.Base3.48",
          "displayName": "testDisplayName"
       }
  }
}
```

This is a reference event for Exadata Infrastructure - Maintenance Scheduled:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
    "com.oraclecloud.databaseservice.exaccinfrastructuremaintenanceschedule
d",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
```



```
"eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
       "eventName": "ExaccInfrastructureMaintenanceScheduled",
       "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",
           "description": "ExaCC Infrastructure Maintenance Notifications",
           "message": "detailed message",
           "shape": "ExadataCC.Base3.48",
           "displayName": "testDisplayName"
      }
  }
}
```

This is a reference event for Exadata Infrastructure - Maintenance Reminder:

```
"exampleEvent": {
   "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
   "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancereminder",
   "source": "databaseservice",
   "eventTypeVersion": "2.0",
   "eventTime": "2019-06-27T21:16:04.000Z",
   "contentType": "application/json",
   "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "ExaccInfrastructureMaintenanceReminder",
        "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",
             "description": "ExaCC Infrastructure Maintenance Notifications",
             "message": "detailed message",
             "shape": "ExadataCC.Base3.48",
             "displayName": "testDisplayName"
        ļ
```



```
This is a reference event for Exadata Infrastructure - Maintenance Begin:
```

}

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenance.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "ExaccInfrastructureMaintenance",
        "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": "detailed message",
            "shape": "ExadataCC.Base3.48",
            "displayName": "testDisplayName"
        }
    }
}
```

This is a reference event for Exadata Infrastructure - Maintenance End:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
    "com.oraclecloud.databaseservice.exaccinfrastructuremaintenance.end",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "ExaccInfrastructureMaintenance",
        "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",
          "timeEnded": "2019-08-29T12:30:00.000Z",
          "description": "ExaCC Infrastructure Maintenance Notifications",
          "message": "detailed message",
          "shape": "ExadataCC.Base3.48",
          "displayName": "testDisplayName"
     }
}
```

This is a reference event for Exadata Infrastructure - Maintenance Custom action time Begin:

```
"exampleEvent": {
   "cloudEventsVersion": "0.1",
   "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
   "eventType":
"com.oraclecloud.databaseservice.exaccinframtncustomactiontime.begin",
   "source": "databaseservice",
   "eventTypeVersion": "2.0",
   "eventTime": "2019-06-27T21:16:04.000Z",
   "contentType": "application/json",
   "data": {
      "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "ExaccInfraMtnCustomActionTime",
      "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"
     }
   }
 }
```



}

This is a reference event for Exadata Infrastructure - Maintenance Custom action time End:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
"com.oraclecloud.databaseservice.exaccinframtncustomactiontime.end",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "ExaccInfraMtnCustomActionTime",
      "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"
      }
    }
  }
```

This is a reference event for Exadata Infrastructure - maintenance Network Switches Begin:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b98",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenanceibswitch
.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "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"
  }
}
```

This is a reference event for Exadata Infrastructure - maintenance Network Switches End:

```
"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"
     }
   }
 }
```



}

This is a reference event for Exadata Infrastructure - maintenance Storage servers Start:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b55",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancestorages
ervers.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"
      }
    }
  1
```

This is a reference event for Exadata Infrastructure - maintenance Storage servers End:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b55",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancestorages
ervers.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": {},
    "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"
   }
 }
}
```

Example 7-76 Exadata Cloud@Customer Infrastructure Monthly Security Maintenance and Quaterly Maintenance Events Examples

Exadata Infrastructure - Maintenance Reminder - maintenanceSubType: QUARTERLY

```
"exampleEvent": {
   "cloudEventsVersion": "0.1",
   "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
   "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancereminder",
   "source": "databaseservice",
    "eventTypeVersion": "2.0",
   "eventTime": "2019-06-27T21:16:04.000Z",
   "contentType": "application/json",
   "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
       "eventName": "ExaccInfrastructureMaintenanceReminder",
        "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",
             "description": "ExaCC Infrastructure Maintenance Notifications",
             "message": "detailed message",
             "shape": "ExadataCC.Base3.48",
             "displayName": "testDisplayName"
             "maintenanceSubType": "QUARTERLY"
```

```
Exadata Infrastructure - Maintenance Reminder - maintenanceSubType: SECURITY_MONTHLY
```

}

}

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancereminder
",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "ExaccInfrastructureMaintenanceReminder",
        "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",
             "description": "ExaCC Infrastructure Maintenance
Notifications",
             "message": "detailed message",
             "shape": "ExadataCC.Base3.48",
             "displayName": "testDisplayName"
             "maintenanceSubType": "SECURITY MONTHLY"
        }
   }
}
```

Exadata Infrastructure - Maintenance Begin - maintenanceSubType: QUARTERLY

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenance.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
```

```
"eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
    "eventName": "ExaccInfrastructureMaintenance",
    "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": "detailed message",
        "shape": "ExadataCC.Base3.48",
        "displayName": "testDisplayName"
        "maintenanceSubType": "QUARTERLY"
    }
}
```

Exadata Infrastructure - Maintenance Begin - maintenanceSubType: SECURITY_MONTHLY

```
"exampleEvent": {
   "cloudEventsVersion": "0.1",
   "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
   "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenance.begin",
   "source": "databaseservice",
   "eventTypeVersion": "2.0",
   "eventTime": "2019-06-27T21:16:04.000Z",
   "contentType": "application/json",
   "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "ExaccInfrastructureMaintenance",
       "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": "detailed message",
            "shape": "ExadataCC.Base3.48",
```



}

```
"displayName": "testDisplayName"
"maintenanceSubType": "SECURITY_MONTHLY"
}
}
```

}

Exadata Infrastructure - Maintenance End - maintenanceSubType: QUARTERLY

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenance.end",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "ExaccInfrastructureMaintenance",
        "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",
             "timeEnded": "2019-08-29T12:30:00.000Z",
             "description": "ExaCC Infrastructure Maintenance
Notifications",
             "message": "detailed message",
             "shape": "ExadataCC.Base3.48",
             "displayName": "testDisplayName"
             "maintenanceSubType": "QUARTERLY"
        }
   }
}
```

Exadata Infrastructure - Maintenance End - maintenanceSubType: SECURITY_MONTHLY

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenance.end",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
```

```
"contentType": "application/json",
    "data": {
        "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "ExaccInfrastructureMaintenance",
        "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",
             "timeEnded": "2019-08-29T12:30:00.000Z",
             "description": "ExaCC Infrastructure Maintenance Notifications",
             "message": "detailed message",
             "shape": "ExadataCC.Base3.48",
             "displayName": "testDisplayName"
             "maintenanceSubType": "SECURITY MONTHLY"
        }
  }
}
```

Exadata Infrastructure - Maintenance Scheduled - maintenanceSubType: QUARTERLY

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
  "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenancescheduled",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "data": {
       "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "ExaccInfrastructureMaintenanceScheduled",
       "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",
           "description": "ExaCC Infrastructure Maintenance Notifications",
            "message": "detailed message",
```

```
"shape": "ExadataCC.Base3.48",
            "displayName": "testDisplayName"
            "maintenanceSubType": "QUARTERLY"
       }
   }
}
Exadata Infrastructure - Maintenance Scheduled - maintenanceSubType:
SECURITY_MONTHLY
"exampleEvent": {
   "cloudEventsVersion": "0.1",
   "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
   "eventType":
"com.oraclecloud.databaseservice.exaccinfrastructuremaintenanceschedule
d",
   "source": "databaseservice",
   "eventTypeVersion": "2.0",
   "eventTime": "2019-06-27T21:16:04.000Z",
   "contentType": "application/json",
   "data": {
       "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
       "eventName": "ExaccInfrastructureMaintenanceScheduled",
       "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",
            "description": "ExaCC Infrastructure Maintenance
Notifications",
            "message": "detailed message",
            "shape": "ExadataCC.Base3.48",
            "displayName": "testDisplayName"
            "maintenanceSubType": "SECURITY_MONTHLY"
       }
   }
}
```

Storage Expansion Event Types

Review the list of event types that storage expansion emits.

Friendly Name	Event Type
Exadata Infrastructure - Add Storage Capacity Begin	com.oraclecloud.databaseservice.addstor agecapacityexadatainfrastructure.begin
Exadata Infrastructure - Add Storage Capacity End	com.oraclecloud.databaseservice.addstor agecapacityexadatainfrastructure.end

Table 7-17 Storage Expansion Events

Example 7-77 Storage Expansion Events Examples

This is a reference event for Exadata Infrastructure - Add Storage Capacity Begin:

```
"exampleEvent": {
 "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.addstoragecapacityexadatainfrastructure.begi
n",
  "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.addstoragecapacityexadatainfrastructur
e.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": {
      "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.

Action **Event Name** Event Type Create Database Software CreateDatabaseSoftwareI com.oraclecloud.Databas Image eService.CreateDatabase mage CreateDatabaseSoftwareI SoftwareImage.begin com.oraclecloud.Databas mage eService.CreateDatabase SoftwareImage.end **Delete Database Software** DeleteDatabaseSoftwareI com.oraclecloud.Databas Image eService.DeleteDatabase mage DeleteDatabaseSoftwareI SoftwareImage.begin com.oraclecloud.Databas mage eService.DeleteDatabase SoftwareImage.end ListDatabaseSoftwareIma com.oraclecloud.Databas List Database Software Images ges eService.ListDatabaseSo ftwareImages

Table 7-18 Database Software Image Events



Action	Event Name	Event Type
Get Database Software Image	GetDatabaseSoftwareImag e	<pre>com.oraclecloud.Databas eService.GetDatabaseSof twareImage</pre>
Change compartment of Database Software Image	MoveDatabaseSoftwareIma ge MoveDatabaseSoftwareIma	com.oraclecloud.Databas eService.MoveDatabaseSo ftwareImage.begin
	ge	com.oraclecloud.Databas eService.MoveDatabaseSo ftwareImage.end
Update Database Software Image	UpdateDatabaseSoftwareI mage	com.oraclecloud.Databas eService.UpdateDatabase SoftwareImage

Table 7-18 (Cont.) Database Software Image Events

Example 7-78 Database Software Image Events Examples

This is a reference event for Create Database Software Image - Begin:

```
{
  "eventType":
"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..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xginczgm2z5d
64cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage 19c",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljsbpu3kxb54loym5sgkn2z4briz4xbokho
uxywkir764ht4txx4nya",
    "availabilityDomain": "AD2",
    "freeformTags": null,
    "definedTags": null,
    "identity": {
      "principalName": "xxxx",
      "principalId":
"ocid1.user.region1..aaaaaaaaedlsln4welqy3upwitxqrss3fnu6potqmds4xo3xld5sqyzw
bveq",
      "authType": "natv",
      "callerName": null,
      "callerId": null,
```



```
"tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqincz
qm2z5d64cn15q",
      "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"
      }
    },
    "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..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqincz
```

```
gm2z5d64cn15q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage 19c",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljsbpu3kxb54loym5sgkn2z4briz4xbokho
uxywkir764ht4txx4nya",
    "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
    },
    "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..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xgincz
qm2z5d64cn15q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage 202106152041",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljrzqt4tr326jtdmyudcziz6h5uhc
e36jbndxojgeg5kpwjcipxelga",
    "availabilityDomain": "AD2",
    "freeformTags": null,
    "definedTags": null,
    "identity": {
      "principalName": "xxx",
      "principalId":
"ocid1.user.region1..aaaaaaaaedlsln4welqy3upwitxqrss3fnu6potqmds4xo3xld
5sqyzwbveq",
      "authType": "natv",
      "callerName": null,
      "callerId": null,
      "tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqincz
gm2z5d64cn15q",
      "ipAddress": "160.34.124.245",
      "credentials": ""
    },
    "request": {
      "id": "csidd793a61a40689138814477fa5c6e/
ffe9b65c8dcf425c9b30ae1a30ae5687/B7CA65C1DDBEDB55052051EF1113DB73",
      "path": "/20160918/databaseSoftwareImages/
ocid1.databasesoftwareimage.oc1.sea.abzwkljrzgt4tr326jtdmyudcziz6h5uhce
36jbndxojgeg5kpwjcipxelga",
      "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/
ffe9b65c8dcf425c9b30ae1a30ae5687/B7CA65C1DDBEDB55052051EF1113DB73",
    "eventName": "DeleteDatabaseSoftwareImage",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xginczgm2z5d
64cn15g",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage 202106152041",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljrzgt4tr326jtdmyudcziz6h5uhce36jbn
dxojgeg5kpwjcipxelga",
    "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"
 }
}
```

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..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqincz
gm2z5d64cn15q",
    "compartmentName": "sic-dbaas",
    "resourceName": "",
    "resourceId": null,
    "availabilityDomain": "AD2",
    "freeformTags": null,
```



}

```
"definedTags": null,
    "identity": {
      "principalName": "xxx",
      "principalId":
"ocid1.user.region1..aaaaaaaaedlsln4welqy3upwitxqrss3fnu6potqmds4xo3xld5sqyzw
bveq",
      "authType": "natv",
      "callerName": null,
      "callerId": null,
      "tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d
64cnl5q",
      "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,
      "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..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xgincz
qm2z5d64cn15q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DBImage 19c",
    "resourceId":
"ocid1.databasesoftwareimage.oc1.sea.abzwkljsbpu3kxb54loym5sgkn2z4briz4
xbokhouxywkir764ht4txx4nya",
    "availabilityDomain": "AD2",
    "freeformTags": null,
    "definedTags": null,
    "identity": {
      "principalName": "xxx",
      "principalId":
"ocid1.user.region1..aaaaaaaalkib75hbawg6xh6rlys5ztumfx3cxf2ktwbon5ryv6
xjev7oectq",
      "authType": "natv",
      "callerName": null,
      "callerId": null,
      "tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqincz
gm2z5d64cn15q",
      "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.abzwkljsbpu3kxb54loym5sgkn2z4briz4x
bokhouxywkir764ht4txx4nya",
      "action": "GET",
      "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 7-19 Database Upgrade Events

Friendly Name	Event Type
MoveDatabase	com.oraclecloud.DatabaseService.MoveDat abase.begin
	com.oraclecloud.DatabaseService.MoveDat abase.end
UpgradeDatabase	com.oraclecloud.DatabaseService.Upgrade Database.begin
	com.oraclecloud.DatabaseService.Upgrade Database.end

Example 7-79 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",
  "eventTime": "2021-07-15T11:03:52.582Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csidff49902144a1a3c98bd4e5f49c1c/
37d3a7113f1e462fbd5a741dc5d52ef7/57472323A2FFABC31623FB7660F519AB",
    "eventName": "MoveDatabase",
    "compartmentId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d
64cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DB211507",
    "resourceId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3
pyiqhpq",
    "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"
          }
        },
        "displayName": "DB211507",
        "freeTags": {},
        "lifecycleState": "UPGRADING"
      }
    },
    "additionalDetails": {
      "databaseEdition": "ENTERPRISE EDITION EXTREME",
      "databaseId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o
4trsr3pyiqhpq",
      "dbHomeId":
"ocid1.dbhome.oc1.sea.abzwkljskjhqcc4bkzuai6kybx3hp3jznvxcgwlghqjnklpch
rktcynwnlog",
      "dbUniqueName": "Audit check",
```

```
"dbVersion": "18.14.0.0",
    "exadataInfrastructureId":
"ocidl.exadatainfrastructure.ocl.sea.abzwkljsqtwpumvorlttjuxcx6qtkxhuasvmyu4q
y3bng74sqolpwjunsrqa",
    "lifecycleState": "UPGRADING",
    "timeCreated": "2021-07-15T11:01:24Z",
    "timeUpdated": "2021-07-15T11:02:34Z",
    "vmClusterId":
"ocidl.vmcluster.ocl.sea.abzwkljsghfg7e3q3xh3gzqqxijt72fepczowmk35t554wxbpew7
txnooprq"
    }
}
```

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..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xginczgm2z5d
64cn15g",
    "compartmentName": "sic-dbaas",
    "resourceName": "DB211507",
    "resourceId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3
pyiqhpq",
    "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"
          }
        },
        "displayName": "DB211507",
        "freeTags": {},
        "lifecycleState": "AVAILABLE"
      }
    },
    "additionalDetails": {
      "databaseEdition": "ENTERPRISE EDITION EXTREME",
      "databaseId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o
4trsr3pyighpg",
      "dbHomeId":
"ocid1.dbhome.oc1.sea.abzwkljs7obwmwq2kgkjgeko4voudfcd5chu6iuonnew7uuxc
ulsvdralzda",
      "dbUniqueName": "Audit check",
      "dbVersion": "19.10.0.0",
      "exadataInfrastructureId":
"ocid1.exadatainfrastructure.oc1.sea.abzwkljsqtwpumvorlttjuxcx6qtkxhuas
vmyu4qy3bng74sqolpwjunsrqa",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2021-07-15T11:01:24Z",
      "timeUpdated": "2021-07-15T11:03:56Z",
      "vmClusterId":
"ocid1.vmcluster.oc1.sea.abzwkljsghfg7e3q3xh3gzqqxijt72fepczowmk35t554w
xbpew7txnooprg"
    }
  }
}
```



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..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xginczgm2z5d
64cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DB211507",
    "resourceId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcg7at54bpsmdricyg4o4trsr3
pyiqhpq",
    "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..aaaaaaaac5dskxmhmemrbyhaeabru4kgpituqmvgxjcm5y22mlxre26o
m4zq",
      "authType": "natv",
      "callerName": null,
      "callerId": null,
      "tenantId":
"ocid1.tenancy.region1..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xqinczgm2z5d
64cn15g",
      "ipAddress": "160.34.92.141",
      "credentials":
"ST$eyJraWQiOiJhc3dfcmVnaW9uMV8yMDIxLTA1LTE0IiwiYWxnIjoiUlMyNTYifQ.eyJzdWIiOi
JVY21kMS51c2VyLnJ1Z21vbjEuLmFhYWFhYWFhYzVkc2t4bWhtZW1yYnloYWVhYnJ1NGtncGl0dXF
tdmd4amNtNXkyMm1seHJlMjZvbTR6cSIsIm1mYV92ZXJpZmllZCI6ImZhbHNlIiwiaXNzIjoiYXV0
aFNlcnZpY2Uub3JhY2xlLmNvbSIsInB0eXBlIjoidXNlciIsInNlc3NfZXhwIjoiRnJpLCAxNiBKd
WwgMjAyMSAxMDo1OToxNyBVVEMiLCJhdWOiOiJvY2kiLCJwc3R5cGUiOiJuYXR2IiwidHR5cGUiOi
Jsb2dpbiIsImV4cCI6MTYyNjM1MDM1OSwiaWF0IjoxNjI2MzQ2NzU5LCJqdGkiOiJhNWE2YmRhNC0
20GJkLTQ0Zjkt0DFlNS0yZGRiZGMyZWU4MTQiLCJ0ZW5hbnQi0iJvY21kMS50ZW5hbmN5LnJ1Z21v
bjEuLmFhYWFhYWFhZTM0Z3diZXozZGJvN3VydGNmNDYyd2ozbW5tYW5iYmE1eHFpbmN6Z20yejVkN
jRjbmw1cSIsImp3ayI6IntcImtpZFwiOlwicHVia2V5LWU5Y2RmMWY4LTJjNTEtNGEyNi1hZjkyLT
```

```
NlZjlkNTlhNWEzN1wiLFwiblwiOlwicjE1azFxR0s1MXJaVUZzLVdFNDFrU2ZLRWozan1nd
1Z2bG9pZVM5Y3RLODJ4QmxZT1F3SkZwX19HUXVGTFktWk45MVZabXdSQkpFeHJYbE1hNHV6
aE5VS1BNQ3M50G9XTENpQmo2U1VMekFLZHVtbHZ1bHEwTUcwVmtMMjVXb1hONzU2N19wUUx
3NDc0T2w2VzNJU19Jd2tBVy1QZnF5QzdpY1NOekI2UXh2aG1VMXRpTUFUZ09RUjQ1Ty1md3
NnVn15WW14eVNCbm1sN1BDUEtoaGpBLT1FOFozeV94cn1PWFBuYj11Q1FrczZOblqzX08yR
kZ1UGxualN5QXdlSEZLSzd6MUJNVVdONUMtLV9Uemw5T0VaOENMeGdKNkpzeUUtRUVMeEZE
UkZ0MlVHZGNubUVBY2hNQ1qxNEtnWEFELVk5a2dkVS1HczNaUU1VTW5pYmFRXCIsXCJ1XCI
6XCJBUUFCXCIsXCJrdHlcIjpcIlJTQVwiLFwiYWxnXCI6XCJSUzI1NlwiLFwidXNlXCI6bn
VsbH0ifQ.GGPB9TMU Ckq0fLNCq02J6-
QwvHEEaY6jky5T92PBoPOBUkXb ruNOe6C8RwQhs2pmIsIBybTP7Ueene4cJ9nNKzsoTzNA
UrSp9714Be2noQgqJL-ZtFIApThz1vXdfFWZtQQInxx-jALpsvkhvD EVdbfx6FvlwDfm-
ReNDSnrUoFkfx8NpIrkqP8KUFhnX-
TRyOnUvqlaiB9C1L8FZGOJbBnOuiyloZe4DNP3Z39jViy37UntHeuMPnIqOYEsmPrSMTiJU
BAypy4 I5irWuR0PpxfrttP6nlEi80V5YrmO312-
ZJBRDKz8TdMkxclw3IT0qulTzFPqB3cXq45PGQ",
      "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.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4
trsr3pyiqhpq/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"
          }
        },
        "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.abzwkljstw7dymj6mfnaotkxpuie2bdxmlizqr5ie46de
wcuzdh35md7fuhq"
    }
}
```

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..aaaaaaaae34gwbez3dbo7urtcf462wj3mnmanbba5xginczgm2z5d
64cnl5q",
    "compartmentName": "sic-dbaas",
    "resourceName": "DB211507",
    "resourceId":
"ocid1.database.oc1.sea.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o4trsr3
pyiqhpq",
    "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": "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.abzwkljs5xinda6nughtf5tzapxhcq7at54bpsmdricyq4o
4trsr3pyiqhpq",
      "dbHomeId":
"ocid1.dbhome.oc1.sea.abzwkljs7obwmwq2kgkjgeko4voudfcd5chu6iuonnew7uuxc
ulsvdralzda",
      "dbUniqueName": "Audit check",
      "dbVersion": "19.10.0.0",
      "exadataInfrastructureId":
"ocid1.exadatainfrastructure.oc1.sea.abzwkljsqtwpumvorlttjuxcx6qtkxhuas
vmyu4qy3bnq74sqolpwjunsrqa",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2021-07-15T11:01:24Z",
      "timeUpdated": "2021-07-15T11:03:56Z",
      "vmClusterId":
"ocid1.vmcluster.oc1.sea.abzwkljsghfg7e3q3xh3gzqqxijt72fepczowmk35t554w
xbpew7txnooprg"
    }
  }
}
```

Pluggable Database Event Types

Review the list of event types that Pluggable Databases emit.



Friendly Name	Event Type
Pluggable Database - Create Begin	com.oraclecloud.databaseservice.createp luggabledatabase.begin
Pluggable Database - Create End	com.oraclecloud.databaseservice.createp luggabledatabase.end
Pluggable Database - Delete Begin	com.oraclecloud.databaseservice.deletep luggabledatabase.begin
Pluggable Database - Delete End	com.oraclecloud.databaseservice.deletep luggabledatabase.end
Pluggable Database - Local Clone Begin	com.oraclecloud.databaseservice.localcl onepluggabledatabase.begin
Pluggable Database - Local Clone End	<pre>com.oraclecloud.databaseservice.localcl onepluggabledatabase.end</pre>
Pluggable Database - Remote Clone Begin	<pre>com.oraclecloud.databaseservice.remotec lonepluggabledatabase.begin</pre>
Pluggable Database - Remote Clone End	<pre>com.oraclecloud.databaseservice.remotec lonepluggabledatabase.end</pre>
Start Pluggable Database - Begin	<pre>com.oraclecloud.databaseservice.startpl uggabledatabase.begin</pre>
Start Pluggable Database - End	com.oraclecloud.databaseservice.startpl uggabledatabase.end
Stop Pluggable Database - Start	com.oraclecloud.databaseservice.stopplu ggabledatabase.begin
Stop Pluggable Database - End	com.oraclecloud.databaseservice.stopplu ggabledatabase.end
Pluggable Database - Convert to Regular Begin	com.oraclecloud.databaseservice.pluggab ledatabase.converttoregular.begin
Pluggable Database - Convert to Regular End	com.oraclecloud.databaseservice.pluggab ledatabase.converttoregular.end
Pluggable Database - Inplace Restore Begin	com.oraclecloud.databaseservice.pluggab ledatabase.inplacerestore.begin
Pluggable Database - Inplace Restore End	com.oraclecloud.databaseservice.pluggab ledatabase.inplacerestore.end
Pluggable Database - Local Clone Begin	com.oraclecloud.databaseservice.pluggab ledatabase.localclone.begin
Pluggable Database - Local Clone End	<pre>com.oraclecloud.databaseservice.pluggab ledatabase.localclone.end</pre>
Pluggable Database - Refresh Begin	<pre>com.oraclecloud.databaseservice.pluggab ledatabase.refresh.begin</pre>
Pluggable Database - Refresh End	<pre>com.oraclecloud.databaseservice.pluggab ledatabase.refresh.end</pre>
Pluggable Database - Relocate Begin	com.oraclecloud.databaseservice.pluggab ledatabase.relocate.begin

Table 7-20	Pluggable	Database Events
------------	-----------	-----------------



Table 7-20	(Cont.) Pluggable Database Events
------------	-----------------------------------

Friendly Name	Event Type
Pluggable Database - Relocate End	<pre>com.oraclecloud.databaseservice.pluggab ledatabase.relocate.end</pre>
Pluggable Database - Remote Clone Begin	<pre>com.oraclecloud.databaseservice.pluggab ledatabase.remoteclone.begin</pre>
Pluggable Database - Remote Clone End	<pre>com.oraclecloud.databaseservice.pluggab ledatabase.remoteclone.end</pre>

Example 7-80 Pluggable Database Examples

This is a reference event for Pluggable Database - Create Begin:

```
{
  "eventID": "aflcdf4e-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",
      "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":
"com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.end",
     "eventTypeVersion": "2.0",
     "cloudEventsVersion": "0.1",
```

```
"source": "databaseservice",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocid1.compartment.oc1.....unique id",
    "compartmentName": "MyCompartment",
    "resourceName": "11092020 PKS PDB1",
    "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique id",
    "availabilityDomain": "XXIT:PHX-AD-1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.pluggabledatabases.oc1.phx.unique id",
      "isRefreshableClone": false,
      "timeCreated": "2021-03-13T21:15:59.000Z",
     "timeUpdated": "2021-03-13T21:15:59.000Z",
      "databaseId": "ocid1.database.oc1....unique id",
      "lifecycleState": "AVAILABLE",
      "displayName": "Pluggable Database - Convert to Regular End"
    }
  }
},
"activationTime": "2021-03-23T15:00:00.000Z",
"eventTypeVersion": "2.0"
```

This is a reference event for Pluggable Database - Inplace Restore Begin:

```
"exampleEvent": {
    "eventID": "unique id",
    "eventTime": "2021-03-23T00:49:14.123Z",
   "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique id"
   },
   "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.begin
",
    "eventTypeVersion": "2.0",
   "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.....unique id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020 PKS PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
```

}

```
"databaseId": "ocid1.database.oc1....unique_id",
    "lifecycleState": "RESTORE_IN_PROGRESS",
    "isRefreshableClone": false,
    "displayName": "Pluggable Database - Inplace Restore Begin"
    }
  }
},
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Inplace Restore End:

```
"exampleEvent": {
    "eventID": "unique id",
    "eventTime": "2021-03-23T00:49:14.123Z",
   "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.....unique id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020 PKS PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1....unique id",
        "lifecycleState": "AVAILABLE",
        "isRefreshableClone": false,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Inplace Restore End"
      }
   }
 },
 "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Local 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.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",
      "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",
"eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Refresh End:

```
"exampleEvent": {
    "eventID": "unique id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique id"
   },
   "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.refresh.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.....unique id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020 PKS PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1....unique id",
        "lifecycleState": "AVAILABLE",
        "isRefreshableClone": true,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Refresh End"
      }
   }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

This is a reference event for Pluggable Database - Relocate Begin:

```
"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
        "compartmentId": "ocid1.compartment.oc1..unique_id"
     },
     "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.relocate.begin",
```

```
"eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocid1.compartment.oc1.....unique id",
   "compartmentName": "MyCompartment",
    "resourceName": "11092020 PKS PDB1",
    "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique id",
    "availabilityDomain": "XXIT:PHX-AD-1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.pluggabledatabases.oc1.phx.unique id",
      "timeCreated": "2021-03-13T21:15:59.000Z",
      "timeUpdated": "2021-03-13T21:15:59.000Z",
      "databaseId": "ocid1.database.oc1....unique id",
      "lifecycleState": "AVAILABLE",
      "isRefreshableClone": false,
      "lifecycleDetails": "Pluggable Database is available",
      "displayName": "Pluggable Database - Relocate Begin"
   }
 }
},
"activationTime": "2021-03-23T15:00:00.000Z",
"eventTypeVersion": "2.0"
```

This is a reference event for Pluggable Database - Relocate End:

```
"exampleEvent": {
    "eventID": "unique id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique id"
   },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.relocate.end",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocid1.compartment.oc1.....unique id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020 PKS PDB1",
      "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.pluggabledatabases.oc1.phx.unique id",
```

}

```
"timeCreated": "2021-03-13T21:15:59.000Z",
    "timeUpdated": "2021-03-13T21:15:59.000Z",
    "databaseId": "ocid1.database.oc1....unique_id",
    "lifecycleState": "AVAILABLE",
    "lifecycleDetails": "Pluggable Database is available",
    "displayName": "Pluggable Database - Relocate End"
    }
  }
},
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}
```

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"
    },
    "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 7-21	VM Node Subsetting Events
------------	---------------------------

Friendly Name	Event Type
VM Cluster - Add Virtual Machine Begin	<pre>com.oraclecloud.databaseservice.vmclust eraddvirtualmachine.begin</pre>
VM Cluster - Add Virtual Machine End	<pre>com.oraclecloud.databaseservice.vmclust eraddvirtualmachine.end</pre>
VM Cluster - Terminate Virtual Machine Begin	com.oraclecloud.databaseservice.vmclust erterminatevirtualmachine.begin



Friendly Name	Event Type
VM Cluster - Terminate Virtual Machine End	<pre>com.oraclecloud.databaseservice.vmclust erterminatevirtualmachine.end</pre>

Table 7-21	(Cont.) VM Node Subsetting Events
------------	-----------------------------------

Example 7-81 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",
  "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",
    "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",
      "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 tfact1 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.

Manage Oracle Trace File Analyzer

• To check the run status of Oracle Trace File Analyzer, run the tfact1 status command as root or a non-root user:

```
# tfactl status
_____
-----.
| Host | Status of TFA | PID | Port | Version | Build
     | Inventory Status|
ID
+----+
| 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 tfact1 stop command as root:

```
# tfactl stop
Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..
```

Manage sysLens

 If sysLens is running, then once every 15 minutes data is collected in the local domU to discover the events to be reported.
 To check if sysLens is running, run the systemctl status syslens command as root in the domU:

```
# systemctl status syslens
\u25cf syslens.service
Loaded: loaded (/etc/systemd/system/syslens.service; disabled; vendor
preset: disabled)
Active: active (running) since Wed 2022-03-16 18:08:59 UTC; 34s ago
Main PID: 358039 (python3)
Memory: 31.6M
CGroup: /system.slice/syslens.service
\u2514\u2500358039 /usr/bin/python3 /var/opt/oracle/syslens/bin/
syslens_main.py --archive /var/opt/oracle/log/...
Mar 16 18:08:59 node1 systemd[1]: Started syslens.service.
Mar 16 18:09:09 node1 su[360495]: (to oracle) root on none
Mar 16 18:09:09 node1 su[360539]: (to grid) root on none
```

```
    If sysLens is enabled, when there is a reboot of the domU, then sysLens starts
automatically.
```

Mar 16 18:09:10 nodel su[360611]: (to grid) root on none Mar 16 18:09:11 nodel su[360653]: (to oracle) root on none

To validate if sysLens is enabled to collect telemetry, run the systemctl is-enabled syslens command as root in the domU:

```
# systemctl is-enabled syslens
enabled
```

• To validate if sysLens is configured to notify events, run the /usr/bin/syslens -- config /var/opt/oracle/syslens/data/exacc.syslens.config --get-key enable telemetry command as root in the domU:

```
# /usr/bin/syslens --config /var/opt/oracle/syslens/data/
exacc.syslens.config --get-key enable_telemetry
syslens Collection 2.3.3
on
```



Manage Database Service Agent

View the /opt/oracle/dcs/log/dcs-agent.log file to identify issues with the agent.

• To check the status of the Database Service Agent, run the systemctl status command:

```
# systemctl status dbcsagent.service
dbcsagent.service
Loaded: loaded (/usr/lib/systemd/system/dbcsagent.service; enabled;
vendor preset: disabled)
Active: active (running) since Fri 2022-04-01 13:40:19 UTC; 6min ago
Process: 9603 ExecStopPost=/bin/bash -c kill `ps -fu opc |grep
"java.*dbcs-agent.*jar" |awk '{print $2}' ` (code=exited, status=0/
SUCCESS)
Main PID: 10055 (sudo)
CGroup: /system.slice/dbcsagent.service
□ 10055 sudo -u opc /bin/bash -c umask 077; /bin/java -
Doracle.security.jps.config=/opt/oracle/...
```

• To start the agent if it is not running, run the systemctl start command as the root user:

systemctl start dbcsagent.service

Related Topics

- Using the Console to Create a VM Cluster
 To create your 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
- 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.



Friendly Name	Event Name	Remediation	Event Type	Threshold
Resource Jtilization - Disk Jsage	HEALTH.DB_GUE ST.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	HEALTH.DB_GU EST.FILESYSTE M.FREE_SPACE	com.oracleclo ud.databasese rvice.dbnode. critical	Critical threshold 90%
Disk Issues	only) HEALTH.DB_GUE ST.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_GU EST.FILESYSTE M.CORRUPTION	com.oracleclo ud.databasese rvice.dbnode. critical	N/A

 Table 7-22
 Database Service Events



Friendly Name	Event Name	Remediation	Event Type	Threshold
Memory Usage	CONFIGURATION .DB_GUEST.MEM ORY.HUGEPAGES _TOO_LARGE	CONFIGURATIO N.DB_GUEST.M EMORY.HUGEPA GES_TOO_LAR GE	<pre>com.oracleclo ud.databasese rvice.dbnode. critical</pre>	90%
	An event of type CRITICAL is created when the amount of memory in the VM configured for HugePages is 90% or more of the total VM memory.	-		
sshd Configuration	CONFIGURATION .DB_GUEST.SSH D.INVALID	CONFIGURATIO N.DB_GUEST.SS HD.INVALID	com.oracleclo ud.databasese rvice.dbnode.	N/A
	An event of type CRITICAL is created if unexpected values are set in the /etc/ssh/ sshd_config file.		critical	

 Table 7-22
 (Cont.) Database Service Events



Friendly Name	Event Name	Remediation	Event Type	Threshold
Oracle EXAchk Reported Issues	HEALTH.DB_CLU STER.EXACHK.C RITICAL_ALERT Oracle EXAchk is Exadata database platform's holistic health check that	HEALTH.DB_CL USTER.EXACHK .CRITICAL_ALE RT	com.oracleclo ud.databasese rvice.dbnode. critical	N/A
CRS status Un/	includes software, infrastructure and database configuration checks. CRITICAL check alerts should be addressed in 24 hours to maintain the maximum stability and availability of your system. This database service event alerts every 24 hours whenever there are any CRITICAL checks that are flagged in the most recent Oracle EXAchk report. The event will point to the latest Oracle EXAchk zip report.		com oraclecio	Ν/Α
CRS status Up/ Down	AVAILABILITY. DB_GUEST.CRS_ INSTANCE.DOWN.	AVAILABILITY.DB _GUEST.CRS_IN STANCE.DOWN	ud.databasese rvice.dbnode.	N/A
	An event of type CRITICAL is created when the Cluster Ready Service (CRS) is detected to be down.		critical (if .DOWN and NOT "user_action")	

 Table 7-22
 (Cont.) Database Service Events



Friendly Name Ev	vent Name	Remediation	Event Type	Threshold
DE IN C Ar IN cr de th do	VAILABILITY. 3_GUEST.CRS_ ISTANCE.DOWN CLEARED IFORMATION is eated once it is eatermined that e event for CRS own has eared.	N/A	<pre>com.oracleclo ud.databasese rvice.dbnode. information (if.DOWN_CLEA RED)</pre>	

 Table 7-22
 (Cont.) Database Service Events



Friendly Name	Event Name	Remediation	Event Type	Threshold
SCAN Listener Jp/Down	AVAILABILITY. DB_CLUSTER.SC AN_LISTENER.D OWN 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_CLEARE D event is created when a SCAN listener is started. There are three SCAN listeners per cluster called LISTENER_SCA N[1,2,3].	AVAILABILITY.DB _CLUSTER.SCA N_LISTENER.DO WN		N/A

Table 7-22 ((Cont.)	Database	Service	Events
--------------	---------	----------	---------	--------

Friendly Name	Event Name	Remediation	Event Type	Threshold
	AVAILABILITY. DB_CLUSTER.SC AN_LISTENER.D OWN_CLEARED An event of type INFORMATION is created once it is determined that the event for SCAN Listener down has cleared.	N/A	com.oracleclo ud.databasese rvice.dbnode. information (if.DOWN_CLEA RED)	

 Table 7-22
 (Cont.) Database Service Events



Friendly Name	Event Name	Remediation	Event Type	Threshold
Net Listener Up/ Down	AVAILABILITY. DB_GUEST.CLIE NT_LISTENER.D OWN 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 (srvct1) or Listener Control (lsnrct1) 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_CLEARE D event is created when a client listener is started. There is one client listener per node, each called LISTENER.	AVAILABILITY.DB _GUEST.CLIENT _LISTENER.DO WN		N/A

Table 7-22	(Cont.)	Database	Service	Events
------------	---------	----------	---------	--------



Friendly Name	Event Name	Remediation	Event Type	Threshold
	AVAILABILITY. DB_GUEST.CLIE NT_LISTENER.D OWN_CLEARED An event of type INFORMATION is created once it is determined that the event for Client Listener down has cleared.	N/A	com.oracleclo ud.databasese rvice.databas e.information (if.DOWN_CLEA RED)	
CDB Up/Down	Cleared. 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_CLEARE D event is created when a database instance is started.	AVAILABILITY.DB GUEST.CDB_IN STANCE.DOWN		N/A

 Table 7-22
 (Cont.) Database Service Events



Friendly Name	Event Name	Remediation	Event Type	Threshold
	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.oracleclo ud.databasese rvice.databas e.information (if.DOWN_CLEA RED)	
Critical DB Errors	HEALTH.DB_CLU STER.CDB.CORR UPTION	HEALTH.DB_CL USTER.CDB.CO RRUPTION	com.oracleclo ud.databasese rvice.databas e.critical	N/A
	Database corruption has been detected on your primary or standby database. The database alert.log is parsed for any specific errors that are indicative of physical block corruptions, logical block corruptions, or logical block corruptions caused by lost writes.			
Other DB Errors	HEALTH.DB_CLU STER.CDB.ARCH IVER_HANG	HEALTH.DB_CL USTER.CDB.AR CHIVER_HANG	com.oracleclo ud.databasese rvice.databas e.critical	N/A
	An event of type CRITICAL is created if a CDB is either unable to archive the active online redo log or unable to archive the active online redo log fast enough to the log archive destinations.			

 Table 7-22
 (Cont.) Database Service Events



Friendly Name	Event Name	Remediation	Event Type	Threshold
	HEALTH.DB_CLU STER.CDB.DATA BASE_HANG An event of type CRITICAL is created when a process/session hang is detected in the CDB.	HEALTH.DB_CL USTER.CDB.DA TABASE_HANG		
Backup Failures	HEALTH.DB_CLU STER.CDB.BACK UP_FAILURE	HEALTH.DB_CL USTER.CDB.BA CKUP_FAILURE	com.oracleclo ud.databasese rvice.databas e.critical	N/A
	An event of type CRITICAL is created if there is a CDB backup with a FAILED status reported in the v\$rman_status view.		e.criticai	
Disk Group Usage	HEALTH.DB_CLU STER.DISK_GRO UP.FREE_SPACE	HEALTH.DB_CL USTER.DISK_G ROUP.FREE_SP	com.oracleclo ud.databasese rvice.dbsyste	Notifications are sent when the usage hits 70%,
	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%.	ACE	<pre>m.critical com.oracleclo ud.databasese rvice.dbsyste m.information (if < 90%)</pre>	80%, 90%, and 100% with a corresponding severity of 4, 3, 2, and 1.

Table 7-22 (Cont.) Database Service Events

Example 7-82 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",
   "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",
```



```
"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 ${\tt tfactl blackout}$ command to temporarily suppress automatic diagnostic collections.

If you set blackout for a target, then Oracle Trace File Analyzer stops automatic diagnostic collections if it finds events in the alert logs for that target while scanning. By default, blackout will be in effect for 24 hours.

You can also restrict automatic diagnostic collection at a granular level, for example, only for **ORA-00600** or even only **ORA-00600** with specific arguments.

Syntax

```
tfactl blackout add|remove|print
-targettype host|crs|asm|asmdg|database|dbbackup|db_dataguard|
db_tablespace|pdb_tablespace|pdb|listener|service|os
-target all|name
[-container name]
[-pdb pdb_name]
-event all|"event_str1, event_str2"|availability
[-timeout nm|nh|nd|none]
[-c|-local|-nodes "node1, node2"]
[-reason "reason for blackout"]
[-docollection]
```

Parameters

Table 7-23	tfactl blackout Command Parameters
------------	------------------------------------

Parameter	Description
add remove print	Adds, removes, or prints blackout conditions.
targettype <i>type</i>	Limits blackout only to the specified target type.
Target type: host crs asm asmdg database dbbackup	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.
db_dataguard db_tablespace	crs: Blackout the availability of the Oracle Clusterware resource or events in the Oracle Clusterware logs.
pdb_tablespace pdb listener service os	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 <i>name</i>	Specify the target for blackout. You can specify a comma-delimited list of targets.
	By default, the target is set to all.

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 " <i>str1,str2</i> "	Limits blackout only to the availability events, or event strings, which should not trigger auto collections, or be marked as blacked out in telemetry JSON.
	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 <i>nh nd</i> none	Specify the duration for blackout in number of hours or days before timing out. By default, the timeout is set to 24 hours (24h).
c local	Specify if blackout should be set to cluster-wide or local.
	By default, blackout is set to local.
reason <i>comment</i>	Specify a descriptive reason for the blackout.
docollection	Use this option to do an automatic diagnostic collection even if a blackout is set for this target.

Table 7-23 (Cont.) tfactl blackout Command Parameters

Example 7-83 tfactl blackout

• To blackout event: ORA-00600 on target type: database, target: mydb

tfactl blackout add -targettype database -target mydb -event "ORA-00600"

To blackout event: ORA-04031 on target type: database, target: all

tfactl blackout add -targettype database -target all -event "ORA-04031" - timeout 1h

• To blackout db backup events on target type: dbbackup, target: mydb

tfactl blackout add -targettype dbbackup -target mydb

To blackout db dataguard events on target type: db_dataguard, target: mydb

tfactl blackout add -targettype db dataguard -target mydb -timeout 30m

 To blackout db tablespace events on target type: db_tablespace, target: system, container: mydb

tfactl blackout add -targettype db_tablespace -target system -container mydb -timeout 30m



• To blackout ALL events on target type: host, target: all

tfactl blackout add -targettype host -event all -target all timeout 1h -reason "Disabling all events during patching"

To print blackout details

tfactl blackout print

_____ _____ _____. myhostname +-----+-----+ | Target Type | Target | Events | Start | End Time | Status | Do Time 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 | | NA false | DB DATAGUARD | MYDB | Thu Mar 24 ALL 16:41:38 UTC 2022 | Thu Mar 24 17:11:38 UTC 2022 | ACTIVE | false | NA | DBBACKUP | MYDB | ALL | Thu Mar 24 16:40:47 UTC 2022 | Fri Mar 25 16:40:47 UTC 2022 | ACTIVE | false | NA | DB TABLESPACE | SYSTEM CDBNAME MYDB | ALL | Thu Mar 24 16:45:56 UTC 2022 | Thu Mar 24 17:15:56 UTC 2022 | ACTIVE | false | NA +----+

 To remove blackout for event: ORA-00600 on target type: database, target: mydb

tfactl blackout remove -targettype database -event "ORA-00600" - target mydb



To remove blackout for db backup events on target type: dbbackup, target: mydb

tfactl blackout remove -targettype dbbackup -target mydb

 To remove blackout for db tablespace events on target type: db_tablespace, target: system, container: mydb

tfactl blackout remove -targettype db_tablespace -target system container mydb

To remove blackout for host events on target type: host, target: all

tfactl blackout remove -targettype host -event all -target all

Remediation

These topics cover some common issues you might run into and how to address them.

- HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE
- 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.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:

 Remove unneeded files and/or directories created manually or by customerinstalled 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 Exadata Database Service on Cloud@Customer Systems Learn to manage Oracle Database homes on 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:

- 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
AuthorizedKeysFile	.ssh/authorized_keys This file must exist in root user home directory.
HostbasedAuthentication	no
IgnoreUserKnownHosts	yes
IgnoreRhosts	yes



SSHD Configuration Setting	Expected Value
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_TELEPHONE LC_MEASUREMENT LC_IDENTIFICATION LC_ALL LANGUAGE XMODIFIERS
Subsystem	<pre>sftp /usr/libexec/openssh/sftp- server</pre>
Protocol	2
AddressFamily	inet

Risk: SSHD configuration is unexpected which may cause Oracle Cloud automation failure or prevent customer SSH access to the VM.

Action: Change SSHD to match expected configuration.

1. Verify SSHD service is active.

```
$ sudo systemctl is-active sshd.service
active
```



If SSHD service is inactive, then start it.

\$ sudo systemctl start sshd.service

2. Verify SSHD service is enabled.

```
$ sudo /opt/oracle.cellos/host_access_control ssh-service --status
[INFO] [IMG-SEC-1201] Service sshd is enabled {1}
```

If SSHD service is disabled, then enable it.

- \$ sudo /opt/oracle.cellos/host access control ssh-service --enable
- 3. Change SSHD configuration to match the expected values according to the table shown in the Problem Statement section above.

SSHD Configuration Setting	How to Change Current setting
Ciphers	/opt/oracle.cellos/ host_access_control sshciphershelp
MACs	/opt/oracle.cellos/ host_access_control ssh-macshelp
PermitRootLogin	/opt/oracle.cellos/ host_access_control rootsshhelp
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-listenhelp
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
- 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.
- If CRS is still down, then investigate the cause of the failure by referring to the alert.log found in /u01/app/grid/diag/crs/<node_name>/crs/trace. Review the log entries corresponding to the date/time of the down event. Act on any potential remediation.
- 4. Restart the CRS, by issuing the crsctl start crs command.



5. A successful restart of CRS will generate the clearing event: AVAILABILITY.DB GUEST.CRS INSTANCE.DOWN CLEARED.

AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN

Problem Statement: A SCAN listener is down and unable to accept application connections.

Risk: If all SCAN listeners are down, then application connections to the database through the SCAN listener will fail.

Action:

Start the SCAN listener to receive the DOWN CLEARED event.

DOWN event of type INFORMATION

- 1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Grid Infrastructure software update, then no action is required. The affected SCAN listener will automatically failover to an available instance.
- 2. If the event was caused by user action, then start the SCAN listener at the next opportunity.

DOWN event of type CRITICAL

Check SCAN status and restart the SCAN listener.

1. Login to the VM as opc user and sudo to the grid user:

[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.

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
```

- Recheck the client listener status on any node: If the client listener is still down, then investigate the cause of the client listener failure:
 - a. Use tfactl to collect both the CRS and operating system logs 30 minutes prior and 10 minutes for the <hostName> indicated in the log. Note the time in the event payload is always provided in UTC. For tfactl collection, adjust the time to the timezone of the VM Cluster.

[grid@vm ~] tfactl diagcollect -crs -os -node <hostName> -from "<eventTime adjusted for local vm timezone> - 30 minute " -to "<eventTime adjusted for local vm timezone> + 10 minutes"

b. Review the listener log located under /u01/app/grid/diag/tnslsnr/ <hostName>/<listenerName>/trace

AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN



Problem Statement: A database instance has gone down.

Risk: A database instance has gone down, which may result in reduced performance if database instances are available on other nodes in the cluster, or complete downtime if database instances on all nodes are down.

Action:

Start the database instance to receive the DOWN CLEARED event.

DOWN event of type INFORMATION

- 1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Database Home software update, then no action is required. The affected database instance will automatically restart when maintenance affecting the instance is complete.
- 2. If the event was caused by user action, then start the affected database instance at the next opportunity.

DOWN event of type CRITICAL

- 1. Check database status and restart the down database instance.
 - a. Login to the VM as oracle user:
 - b. Set the environment:

```
[oracle@vm ~] . <dbName>.env
```

c. Check the database status:

[oracle@vm ~] srvctl status database -db <dbName>

d. Start the database instance:

```
[oracle@vm ~] srvctl start instance -db <dbName> -instance
<instanceName>
```

- 2. Investigate the cause of the database instance failure.
 - a. Review Trace File Analyzer (TFA) events for the database:

```
[oracle@vm ~] tfactl events -database <dbName> -instance
<instanceName>
```

b. Review the database alert log located at \$ORACLE_BASE/diag/rdbms/ <dbName>/<instanceName>/trace/alert <instanceName>.log

HEALTH.DB_CLUSTER.CDB.CORRUPTION

Problem Statement: Corruptions can lead to application or database errors and in worse case result in significant data loss if not addressed promptly.

A corrupt block is a block that was changed so that it differs from what Oracle Database expects to find. Block corruptions can be categorized as physical or logical:

• In a physical block corruption, which is also called a media corruption, the database does not recognize the block at all; the checksum is invalid or the block contains all zeros. An



example of a more sophisticated block corruption is when the block header and footer do not match.

 In a logical block corruption, the contents of the block are physically sound and pass the physical block checks; however, the block can be logically inconsistent. Examples of logical block corruption include incorrect block type, incorrect data or redo block sequence number, corruption of a row piece or index entry, or data dictionary corruptions.

For more information, see *Physical and Logical Block Corruptions*. All you wanted to know about it. (Doc ID 840978.1).

Block corruptions can also be divided into interblock corruption and intrablock corruption:

- In an intrablock corruption, the corruption occurs in the block itself and can be either a physical or a logical block corruption.
- In an interblock corruption, the corruption occurs between blocks and can only be a logical block corruption.

Oracle checks for the following errors in the alert.log:

- ORA-01578
- ORA-00752
- ORA-00753
- ORA-00600 [3020]
- ORA-00600 [kdsgrp1]
- ORA-00600 [kclchkblk_3]
- ORA-00600 [13013]
- ORA-00600 [5463]

Risk: A data corruption outage occurs when a hardware, software, or network component causes corrupt data to be read or written. The service-level impact of a data corruption outage may vary, from a small portion of the application or database (down to a single database block) to a large portion of the application or database (making it essentially unusable). If remediation action is not taken promptly, then potential downtime and data loss can increase.

Action:

The current event notification currently triggers on physical block corruptions (ORA-01578), lost writes (ORA-00752, ORA-00753 and ORA-00600 with first argument 3020), and logical corruptions (typical detected from ORA-00600 with first argument of kdsgrp1, kdsgrp1, kclchkblk_3, 13013 OR 5463).

Oracle recommends the following steps:

- Confirm that these corruptions were reported in the alert.log trace file. Log a Service Request (SR) with latest EXAchk report, excerpt of the alert.log and trace file containing the corruption errors, any history of recent application, database or software changes and any system, clusterware and database logs for the same time period. For all these cases, a TFA collection should be available and should be attached to the SR.
- 2. For repair recommendations, refer to Handling Oracle Database Corruption Issues (Doc ID 1088018.1).



For physical corruptions or ORA-1578 errors, the following notes will be helpful:

- Doc ID 1578.1 : OERR: ORA-1578 "ORACLE data block corrupted (file # %s, block # %s)" Primary Note
- Doc ID 472231.1 : How to identify all the Corrupted Objects in the Database reported by RMAN
- Doc ID 819533.1 : How to identify the corrupt Object reported by ORA-1578 / RMAN / DBVERIFY
- Depending on the object that has the corruption, follow the guidance in Doc ID 1088018.1. Note RMAN can be used to recover one or many data block that are physically corrupted. Also using Active Data Guard with real time apply, auto block repair of physical data corruptions would have occurred automatically.

For logical corruptions caused by lost writes (ORA-00752, ORA-00753 and ORA-00600 with first argument 3020) on the primary or standby databases, they will be detected on the primary or with standby's redo apply process. The following notes will be helpful:

- Follow the guidance, follow Doc ID 1088018.1.
- If you have a standby and lost write corruption on the primary or standby, refer to Resolving ORA-00752 or ORA-00600 [3020] During Standby Recovery (Doc ID 1265884.1)

For logical corruptions (typical detected from ORA-00600 with arguments of kdsgrp1, 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, <code>\$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.</code>
 - If the number of redo logs groups per thread is less than 4, then consider adding additional logs groups to reach 4, see item 1 below for add redo log steps.
 - The other possible solution is to resize the redo logs, see item 2 below for resizing steps.
- For Data Guard and Non Data Guard review the *Configure Online Redo Logs Appropriately* of section Oracle Database High Availability Overview and Best Practices for sizing guidelines.
- **1.** Add a redo log group for each thread. The additional redo log should equal the current log size.
 - a. Use the following query:

select max(group#) Ending_group_number, thread#, count(*)
number_of_groups_per_thread, bytes redo_size_in_bytes from v\$log
group by thread#,bytes

b. Add one new group per thread using the same size as the current redo logs.

alter database add logfile thread <thread_number> Group <max
group + 1> ('<DATA DISKGROUP>') size <redo size in bytes>

- 2. Resize the online redo logs by adding larger redo logs and dropping the current smaller redo logs.
 - a. Use the following query:

select max(group#) Ending_group_number, thread#, count(*)
number_of_groups_per_thread, bytes redo_size_in_bytes from v\$log
group by thread#,bytes

b. Add the same number of redo logs for each thread <number_of_groups_per_thread> that currently exist. The <new_redo_size_in_bytes> should be based on Configure Online Redo Logs Appropriately of section Oracle Database High Availability Overview and Best Practices.



- i. alter database add logfile thread <thread_number> Group <max group
 + 1> ('<DATA DISKGROUP>') size <new redo size in bytes>
- The original smaller redo logs should be deleted. A redo log can only be deleted if its status is inactive.
 To determine the status of a redo logs issue:

select group#, thread#, status, bytes from v\$log order by bytes, group#, thread#;

To delete the original smaller redo logs:

alter database drop logfile <group#>

 If the database is hung, the primary log archive destination and alternate may be full. Review the HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE for details on freeing space in RECO and DATA disk groups.

Related Topics

- Script To Find Redolog Switch History And Find Archivelog Size For Each Instances In RAC (Doc ID 2373477.1)
- Configure Online Redo Logs Appropriately
- HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

HEALTH.DB_CLUSTER.CDB.DATABASE_HANG

Problem Statement: Hang management detected a process hang and generated a ORA-32701 error message. Additionally, this event may be raised if Diagnostic Process (DIA0) process detects a hang in a critical database process.

Risk: A hang can indicate resource, operating system, or application coding related issues.

Action:

Investigate the cause of the session hang.

 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

 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-mmdd_24hh:mm:ss.zzzzzzzzz.log within the obkup directory. The logs are located on the lowest active node/instance of the database when the backup was initiated.
 - Review the log for any failures:
 - * If the failure is due to an external event outside of RMAN, for example the backup location was full or a networking issue, resolve the external issue.
 - * For other RMAN script errors, collect the diagnostic logs and open a Service Request. See DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check.
 - If the issue is transient or is resolved, take a new incremental backup: See *dbaascli database backup*.
- For Customer owned and managed backup taken through RMAN:
 - Review the RMAN logs for the backup.

Related Topics

- DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check
- dbaascli database backup
 To configure Oracle Database with a backup storage destination, take database
 backups, query backups, and delete a backup, use the dbaascli database
 backup command.

HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

Problem Statement: ASM disk group space usage is at or exceeds 90%.

Risk: Insufficient ASM disk group space can cause database creation failure, tablespace and data file creation failure, automatic data file extension failure, or ASM rebalance failure.

Action:



6.11

ASM disk group used space is determined by the running the following query while connected to the ASM instance.

18874368 17720208

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:

- 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

ora.RECOC1.dq

- 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.

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
AUGICKECENCIONDB	days.
	Default value: 1



Parameter	Description and Default Value
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

Archiving Alert Logs and Listener Logs

When cleaning up alert and listener logs, cleandblogs first archives and compresses the logs, operating as follows:

- **1**. The current log file is copied to an archive file that ends with a date stamp.
- 2. The current log file is emptied.
- 3. The archive file is compressed using gzip.
- 4. Any existing compressed archive files older than the retention period are deleted.

Running the cleandblogs Script Manually

The cleandblogs script automatically runs daily on each compute node, but you can also run the script manually if the need arises.

1. Connect to the compute node as the oracle user to clean log and diagnostic files that are associated with Oracle Database, or connect as the grid user to clean log and diagnostic files that are associated with Oracle Grid Infrastructure.



For detailed instructions, see Connecting to a Virtual Machine with SSH.

Change to the directory containing the cleandblogs script:

```
$ cd /var/opt/oracle/cleandb
```

2. Run the cleandblogs script:

```
$ ./cleandblogs.pl
```

When running the script manually, you can specify an alternate configuration file to use instead of cleandblogs.cfg by using the --pfile option:

```
$ ./cleandblogs.pl --pfile config-file-name
```

3. Close your connection to the compute node:

\$ exit

Related Topics

- Automatic Diagnostic Repository (ADR)
- Connecting to a Virtual Machine with SSH You can connect to the virtual machines in an Exadata Database Service on Cloud@Customer system by using a Secure Shell (SSH) connection.

Compute Expansion Event Types

Review the list of event types that compute expansion emits.

Table 7-24 Compute Expansion Events

Friendly Name	Event Type
VM Cluster Network - Add DB Server Network Begin	<pre>com.oraclecloud.databaseservice.addd bservervmclusternetwork.begin</pre>
VM Cluster Network - Add DB Server Network Begin	<pre>com.oraclecloud.databaseservice.addd bservervmclusternetwork.end</pre>
VM Cluster Network - Remove DB Server Network Begin	<pre>com.oraclecloud.databaseservice.remo vedbservervmclusternetwork.begin</pre>
VM Cluster Network - Remove DB Server Network End	<pre>com.oraclecloud.databaseservice.remo vedbservervmclusternetwork.end</pre>

Example 7-84 Compute Expansion Event Types Examples

This is a reference event for VM Cluster Network - Add DB Server Network Begin:

```
{
   "id":
   "ocid1.eventschema.oc1.phx.abyhqljrdsa7ybmw6mjndfutmibmw6a4ztvu7sgm562o
   jbudo3le6qwsspdq",
        "exampleEvent": {
            "cloudEventsVersion": "0.1",
            "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",
  "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"] }
 1,
  "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.abyhqljrupbmw5ksic23bmwnqc4rssbmwc23eslydoumcnr6rc
um6mtmgg3a",
   "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",
  "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.abyhqljrdsa7ybmw6mjndfutmifv46a4ztvu7sgm562o
   jbudo3le6qwsspdq",
    "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" },
   { "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"] }
 1,
  "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.abyhqljrupbmw5ksic23bmwnqc4rsskxwc23eslydoumcnr6rc
um6mtmgg3a",
   "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" },
   { "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 7-25
 Serial Console Connection Events

Friendly Name	Event Type
DB Node Console Connection - Create Begin	com.oraclecloud.databaseservice.crea tedbnodeconsoleconnection.begin



Friendly Name	Event Type
DB Node Console Connection - Create End	com.oraclecloud.databaseservice.crea tedbnodeconsoleconnection.end
DB Node Console Connection - Delete Begin	<pre>com.oraclecloud.databaseservice.dele tedbnodeconsoleconnection.begin</pre>
DB Node Console Connection - Delete End	com.oraclecloud.databaseservice.dele tedbnodeconsoleconnection.end
DB Node Console Connection - Update	com.oraclecloud.databaseservice.upda tedbnodeconsoleconnection
DB Node - Update	com.oraclecloud.databaseservice.upda tedbnode

Table 7-25	(Cont.)	Serial	Console	Connection Events
------------	---------	--------	---------	--------------------------

Example 7-85 Serial Console Connection Event Types Examples

This is a reference event for DB Node Console Connection - Create Begin:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsoleconnection.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique ID",
   "freeFormTags": { },
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique ID",
      "lifecycleState": "CREATING",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique ID",
      "tenantId": "ocid1.tenant.oc1..unique ID",
      "compartmentId": "ocid1.compartment.oc1..unique ID"
   }
 }
}
```



This is a reference event for DB Node Console Connection - Create End:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsoleconnection.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique ID"
  },
  "data": {
   "compartmentId": "ocid1.compartment.oc1..unique ID",
   "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique ID",
   "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique ID",
      "lifecycleState": "ACTIVE",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique ID",
      "tenantId": "ocid1.tenant.oc1..unique ID",
      "compartmentId": "ocid1.compartment.oc1..unique ID"
   }
  }
}
```

This is a reference event for DB Node Console Connection - Delete Begin:

```
"exampleEvent": {
 "cloudEventsVersion": "0.1",
 "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
 "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.begin",
 "source": "databaseservice",
 "eventTypeVersion": "1.0",
 "eventTime": "2019-08-29T21:16:04.000Z",
 "contentType": "application/json",
 "extensions": {
   "compartmentId": "ocid1.compartment.oc1..unique ID"
 },
 "data": {
   "compartmentId": "ocid1.compartment.oc1..unique ID",
   "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique ID",
   "freeFormTags": {},
   "definedTags": {},
   "additionalDetails": {
     "id": "ocid1.dbnodeconsoleconnection.oc1..unique ID",
      "lifecycleState": "DELETING",
```

```
"timeCreated": "2019-08-29T12:00:00.000Z",
    "timeUpdated": "2019-08-29T12:30:00.000Z",
    "lifecycleDetails": "detail message",
    "dbnodeId": "ocid1.dbnode.oc1..unique_ID",
    "tenantId": "ocid1.tenant.oc1..unique_ID",
    "compartmentId": "ocid1.compartment.oc1..unique_ID"
}
```

This is a reference event for DB Node Console Connection - Delete End:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique ID"
  },
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique ID",
    "resourceId": "ocid1.dbnodeconsoleconnection.oc1..unique ID",
   "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.dbnodeconsoleconnection.oc1..unique ID",
      "lifecycleState": "DELETED",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocid1.dbnode.oc1..unique ID",
      "tenantId": "ocid1.tenant.oc1..unique ID",
      "compartmentId": "ocid1.compartment.oc1..unique ID"
   }
  }
}
```

This is a reference event for DB Node Console Connection - Update:

```
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
    "eventType":
    "com.oraclecloud.databaseservice.updatedbnodeconsoleconnection",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "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",
      "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 7-86 Serial Console Connection Audit Event Example

This is a reference event for Serial Console Connection:

```
{
 "eventType": "VmConsoleConnected",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "VmConsoleConnectionAPI",
 "eventId": "2367d627-cff8-11ed-bfd3-02001714f979",
 "eventTime": "2023-03-31T19:13:37.120Z",
 "contentType": "application/json",
 "data": {
   "eventGroupingId": "2367d62d-cff8-11ed-bfd3-02001714f979",
   "eventName": "VmConsoleConnected",
   "compartmentId": "ocid1.compartment.oc1..<TRUNCATED>aaaaxxxxx",
   "compartmentName": "exacc-dev",
   "resourceName": "",
   "resourceId":
"ocid1.dbnodeconsoleconnection.oc1.iad.<TRUNCATED>aaaaaaxxxxx",
   "availabilityDomain": null,
   "freeformTags": null,
   "definedTags": null,
   "identity": {
      "principalName": "dsaes",
      "principalId": "ocid1.user.oc1..<TRUNCATED>aaaaaaaaaaaxxxxxxxxx,
      "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.

Table 7-26 Serial Console History Events

User Action	Event Type	Friendly Name	Event Type
Create Console History	Async	DB Node Console History - Create Begin	<pre>com.oraclecloud.d atabaseservice.cr eatedbnodeconsole history.begin</pre>



User Action	Event Type	Friendly Name	Event Type
		DB Node Console History - Create End	<pre>com.oraclecloud.d atabaseservice.cr eatedbnodeconsole history.end</pre>
Terminate Console History	Async	DB Node Console History - Delete Begin	com.oraclecloud.d atabaseservice.de letedbnodeconsole history.begin
		DB Node Console History - Delete End	com.oraclecloud.d atabaseservice.de letedbnodeconsole history.end
Update Console History	Sync	DB Node Console History - Update	com.oraclecloud.d atabaseservice.up datedbnodeconsole history
Get Console History Content	Sync	DB Node Console History - Get Content	com.oraclecloud.d atabaseservice.ge tdbnodeconsolehis torycontent

Table 7-26 (Cont.) Serial Console History Events

Example 7-87 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",
        "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",
    "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": "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",
    "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"
}
```

Interim Software Updates Event Types

Friendly Name	Event Type
Oneoff Patch - Create Begin	<pre>com.oraclecloud.databaseservice.createo neoffpatch.begin</pre>
Oneoff Patch - Create End	<pre>com.oraclecloud.databaseservice.createo neoffpatch.end</pre>
Oneoff Patch - Delete Begin	<pre>com.oraclecloud.databaseservice.deleteo neoffpatch.begin</pre>
Oneoff Patch - Delete End	<pre>com.oraclecloud.databaseservice.deleteo neoffpatch.end</pre>
Oneoff Patch - Download Begin	<pre>com.oraclecloud.databaseservice.downloa doneoffpatch.begin</pre>
Oneoff Patch - Download End	com.oraclecloud.databaseservice.downloa doneoffpatch.end

Table 7-27 Interim Software Updates Events

Example 7-88 One-Off Patches Event Types Examples

This is a reference event for Oneoff Patch - Create Begin:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhgljrsllp7rfneajgg2knxbgopwux24za7gzoe3mfj2bzfx
tnwqcxpbcq",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique ID",
      "compartmentName": "example name",
      "resourceName": "my oneoffpatch",
      "resourceId": "OneOffPatch-unique ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
```



```
"id": "ocid1.id..oc1...unique ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test patch set"
      }
   }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Create Begin",
  "eventType":
"com.oraclecloud.databaseservice.createoneoffpatch.begin",
  "additionalDetails": [
    { "name": "id", "type": "string" },
   { "name": "lifecycleState", "type": "string" },
   { "name": "timeCreated", "type": "string" },
   { "name": "displayName", "type": "string" },
   { "name": "dbVersion", "type": "string" },
   { "name": "patchType", "type": "string" },
   { "name": "patchShapeFamily", "type": "string" },
   { "name": "releaseUpdate", "type": "string" }
 1,
  "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

This is a reference event for Oneoff Patch - Create End:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrj4vvuph4qvj5eateeel6axblhkq3caqndgmj
vwl3sldpgb255j2q",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.createoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique ID",
      "compartmentName": "example name",
      "resourceName": "my oneoffpatch",
      "resourceId": "OneOffPatch-unique ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique ID",
```

```
"lifecycleState": "AVAILABLE",
      "timeCreated": "2020-08-26T12:00:00.000Z",
      "displayName": "testDisplayName",
      "databaseVersion": "19.6.0.0",
      "patchSet": "test patch set"
    }
  }
},
"serviceName": "Database",
"displayName": "Oneoff Patch - Create End",
"eventType": "com.oraclecloud.databaseservice.createoneoffpatch.end",
"additionalDetails": [
  { "name": "id", "type": "string" },
  { "name": "lifecycleState", "type": "string" },
  { "name": "timeCreated", "type": "string" },
  { "name": "displayName", "type": "string" },
  { "name": "dbVersion", "type": "string" },
  { "name": "patchType", "type": "string" },
  { "name": "patchShapeFamily", "type": "string" },
  { "name": "releaseUpdate", "type": "string" }
1,
"timeCreated": "2020-06-26T13:31:31.979Z"
```

```
This is a reference event for Oneoff Patch - Delete Begin:
```

}

```
{
 "id":
"ocid1.eventschema.oc1.phx.abyhqljrdripga5rryplwmv4ws6hqzr3pjyl7wfvoaqutvg2ey
2vtycn5ong",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.begin",
   "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique ID",
      "compartmentName": "example name",
      "resourceName": "my oneoffpatch",
      "resourceId": "OneOffPatch-unique ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
```



```
"databaseVersion": "19.6.0.0",
       "patchSet": "test patch set"
     }
   }
 },
 "serviceName": "Database",
 "displayName": "Oneoff Patch - Delete Begin",
 "eventType":
"com.oraclecloud.databaseservice.deleteoneoffpatch.begin",
 "additionalDetails": [
   { "name": "id", "type": "string" },
   { "name": "lifecycleState", "type": "string" },
   { "name": "timeCreated", "type": "string" },
   { "name": "displayName", "type": "string" },
   { "name": "dbVersion", "type": "string" },
   { "name": "patchType", "type": "string" },
   { "name": "patchShapeFamily", "type": "string" },
   { "name": "releaseUpdate", "type": "string" }
 ],
 "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

```
This is a reference event for Oneoff Patch - Delete End:
```

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrgwk2gvx5lmx6fiwotqdy32mdmrnkyzznz37d
pb4mmehgzt37vl7a",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType":
"com.oraclecloud.databaseservice.deleteoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique ID",
      "compartmentName": "example name",
      "resourceName": "my oneoffpatch",
      "resourceId": "OneOffPatch-unique ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
```



```
"patchSet": "test patch set"
     }
   }
 },
 "serviceName": "Database",
 "displayName": "Oneoff Patch - Delete End",
 "eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.end",
 "additionalDetails": [
   { "name": "id", "type": "string" },
   { "name": "lifecycleState", "type": "string" },
   { "name": "timeCreated", "type": "string" },
   { "name": "displayName", "type": "string" },
   { "name": "dbVersion", "type": "string" },
   { "name": "patchType", "type": "string" },
   { "name": "patchShapeFamily", "type": "string" },
   { "name": "releaseUpdate", "type": "string" }
 1,
 "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

```
This is a reference event for Oneoff Patch - Download Begin:
```

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljr3vkb7klt5hkbsnqzjaxmszsqomanlbqmr2tsrcq7xa
fcv2c7412q",
  "exampleEvent": {
   "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.begin",
    "source": "databaseservice",
   "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique ID"
   },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique ID",
      "compartmentName": "example name",
      "resourceName": "my oneoffpatch",
      "resourceId": "OneOffPatch-unique ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test patch set"
      }
   }
```



```
},
 "serviceName": "Database",
 "displayName": "Oneoff Patch - Download Begin",
 "eventType":
"com.oraclecloud.databaseservice.downloadoneoffpatch.begin",
 "additionalDetails": [
   { "name": "id", "type": "string" },
   { "name": "lifecycleState", "type": "string" },
   { "name": "timeCreated", "type": "string" },
   { "name": "displayName", "type": "string" },
   { "name": "dbVersion", "type": "string" },
   { "name": "patchType", "type": "string" },
   { "name": "patchShapeFamily", "type": "string" },
   { "name": "releaseUpdate", "type": "string" }
 ],
 "timeCreated": "2020-06-26T13:31:31.979Z"
}
```

This is a reference event for Oneoff Patch - Download End:

```
{
  "id":
"ocid1.eventschema.oc1.phx.abyhqljrn2lruez55ah56kqksi5qfq6m7iqvven7o2qk
ahlk5tkwrj5ll3oa",
  "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": "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"
```

Policy Details for Exadata Database Service on Cloud@Customer

Learn to write policies to control access to 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 Exadata Database Service on Cloud@Customer instances".

- About Resource-Types Learn about resource-types you can use in your policies.
- Resource-Types for Exadata Database Service on Cloud@Customer Review the list of resource-types specific to 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 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 databasefamily is equivalent to writing eight separate policies for the group that would grant access to the exadata-infrastructures, vmcluster-networks, vmclusters, backupdestinations, 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 autonomousdatabase-family is equivalent to writing four separate policies for the group that would grant access to the autonomous-databases, autonomous-backups, autonomouscontainer-databases, and cloud-autonomous-vmclusters resource-types.

For more information, see Resource-Types.

Resource-Types for Exadata Database Service on Cloud@Customer

Review the list of resource-types specific to Exadata Database Service on Cloud@Customer.

Aggregate Resource-Type

database-family



Individual Resource-Types

exadata-infrastructures

vmclusters

backup-destinations

db-nodes

db-homes

databases

backups

database-software-images

autonomous-vmclusters

autonomous-container-databases

autonomous-databases

```
key-stores
```

 $autonomous {\tt ContainerDatabaseDataguardAssociations}$

AutonomousDatabaseDataguardAssociation

dbnode-console-connection

dbnode-console-history

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 Understand the level of access of each verb.
- exadata-infrastructures Review the list of permissions and API operations for exadata-infrastructures resource-type.
- vmcluster-networks
 Review the list of permissions and API operations for vmcluster-networks
 resource-type.
- vmclusters Review the list of permissions and API operations for vmclusters resource-type.
- backup-destinations Review the list of permissions and API operations for backup-destinations resource-type.
- db-nodes Review the list of permissions and API operations for db-nodes resource-type.
- db-homes Review the list of permissions and API operations for db-homes resource-type.
- databases
 Review the list of permissions and API operations for databases resource-type.
- backups Review the list of permissions and API operations for backups resource-type.
- database-software-image Review the list of permissions and API operations for database-software-image resource-type.
- autonomous-databases Review the list of permissions and API operations for autonomous-databases resource-type.
- autonomous-backups Review the list of permissions and API operations for autonomous-backups resource-type.
- autonomous-container-databases Review the list of permissions and API operations for autonomous-containerdatabases resource-type.
- autonomous-vmclusters Review the list of permissions and API operations for autonomous-vmclusters resource-type.



- autonomousContainerDatabaseDataguardAssociations Review the list of permissions and API operations for autonomousContainerDatabaseDataguardAssociations resource-type.
- AutonomousDatabaseDataguardAssociation Review the list of permissions and API operations for AutonomousDatabaseDataguardAssociation resource-type.
- autonomous-virtual-machine Review the list of permissions and API operations for autonomous-virtual-machine resource-type.
- key-stores Review the list of permissions and API operations for key-store resource-type.
- pluggable-databases (PDBs) Review the list of permissions and API operations for pluggable-databases resourcetype.
- dbServers
 Review the list of permissions and API operations for dbServers resource-type.
- dbnode-console-connection Review the list of permissions and API operations for dbnode-console-connection resource-type.
- dbnode-console-history Review the list of permissions and API operations for dbnode-console-history resourcetype.
- oneoffPatch

Related Topics

- Permissions
- Verbs
- Resource-Types

Database-Family Resource Types

Understand the level of access of each verb.

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.

exadata-infrastructures

Review the list of permissions and API operations for exadata-infrastructures resource-type.

Granting permissions on exadata-infrastructure resources grants permissions on associated vmcluster-network resources.



Permission	APIs Fully Covered	APIs Partially Covered
EXADATA_INFRASTRUCTURE_ INSPECT	ListExadataInfrastructu res	ChangeExadataInfrastruc tureCompartment
	GetExadataInfrastructur e	
	GenerateRecommendedNetw orkDetails	
	ListVmClusterNetworks	
	GetVmClusterNetwork	
	ValidateVmClusterNetwor k	
	DownloadExadataInfrastr uctureConfigFile	
	DownloadVmClusterNetwor kConfigFile	

Table 7-28 INSPECT

Table 7-29 READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT+	none	none
EXADATA_INFRASTRUCTURE_ CONTENT_READ		

Table 7-30 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + EXADATA INFRASTRUCTURE	ActivateExadataInfrastr ucture	CreateVmCluster (also needs manage vmclusters)
UPDATE	UpdateExadataInfrastruc ture	UpdateVmCluster (also needs use vmclusters)
	ChangeExadataInfrastruc tureCompartment	ChangeExadataInfrastruc tureCompartment
	AddStorageCapacityExada taInfrastructure	
	CreateVmClusterNetwork	
	UpdateVmClusterNetwork	
	DeleteVmClusterNetwork	

Permissions	APIs Fully Covered	APIs Partially Covered
USE + EXADATA INFRASTRUCTURE CR	CreateExadataInfrastructu re	none
EATE	DeleteExadataInfrastructu re	
LETE — —	downloadExadataInfrastruc tureConfigFile	

Table 7-31 MANAGE

vmcluster-networks

Review the list of permissions and API operations for vmcluster-networks resource-type.

vmcluster-network resources inherit permissions from the exadata-infrastructure resources with which they are associated. You cannot grant permissions to vmcluster-network resources explicitly.

Table 7-32 INSPECT

Permission	APIs Fully Covered	APIs Partially Covered
EXADATA_INFRASTRUCTURE_IN	ListVmClusterNetworks	none
SPECT	GetVmClusterNetwork	
	ValidateVmClusterNetwork	

Table 7-33 READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + EXADATA_INFRASTRUCTURE_CO NTENT_READ	DownloadVmClusterNetworkC onfigFile	none

Table 7-34 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ +	CreateVmClusterNetwork	none
EXADATA_INFRASTRUCTURE_UP	UpdateVmClusterNetwork	
DATE	DeleteVmClusterNetwork	



Permissions	APIs Fully Covered	APIs Partially Covered
USE+	none	none
EXADATA_INFRASTRUCTURE_ CREATE		
EXADATA_INFRASTRUCTURE_ DELETE		

Table 7-35 MANAGE

vmclusters

Review the list of permissions and API operations for vmclusters resource-type.

Permission	APIs Fully Covered	APIs Partially Covered
VM_CLUSTER_INSPECT	ListVmClusters	none
	GetVmCluster	
	ListVmClusterPatches	
	ListVmClusterPatchHisto ryEntries	
	GetVmClusterPatch	
	GetVmClusterPatchHistor yEntry	
	ListVmClusterUpdates	
	ListVmClusterUpdateHist oryEntries	
	GetVmClusterUpdate	
	GetVmClusterUpdateHisto ryEntry	

Table 7-36 INSPECT

Table 7-37 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	No extra	none

Permissions	APIs Fully Covered	APIs Partially Covered
READ + ChanVM_CLUSTER_UPDATEt	ChangeVmClusterCompartmen t	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

Table 7-38 USE

Table 7-39 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE + VM_CLUSTER_CREATE VM_CLUSTER_DELETE	No extra	CreateVmCluster (also needs use exadata- infrastructures)
		DeleteVmCluster (also needs use exadata- infrastructures)

backup-destinations

Review the list of permissions and API operations for backup-destinations resource-type.

Table 7-40 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
BACKUP_DESTINATION_INSPEC	ListBackupDestinations	none
Т	GetBackupDestination	

Table 7-41 READ

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none



Permissions	APIs Fully Covered	APIs Partially Covered
READ +	UpdateBackupDestination	none
BACKUP_DESTINATION_UPDA TE	ChangeBackupDestination Compartment	

Table 7-42 USE

Table 7-43 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE+	CreateBackupDestination	none
BACKUP_DESTINATION_CREA TE	DeleteBackupDestination	
BACKUP_DESTINATION_DELE		

db-nodes

Review the list of permissions and API operations for db-nodes resource-type.

Table 7-44 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
DB_NODE_INSPECT	GetDbNode	none
DB_NODE_QUERY		

Table 7-45 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	No extra	none

Table 7-46 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ +	UpdateDbNode	none
DB_NODE_UPDATE		

Table 7-47 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE+	DbNodeAction	none
DB_NODE_POWER_ACTIONS		



db-homes

Review the list of permissions and API operations for db-homes resource-type.

Table 7-48 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
DB_HOME_INSPECT	ListDBHome	none
	GetDBHome	
	ListDbHomePatches	
	ListDbHomePatchHistoryEnt ries	
	GetDbHomePatch	
	GetDbHomePatchHistoryEntr	
	У	

Table 7-49 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	No extra	none

Table 7-50 USE

Permissions	APIs Fully Covered	APIs Partially Covered
DB_HOME_UPDATE	UpdateDBHome	none

Table 7-51 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
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.



databases

Review the list of permissions and API operations for databases resource-type.

Table 7-52 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
DATABASE_INSPECT	ListDatabases	ListDatabaseUpgradeHist
	GetDatabase oryEntries	oryEntries
	getDatabaseUpgradeHisto ryEntry	UpgradeDatabase
	ListDataGuardAssociatio ns	
	GetDataGuardAssociation	

Table 7-53 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	No extra	none

Table 7-54 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ +	UpdateDatabase	If enabling automatic backups, also needs manage backups.
DATABASE_UPDATE DB HOME UPDATE	SwitchoverDataGuardAsso ciation	
	FailoverDataGuardAssoci ation	
	ReinstateDataGuardAssoc iation	



Permissions	APIs Fully Covered	APIs Partially Covered
USE +	No extra	CreateDbHome, (also needs use vmclusters and manage db-
DATABASE_CREATE		homes). If automatic backups are
DATABASE_DELETE		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

Table 7-55 MANAGE

backups

Review the list of permissions and API operations for backups resource-type.

Table 7-56 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
DB_BACKUP_INSPECT	GetBackup	none
	ListBackups	

Table 7-57 READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT+	none	RestoreDatabase (also needs
DB_BACKUP_CONTENT_READ		use databases)

Table 7-58 USE

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

Table 7-59 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE+	no extra	none
DB_BACKUP_CREATE		
DB_BACKUP_DELETE		

database-software-image

Review the list of permissions and API operations for ${\tt database-software-image}$ resource-type.



Permissions	APIs Fully Covered	APIs Partially Covered
DB_SOFTWARE_IMG_INSPECT	ListDatabaseSoftwareIma ges	none
	GetDatabaseSoftwareImag e	

Table 7-60 INSPECT

Table 7-61 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	No extra	none

Table 7-62 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + DB SOFTWARE IMG UPDATE	UpdateDatabaseSoftwareI mage	none
	ChangeDatabaseSoftwareI mageCompartment	

Table 7-63 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	CreateDatabaseSoftwareI	none
DB_SOFTWARE_IMG_CREATE +	mage	
DB_SOFTWARE_IMG_DELETE	DeleteDatabaseSoftwareI mage	

autonomous-databases

Review the list of permissions and API operations for autonomous-databases resource-type.

Table 7-64 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
AUTONOMOUS_DATABASE_INS PECT	GetAutonomousDatabase, ListAutonomousDatabases	no extra



Table 7-65 READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + AUTONOMOUS_DATABASE_CONTE NT_READ	no extra	CreateAutonomousDatabaseB ackup (also needs manage autonomous-backups)

Table 7-66 USE

Permissions	APIs Fully Covered	APIs Partially Covered
<pre>READ + AUTONOMOUS_DATABASE_CONTE NT_WRITE +</pre>	UpdateAutonomousDatabase	RestoreAutonomousDatabase (also needs read autonomous- backups)
AUTONOMOUS_DATABASE_UPDAT E		ChangeAutonomousDatabaseC ompartment (also needs read autonomous-backups)

Table 7-67 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE+	CreateAutonomousDatabase	none
AUTONOMOUS_DATABASE_CREAT E		
AUTONOMOUS_DATABASE_DELET E		

autonomous-backups

Review the list of permissions and API operations for autonomous-backups resource-type.

Table 7-6	8 INS	SPECT
-----------	-------	-------

Permissions	APIs Fully Covered	APIs Partially Covered
AUTONOMOUS_DB_BACKUP_INSP ECT	ListAutonomousDatabaseBac kups, GetAutonomousDatabaseBack up	none

Table 7-69 READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + AUTONOMOUS_DB_BACKUP_CONT ENT_READ	no extra	RestoreAutonomousDatabase (also needs use autonomous- databases)
		ChangeAutonomousDatabaseC ompartment (also needs use autonomous-databases)



Table 7-70 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + no extra	no extra	none

Table 7-71 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE + AUTONOMOUS_DB_BACKUP_CR EATE	DeleteAutonomousDatabas eBackup	CreateAutonomousDatabas eBackup (also needs read autonomous-databases)
AUTONOMOUS_DB_BACKUP_DE LETE		

autonomous-container-databases

Review the list of permissions and API operations for autonomous-containerdatabases resource-type.

Table 7-72 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
AUTONOMOUS_CONTAINER_DA TABASE_INSPECT	ListAutonomousContainer Databases, GetAutonomousContainerD atabase	none

Table 7-73 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	No extra	none

Table 7-74 USE

Permissions	APIs Fully Covered	APIs Partially Covered
AUTONOMOUS_CONTAINER_DA TABASE_UPDATE	UpdateAutonomousContain erDatabase	e (also needs manage
	ChangeAutonomousContain erDatabaseCompartment	autonomous-databases)



Permissions	APIs Fully Covered	APIs Partially Covered
USE+	No extra	CreateAutonomousContainer
AUTONOMOUS CONTAINER DATA		Database,
BASE CREATE		TerminateAutonomousContai
AUTONOMOUS_CONTAINER_DATA BASE_DELETE		nerDatabase (both also need use autonomous-VmCluster)

Table 7-75 MANAGE

autonomous-vmclusters

Review the list of permissions and API operations for autonomous-vmclusters resource-type.

Table 7-76 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
AUTONOMOUS_VM_CLUSTER_INS	ListAutonomousVmClusters	ChangeAutonomousVmCluster
PECT	GetAutonomousVmCluster	Compartment

Table 7-77 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	No extra	none

Table 7-78 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ+	ChangeAutonomousVmCluster	UpdateAutonomousVmCluster
AUTONOMOUS_VM_CLUSTER_UPD ATE	Compartment	CreateAutonomousContainer Database
		TerminateAutonomousContai nerDatabase

Table 7-79 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE+	DeleteAutonomousVmCluster	CreateAutonomousVmCluster
AUTONOMOUS_VM_CLUSTER_CRE ATE +		
AUTONOMOUS_VM_CLUSTER_DEL ETE		



$autonomous Container {\tt Database} {\tt Dataguard} {\tt Associations}$

Review the list of permissions and API operations for

autonomousContainerDatabaseDataguardAssociations resource-type.

	Table	7-80	INSPECT
--	-------	------	---------

Permissions	APIs Fully Covered	APIs Partially Covered
AUTONOMOUS_VM_CLUSTER_I NSPECT	GetAutonomousContainerD atabase	CreateAutonomousContain erDatabase
AUTONOMOUS_CONTAINER_DA TABASE_INSPECT	ListAutonomousContainer DatabaseDataguardAssoci ations	FailoverAutonomousConta inerDatabaseDataguardAs sociation
	GetAutonomousContainerD atabaseDataguardAssocia tion	SwitchoverAutonomousCon tainerDatabaseDataguard Association
		ReinstateAutonomousCont ainerDatabaseDataguardA ssociation

Table 7-81 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	No extra	none

Table 7-82 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + AUTONOMOUS VM CLUSTER U	none	CreateAutonomousContain erDatabase
PDATE + AUTONOMOUS CONTAINER DA		deleteAutonomouContaine rDatabase
TABASE_UPDATE	Fai	FailoverAutonomousConta inerDatabaseDataguardAs sociation
		SwitchoverAutonomousCon tainerDatabaseDataguard Association
		ReinstateAutonomousCont ainerDatabaseDataguardA ssociation



Permissions	APIs Fully Covered	APIs Partially Covered
USE +	none	CreateAutonomousContainer Database
AUTONOMOUS_CONTAINER_DATA BASE_CREATE +		deleteAutonomouContainerD
AUTONOMOUS_CONTAINER_DATA BASE_DELETE		atabase

Table 7-83 MANAGE

$\label{eq:autonomous} Autonomous {\sf Database} {\sf Dataguard} {\sf Association}$

Review the list of permissions and API operations for AutonomousDatabaseDataguardAssociation resource-type.

Table 7-84 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
AUTONOMOUS_DATABASE_INSPE CT	GetAutonomousDatabase	none

Table 7-85 READ

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

Table 7-86 USE

Permissions	APIs Fully Covered APIs Partially Covered	
READ+	no extra	none
no extra		

Table 7-87 MANAGE

Permissions	APIs Fully Covered APIs Partially Covere	
USE+	no extra	none
no extra		

autonomous-virtual-machine

Review the list of permissions and API operations for autonomous-virtual-machine resource-type.



Permissions	APIs Fully Covered	APIs Partially Covered
AUTONOMOUS_VIRTUAL_MACH INE_INSPECT	GetAutonomousVirtualMac hine	none
	ListAutonomousVirtualMa chines	

Table 7-88 INSPECT

key-stores

Review the list of permissions and API operations for key-store resource-type.

Table 7-89 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
KEY_STORE_INPSECT	GetKeyStore	ChangeKeyStoreCompartme
AUTONOMOUS_CONTAINER_DA	GetAutonomousContainerD	nt
TABASE_INSPECT	atabase	RotateAutonomousContain
AUTONOMOUS_DATABASE_INS	GetAutonomousDatabase	erDatabaseKey
PECT	GetAutonomousDatabaseBa	
AUTONOMOUS_DB_BACKUP_IN	ckup	
SPECT		

Table 7-90 READ

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	no extra

Table 7-91 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + KEY_STORE_UPDATE +	UpdateKeyStore	ChangeKeyStoreCompartme nt
AUTONOMOUS_VM_CLUSTER_U PDATE +	none none none	CreateAutonomousContain erDatabase
AUTONOMOUS_CONTAINER_DA TABASE_UPDATE	RotateAutonomousDatabas eKey	RotateAutonomousContain erDatabaseKey
AUTONOMOUS_DATABASE_UPD ATE		none



Permissions	APIs Fully Covered	APIs Partially Covered
USE + KEY_STORE_CREATE +	CreateKeyStore	none
KEY_STORE_DELETE +	DeleteKeyStore	none
AUTONOMOUS_CONTAINER_DATA BASE_CREATE	CreateAutonomousContainer Database	none

Table 7-92 MANAGE

pluggable-databases (PDBs)

Review the list of permissions and API operations for pluggable-databases resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	PLUGGABLE_DATABASE_ INSPECT	ListPluggableDataba ses	UpdatePluggableData base
		GetPluggableDatabas e	StartPluggableDatab ase
			StopPluggableDataba se
			LocalClonePluggable Database
			RemoteClonePluggabl eDatabase
			RefreshPluggableDat abase
			ConvertRefreshableP luggableDatabase
	DATABASE_INSPECT	no extra	CreatePluggableData base
			DeletePluggableData base
			LocalClonePluggable Database
			RemoteClonePluggabl eDatabase

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
read	INSPECT + PLUGGABLE_DATABASE_ CONTENT_READ	no extra	CreatePluggableData base (Additional permissions are required if auto-backups are enabled on the CDB and includes this PDB.)
			UpdatePluggableData base (Additional permissions are required if auto-backups are enabled on the CDB and includes this PDB.)
			LocalClonePluggable Database
			RemoteClonePluggabl eDatabase
use	READ + PLUGGABLE DATABASE	no extra	LocalClonePluggable Database
	CONTENT_WRITE		RemoteClonePluggabl eDatabase
	PLUGGABLE_DATABASE_ UPDATE	no extra	UpdatePluggableData base
			StartPluggableDatab ase
			StopPluggableDataba se
			LocalClonePluggable Database
			RemoteClonePluggabl eDatabase
			RefreshPluggableDat abase
			ConvertRefreshableP luggableDatabase
	DATABASE_UPDATE	no extra	CreatePluggableData base
			DeletePluggableData base
			LocalClonePluggable Database
			RemoteClonePluggabl eDatabase
manage	USE + Pluggable database	no extra	CreatePluggableData base
	CREATE		LocalClonePluggable Database
			RemoteClonePluggabl eDatabase



Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
	PLUGGABLE_DATABASE_ DELETE	_ no extra	DeletePluggableData base

dbServers

Review the list of permissions and API operations for dbServers resource-type.

Table 7-93 INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered	
EXADATA_INFRASTRUCTURE_IN SPECT	none	GetDbServer ListDbServers	

Table 7-94 READ

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	none	none

Table 7-95 USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ +	none	AddVirtualMachineToVmClus
VM_CLUSTER_UPDATE EXADATA_INFRASTRUCTURE_UP DATE		ter, RemoveVirtualMachineFromV mCluster

Table 7-96 MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	none	none

dbnode-console-connection

Review the list of permissions and API operations for <code>dbnode-console-connection</code> resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DBNODE_CONSOLE_CON ECTION_INSPECT	N GetDbNodeConsoleCon nection	none
		ListDbNodeConsoleCo nnections	
read	no extra	no extra	none

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
use	READ + DBNODE_CONSOLE_CONN ECTION_UPDATE PLUGGABLE_DATABASE_	UpdateDbNodeConsole Connection	none
manage	UPDATE USE + DBNODE_CONSOLE_CONN ECTION_CREATE DBNODE_CONSOLE_CONN ECTION_DELETE	DeleteDbNodeConsole	none

dbnode-console-history

Review the list of permissions and API operations for <code>dbnode-console-history</code> resource-type.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DBNODE_CONSOLE_HI STORY_INSPECT	GetDbNodeConsoleH istory ListDbNodeConsole Histories	none
read	INSPECT +	DBNODE_CONSOLE_HI STORY_CONTENT_REA D	none
use	READ + DBNODE_CONSOLE_HI STORY_UPDATE PLUGGABLE_DATABAS E_UPDATE	UpdateDbNodeConso leHistory	none
manage	USE + DBNODE_CONSOLE_HI STORY_CREATE DBNODE_CONSOLE_HI STORY_DELETE	CreateDbNodeConso leHistory DeleteDbNodeConso leHistory	none

oneoffPatch

Review the list of permissions and API operations for oneoffPatch resource-type.



Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	ONEOFF_PATCH_INSPEC	DownloadOneoffPatch	CreateOneoffPatch
	Т	GetOneoffPatch	DeleteOneoffPatch
		ListOneoffPatches	UpdateOneoffPatch
			ChangeOneoffPatchCo mpartment
read	INSPECT + no extra	DownloadOneoffPatch	none
use	READ +	none	UpdateOneoffPatch
	ONEOFF_PATCH_UPDATE		ChangeOneoffPatchCo mpartment
manage	USE +	none	CreateOneoffPatch
	ONEOFF_PATCH_CREATE		DeleteOneoffPatch
	ONEOFF_PATCH_DELETE		

Related Topics

OneoffPatch Reference

Permissions Required for Each API Operation

Review the list of API operations for Exadata Database Service on Cloud@Customer resources in a logical order, grouped by resource type.

For information about permissions, see Permissions.

Table 7-97	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
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



API Operation	Permissions Required to Use the Operation
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
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

Table 7-97	(Cont.)	Database	API	Operations
------------	---------	----------	-----	------------



API Operation	Permissions Required to Use the Operation
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
getDatabaseUpgradeHistoryEntry	DATABASE_INSPECT
ListDbVersions	(no permissions required; available to anyone)
GetBackup	DB_BACKUP_INSPECT
ListBackups	DB_BACKUP_INSPECT

Table 7-97 (Cont.) Database API Operations



API Operation	Permissions Required to Use the Operation
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
ListAutonomousVmClusters	AUTONOMOUS_VM_CLUSTER_INSPECT
GetAutonomousVmCluster	AUTONOMOUS_VM_CLUSTER_INSPECT
CreateAutonomousVmCluster	AUTONOMOUS_VM_CLUSTER_CREATE and EXADATA_INFRASTRUCTURE_INSPECT and EXADATA_INFRASTRUCTURE_UPDATE
UpdateAutonomousVmCluster	AUTONOMOUS_VM_CLUSTER_UPDATE and EXADATA_INFRASTRUCTURE_INSPECT and EXADATA_INFRASTRUCTURE_UPDATE
ChangeAutonomousVmClusterCompartment	AUTONOMOUS_VM_CLUSTER_INSPECT and AUTONOMOUS_VM_CLUSTER_UPDATE
DeleteAutonomousVmCluster	AUTONOMOUS_VM_CLUSTER_DELETE
ListAutonomousContainerDatabases	AUTONOMOUS_CONTAINER_DATABASE_INSPECT
GetAutonomousContainerDatabase	AUTONOMOUS_CONTAINER_DATABASE_INSPECT
CreateAutonomousContainerDatabase	EXADATA_INFRASTRUCTURE_UPDATE and AUTONOMOUS_VM_CLUSTER_UPDATE and AUTONOMOUS_CONTAINER_DATABASE_CREATE
TerminateAutonomousContainerDatabase	EXADATA_INFRASTRUCTURE_UPDATE and AUTONOMOUS_VM_CLUSTER_UPDATE and AUTONOMOUS_CONTAINER_DATABASE_DELETE
UpdateAutonomousContainerDatabase	AUTONOMOUS_CONTAINER_DATABASE_UPDATE
ChangeAutonomousContainerDatabaseCompar tment	AUTONOMOUS_CONTAINER_DATABASE_INSPECT and AUTONOMOUS_CONTAINER_DATABASE_UPDATE
RotateAutonomousContainerDatabaseEncryp tionKey	AUTONOMOUS_CONTAINER_DATABASE_UPDATE and AUTONOMOUS_CONTAINER_DATABASE_INSPECT
GetAutonomousDatabase	AUTONOMOUS_DATABASE_INSPECT
ListAutonomousDatabases	AUTONOMOUS_DATABASE_INSPECT
CreateAutonomousDatabase	AUTONOMOUS_DATABASE_CREATE
UpdateAutonomousDatabase	AUTONOMOUS_DATABASE_UPDATE
ChangeAutonomousDatabaseCompartment	AUTONOMOUS_DATABASE_UPDATE and AUTONOMOUS_DB_BACKUP_INSPECT and AUTONOMOUS_DB_BACKUP_CONTENT_READ and AUTONOMOUS_DATABASE_CONTENT_WRITE
DeleteAutonomousDatabase	AUTONOMOUS DATABASE DELETE

Table 7-97	(Cont.)	Database	API O	perations
------------	---------	----------	-------	-----------



API Operation	Permissions Required to Use the Operation
StartAutonomousDatabase	AUTONOMOUS_DATABASE_UPDATE
StopAutonomousDatabase	AUTONOMOUS_DATABASE_UPDATE
RestartAutonomousDatabase	AUTONOMOUS_DATABASE_UPDATE
RestoreAutonomousDatabase	AUTONOMOUS_DB_BACKUP_CONTENT_READ and AUTONOMOUS_DATABASE_CONTENT_WRITE
RotateAutonomousDatabaseEncryptionKey	AUTONOMOUS_DATABASE_UPDATE
CreateAutonomousDatabaseBackup	AUTONOMOUS_DB_BACKUP_CREATE and AUTONOMOUS_DATABASE_CONTENT_READ
ListAutonomousDatabaseBackups	AUTONOMOUS_DB_BACKUP_DELETE
GetAutonomousDatabaseBackup	AUTONOMOUS_DB_BACKUP_DELETE
ListAutonomousContainerDatabaseDataguar dAssociations	AUTONOMOUS_CONTAINER_DATABASE_INSPECT
GetAutonomousContainerDatabaseDataguard Association	AUTONOMOUS_CONTAINER_DATABASE_INSPECT
FailoverAutonomousContainerDatabaseData guardAssociation	AUTONOMOUS_CONTAINER_DATABASE_INSPECT and AUTONOMOUS_CONTAINER_DATABASE_UPDATE
SwitchoverAutonomousContainerDatabaseDa taguardAssociation	AUTONOMOUS_CONTAINER_DATABASE_INSPECT and AUTONOMOUS_CONTAINER_DATABASE_UPDATE
ReinstateAutonomousContainerDatabaseDat aguardAssociation	AUTONOMOUS_CONTAINER_DATABASE_INSPECT and AUTONOMOUS_CONTAINER_DATABASE_UPDATE
UpdateAutonomousContainerDatabaseDatagu ardAssociation	AUTONOMOUS_CONTAINER_DATABASE_INSPECT and AUTONOMOUS_CONTAINER_DATABASE_UPDATE
ListAutonomousDatabaseDataguardAssociat ions	AUTONOMOUS_CONTAINER_DATABASE_INSPECT
GetAutonomousDatabaseDataguardAssociati on	AUTONOMOUS_CONTAINER_DATABASE_INSPECT
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

Table 7-97 (Cont.) Database API Operations



API Operation	Permissions Required to Use the Operation		
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		
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		

Table 7-97 (Cont.) Database API Operations



API Operation Permissions Required to Use the C		
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	
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	

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.
- 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.

"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.

Security Guide for Exadata Database Service on Cloud@Customer Systems

This guide describes security for an Oracle Exadata Cloud@Customer System. It includes information about the best practices for securing the Oracle Exadata 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
- Guest VM Default Security Settings
- Guest VM Default Processes



Guest VM Network Security

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 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:
 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

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 <code>exawatch</code> and <code>dbmsvc</code>. These users should not be altered. These users cannot login to the system as all are set to <code>/bin/false</code>.

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

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 <code>UpdateVmCluster</code> 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 <code>UpdateVmCluster</code> method, the key is only added to the <code>authorized_keys</code> 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 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.
 - It is recommended to leave this setting to permit the subset of cloud automation that uses this access path (for example, customer VM OS patching) to function. Setting PermitRootLogin to no will disable this subset of cloud automation functionality.
- 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 the hand maxloging 10 and it is the recommended secure value.
 - 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, dbcsagent-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.
 - 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

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
Oracle Clusterware	clib0/clre0	1525	Oracle TNS listener
running on each cluster node communicates		3260	Synology DSM iSCSI
through these interfaces.		5054	Oracle Grid Interproces Communication
		7879	Jetty Management Server
		Dynamic Port: 9000-65500	System Monitor service (osysmond)
		Ports are controlled by the configured ephemeral range in the operating system and are dynamic.	
		Dynamic Port: 9000-65500	Cluster Logger service (ologgerd)
		Ports are controlled by the configured ephemeral range in the operating system and are dynamic.	

Table 7-98 Default Port Matrix for Guest VM Services



Type of interface	Name of interface	Port	Process running
	clib1/clre1	5054	Oracle Grid Interprocess communication
		7879	Jetty Management Server
Cluster nodes use these	stib0/stre0	7060	dbcs-admin
interfaces to access storage cells (ASM		7070	dbcs-agent
disks).	stib1/stre1	7060	dbcs-admin
However, the IP/ports 7060/7070 attached to the storage interfaces are used to access DBCS agent from the Control Plane server.		7070	dbcs-agent
Control Plane server to domU	eth0	22	sshd
Loopback	lo	22	sshd
		2016	Oracle Grid Infrastructure
		6100	Oracle Notification Service (ONS), part of Oracle Grid Infrastructure
		7879	Jetty Management Server
		Dynamic Port 9000-65500	Oracle Trace File Analyzer

Table 7-98 (Cont.) Default Port Matrix for Guest VM Services

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

Additional Procedures for Updating Security Posture

- Customer Responsibilities
- Enabling Additional Security Capabilities

Customer Responsibilities

	Oracle Cloud Platform		Customer/Tenant Instances	
Monitoring	Oracle Cloud Ops	Customer	Oracle Cloud Ops	Customer
	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

Table 7-99 Resposibilities



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

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.
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 7-100 host_access_control password-aging

Table 7-101 host_access_control pam-auth

Options	Description	
-h,help	Shows this help message and exits.	



Options	Description
-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)
-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 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 <i>n</i> 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

Table 7-101 (Cont.)	host_acce	ss_control	pam-auth
---------------------	-----------	------------	----------



Table 7-101	(Cont.) host	_access_	_control	pam-auth
--------------------	--------------	----------	----------	----------

Options	Description
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

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 7-102	User Accounts
-------------	---------------

User Name and Password	User Type	Component
opc - key-based login only	Operating system user	Oracle Exadata Database Servers



User Name and Password	User Type	Component
exawatch (release 19.1.0 and later) - no logon privileges	Operating system user	Oracle Exadata Database Servers
		Oracle Exadata Storage Servers
SYS/Welcome\$	Oracle Database user	Oracle Exadata Database Servers
SYSTEM/Welcome\$	Oracle Database user	Oracle Exadata Database Servers
MSUser	ILOM user	Database server ILOMs
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		Oracle Exadata Storage Server ILOMs
password.		
Do not modify this account. This account is to be used by MS only.		

Table 7-102 (Cont.) User Accounts

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 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

Troubleshooting 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 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

Patching Failures on 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 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 Exadata Database Service on Cloud@Customer components.

ORACLE

Determining the Problem

In the Console, you can identify a failed patching operation by viewing the patch history of an 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 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:

- Verify that your DNS configuration is correct so that the relevant virtual machine addresses are resolvable within the VM cluster.
- Refer to the relevant Cloud Tooling logs as instructed in the *Obtaining Further Assistance* section and contact Oracle Support for further assistance.

Oracle Grid Infrastructure Issues

One or more of the following conditions on Oracle Grid Infrastructure can cause patching operations to fail.

Oracle Grid Infrastructure is Down

Oracle Clusterware enables servers to communicate with each other so that they can function as a collective unit. The cluster software program must be up and running on the VM Cluster for patching operations to complete. Occasionally you might need to restart the Oracle Clusterware to resolve a patching failure.



In such cases, verify the status of the Oracle Grid Infrastructure as follows:

```
./crsctl check cluster
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
```

If Oracle Grid Infrastructure is down, then restart by running the following commands:

```
crsctl start cluster -all
```

crsctl check cluster

Oracle Databases Issues

An improper database state can lead to patching failures.

Oracle Database is Down

The database must be active and running on all the active nodes so the patching operations can be completed successfully across the cluster.

Use the following command to check the state of your database, and ensure that any problems that might have put the database in an improper state are resolved:

srvctl status database -d db unique name -verbose

The system returns a message including the database instance status. The instance status must be **Open** for the patching operation to succeed.

If the database is not running, use the following command to start it:

srvctl start database -d db unique name -o open

Obtaining Further Assistance

If you were unable to resolve the problem using the information in this topic, follow the procedures below to collect relevant database and diagnostic information. After you have collected this information, contact Oracle Support.

- Collecting Cloud Tooling Logs Use the relevant log files that could assist Oracle Support for further investigation and resolution of a given issue.
- Collecting Oracle Diagnostics

Related Topics

Oracle Support



Collecting Cloud Tooling Logs

Use the relevant log files that could assist Oracle Support for further investigation and resolution of a given issue.

DBAASCLI Logs

/var/opt/oracle/log/dbaascli

dbaascli.log

Collecting Oracle Diagnostics

To collect the relevant Oracle diagnostic information and logs, run the dbaascli diag collect command.

For more information about the usage of this utility, see DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check.

Related Topics

• DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check

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

/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 <code>rhphelper</code> 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 <code>dbnodeupdate.trc</code> file moves and the Grid Infrastructure stack on the node is shutdown.

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 Guardenabled 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

CPU Offline Scaling Fails



Description: CPU offline scaling fails with the following error:

 ** CPU Scale Update $^{\star\star}An$ 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.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 <code>lsnrctl reload -with ha</code>.
 - 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**.
- 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_name -endpoints endpoints* command. For more information, see srvctl modify listener in the Oracle Real Application Clusters Administration and Deployment Guide.

