

Oracle® Enterprise Manager

Oracle Enterprise Manager for Oracle Exadata Cloud



24ai Release 1 (24.1)

F97226-04

August 2025

ORACLE®

Copyright © 2020, 2025, Oracle and/or its affiliates.

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

Preface

Audience	i
Documentation Accessibility	i
Related Resources	ii
Conventions	ii

Part I Introduction to Monitoring Exadata Cloud Service

Part II Discovery of Exadata Cloud

1 Exadata Cloud Discovery Prerequisites

Configuration Prerequisites	1
Verify Names Resolution	1
Verify Firewall Configuration	1
Create Required Named Credentials	3
Verify ExaCLI Credentials	3
Create IAM Policies to Support Discovery and Monitoring	3
Create Named Credentials for OCI Monitoring	6
Create SSH Based Credentials	8
Create Storage Server Credentials	8
Agent Installation	9
Overview of the Required Agents	9
Install Agents	10
Set Up a Proxy Connection	10
Test the Named Credentials Created for OCI Monitoring	10

2 Discover Exadata Cloud Targets

Discover the Exadata Cloud Target Using Console	1
-------------------------------------------------	---

3 Explore the Exadata Cloud Target Home Pages

Exadata Infrastructure Home Page	1
Target Navigation Tree	4
Exadata VM Cluster Home Page	4
Oracle Autonomous Exadata VM Cluster Home Page	7
Oracle Exadata Database Service on Exascale Infrastructure Home Page	8

Part III Exadata Cloud Lifecycle Management

4 Database Lifecycle Management in Exadata Cloud

Provision a Pluggable Database (PDB)	1
Update and Upgrade PDBs using Fleet Maintenance	3
Relocate a PDB	4
Unplug and Plug a PDB	4
Manage Configuration of PDB	7
Manage Compliance of a PDB	7
Clone a PDB	8
Delete a PDB	9

5 Hybrid Multicloud Management - Exadata Cloud

Setting up Pluggable Database as a Service	1
Create an Empty PDB	2
Create a PDB With Data	3
Unplug and Plug a PDB	4
Dehydrate and Rehydrate a PDB	5
Resize PDB	5

Part IV Administration and Monitoring

6 Administer Exadata Cloud Service Target

Navigate to Individual Targets	1
View Cloud Target Alerts	1
View Metrics of the Individual targets	2
Delete a Target Using the UI	2
Manage Exadata Storage Server Metrics and Alert Settings	2
Accessing Oracle Analytics Server Reports	3

7 Monitor Exadata Cloud Service

Access Exadata Metrics	1
Aggregated Exadata FlashDisk and HardDisk Metric Example	1
Exadata Cell Metric Example	2
Exadata Key Performance Indicators Metrics Examples	3

8 Exadata Warehouse

Setting up Exadata Warehouse and Analyzing Performance Data	2
Simplifying Capacity Planning and Forecasting	2
Exadata Warehouse Sample Queries	3
Configuring Purge Settings for Old Metric Data	7

Part V Troubleshooting

Preface

Oracle Database Exadata Cloud Service provides full Oracle Databases hosted on Oracle Exadata Database Machine.

Exadata Cloud can be one of the following in the context of this book:

- Exadata Cloud is based on Oracle Cloud Infrastructure (OCI). *Exadata Infrastructure* is available in two flavors:
 - *Oracle Exadata Database Service on Dedicated Infrastructure* (ExaDB-D) is offered on Oracle Cloud, using state-of-the-art Oracle-managed data centers.
 - *Oracle Exadata Database Service on Cloud@Customer* (ExaCC) provides Exadata Cloud Service hosted in your data center.
- Exadata Cloud also includes *Oracle Exadata Database Service on Exascale Infrastructure* (ExaDB-XS) which is based on Oracle Cloud Infrastructure (OCI) and provides Exadata Cloud that leverages Exascale storage.
- The discovery of *Exadata Infrastructure* covers both flavors *Oracle Exadata Database Service on Dedicated Infrastructure* and *Oracle Exadata Database Service on Cloud@Customer*. *Oracle Exadata Database Service on Exascale Infrastructure* is discovered as an Exadata VM Cluster.

This book describes how you can use Oracle Enterprise Manager for Exadata Cloud Service to:

- Have a unified view of all the cloud components and their status
- Efficiently manage cloud target resources
- Perform typical administrative tasks on the cloud target and its component targets

Audience

This guide is intended for administrators of *Oracle Enterprise Manager for Exadata Cloud Service* who manage and monitor the resources associated with these services.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers who have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Resources

For more information, see these Oracle resources:

- [Oracle Exadata Database Service on Dedicated Infrastructure](#) (Oracle Database Exadata Cloud Service Deployed on Oracle Cloud Infrastructure)
- [Oracle Exadata Database Service on Cloud@Customer](#)
- [Oracle Exadata Database Service on Exascale Infrastructure](#)

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Part I

Introduction to Monitoring Exadata Cloud Service

Oracle Enterprise Manager offers a complete cloud lifecycle management solution for the various Exadata Cloud Services.

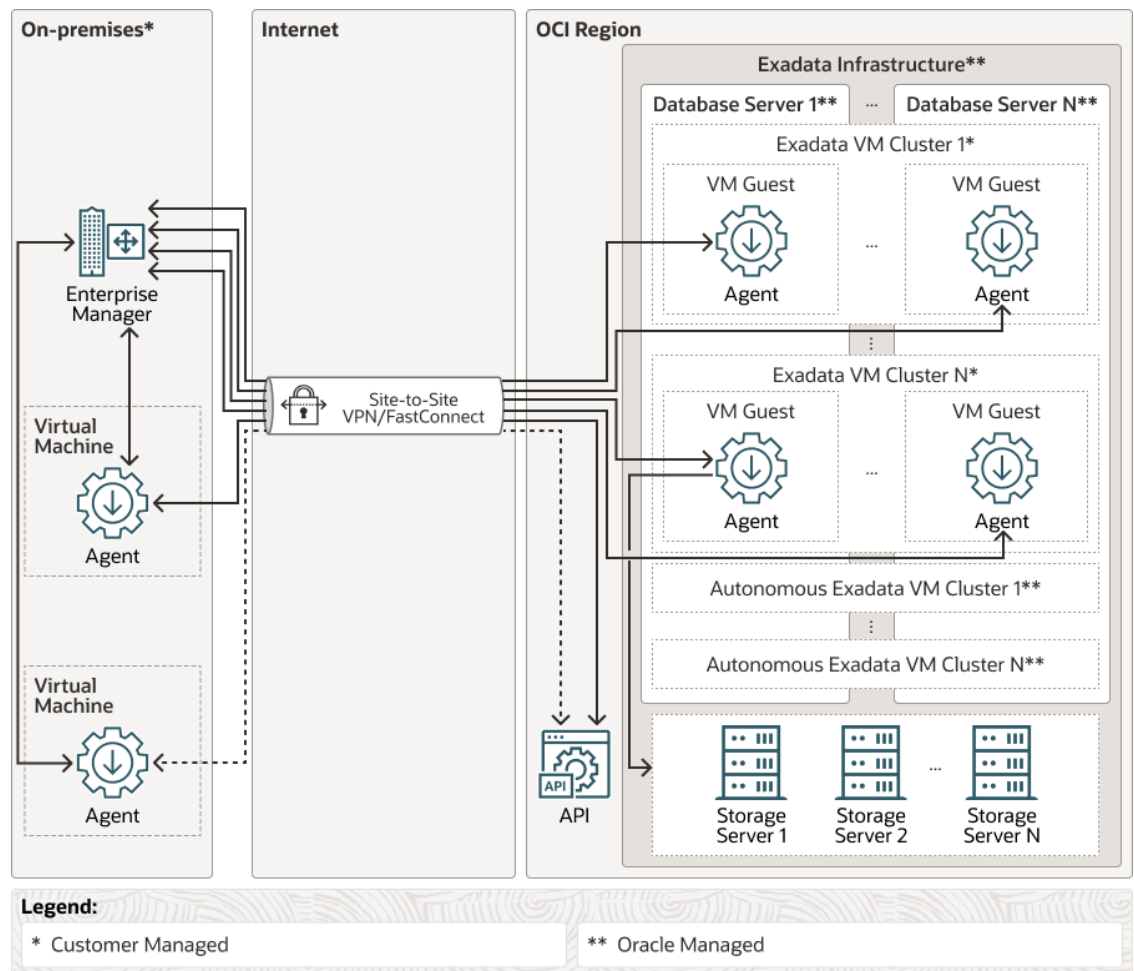
Note

To monitor Oracle Exadata Database Machine using Oracle Enterprise Manager, see *Oracle Exadata Database Machine Getting Started Guide*.

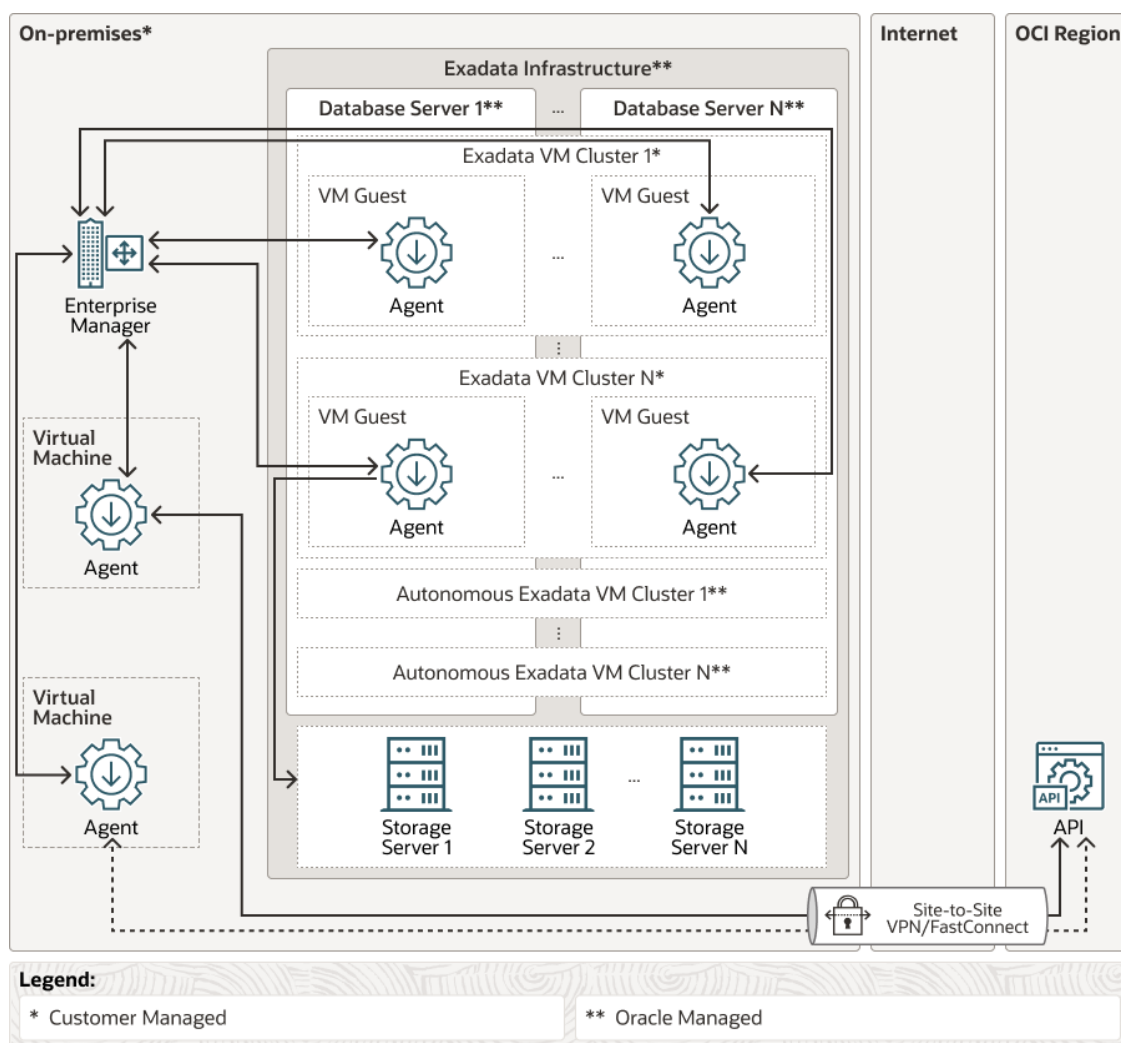
Overview and Architecture

Oracle Enterprise Manager provides an integrated monitoring solution for the various resources in the Exadata Cloud Service. The following diagrams show the various resources of the Exadata Cloud Services and the resource management by *Oracle* or *Customer*.

Oracle Exadata Database Service on Dedicated Infrastructure:



Oracle Exadata Database Service on Cloud@Customer:



As in the on-premises Oracle Exadata Database Machine, an Enterprise Manager Agent is installed on each of the Exadata Cloud virtual machines for each VM cluster. The agent helps in monitoring the resources that reside in that VM Cluster such as host, database, listener, ASM, etc. One of the agents installed on a VM cluster is used to monitor the storage servers through the cloud-provided *ExaCLI* user. In addition to the agents installed in the VM cluster hosts, a standard Enterprise Manager Agent is installed outside of the Exadata Cloud VM clusters. This agent is used for discovery and monitoring the Exadata Cloud OCI resources such as the Exadata Infrastructure and VM Clusters. It is installed in a location that allows communication with the OCI REST API endpoints as well as the Enterprise Manager OMS servers. This same agent can be used to discover and monitor multiple Exadata Cloud targets. For more details regarding the required agents, see [Overview of the Required Agents](#).

Oracle Enterprise Manager discovers Exadata Cloud Service as a single target and automatically identifies and organizes all its dependent components. Using Oracle Enterprise Manager you can then:

- Monitor and manage all Exadata, Exadata Cloud systems along with any other targets, from a single interface
- Visualize storage and compute data
- View performance metrics of your Exadata components

Features

Oracle Enterprise Manager for Oracle Exadata Cloud includes the following features:

Enterprise Manager target for Exadata Cloud

- 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

- Adds Exadata Storage Server and Exadata Storage Grid targets.
- Offers visualization of storage and compute performance for Oracle Exadata Cloud
- Enables use of the same Maximum Availability Architecture (MAA) Key performance Indicators (KPI) developed for Oracle Exadata Database Machine

Provides a single pane of glass for Oracle Exadata Database Machine and Oracle Exadata Cloud

- Monitors and manages Oracle Exadata Database Machine and Oracle Exadata Cloud systems through a common Exadata target menu
- Exadata Cloud Target provides consistent Enterprise Manager experience across Exadata Cloud Service by means of a single target type.

Visualization

- Visualize the database and related targets associated with each Exadata Cloud Service.

Supported Operating Systems

The following operating systems (where OMS and agent is installed on) are supported by the Oracle Exadata plug-in 13.5.1:

- Management Server plug-in (all OMS-certified platforms):
 - IBM AIX on POWER Systems (64-bit)
 - HP-UX Itanium
 - Linux x86 and x86-64
 - Microsoft Windows x64 (64-bit)
 - Oracle Solaris on SPARC (64-bit)
 - Oracle Solaris on x86-64 (64-bit)
- Agent Plug-ins for Exadata and Supercluster
 - Exadata Plug-in + SI plug-in + VI Plug-in for Exadata
 - Exadata Plug-in + SI Plug-in for SSC
 - Linux x86-64
 - Oracle Solaris on x86-64 (64-bit)
 - Oracle Solaris on SPARC (64-bit)

Part II

Discovery of Exadata Cloud

Use the following steps to perform prerequisites and then discover *Oracle Exadata Database Service on Dedicated Infrastructure*, *Oracle Exadata Database Service on Cloud@Customer*, or *Oracle Exadata Database Service on Exascale Infrastructure* target.

Starting with *Oracle Enterprise Manager 24ai Release 1 Update 2 (24.1.0.2)*, fresh discovery of Gen1 Exadata Cloud targets is not supported in Oracle Enterprise Manager. However, you can continue to monitor the already discovered Gen1 Exadata Cloud targets.

Topics:

- [Exadata Cloud Discovery Prerequisites](#)
- [Discover Exadata Cloud Targets](#)
- [Explore the Exadata Cloud Target Home Pages](#)

1

Exadata Cloud Discovery Prerequisites

Complete the following prerequisite tasks before you begin the discovery process for *Oracle Exadata Database Service on Dedicated Infrastructure*, *Oracle Exadata Database Service on Cloud@Customer*, or *Oracle Exadata Database Service on Exascale Infrastructure*.

Topics:

- [Configuration Prerequisites](#)
 - [Verify Names Resolution](#)
 - [Verify Firewall Configuration](#)
- [Create Required Named Credentials](#)
 - [Verify ExaCLI Credentials](#)
 - [Create IAM Policies to Support Discovery and Monitoring](#)
 - [Create Named Credentials for OCI Monitoring](#)
 - [Create SSH Based Credentials](#)
 - [Create Storage Server Credentials](#)
- [Agent Installation](#)
 - [Overview of the Required Agents](#)
 - [Install Agents](#)
 - [Set Up a Proxy Connection](#)
- [Test the Named Credentials Created for OCI Monitoring](#)

Configuration Prerequisites

Before starting the discovery procedure, make sure the following prerequisites are met:

- [Verify Names Resolution](#)
- [Verify Firewall Configuration](#)

Verify Names Resolution

The Oracle Enterprise Manager OMS servers require direct network access to each of the VM Guests. If the names of the VM Guests are not registered in the OMS nodes' DNS, then they must be manually entered in the `/etc/hosts` file for each OMS.

To manage the Exadata Cloud Service components from Oracle Enterprise Manager, the agent on the VM guest hosts must be able to resolve the upload URL utilized by Oracle Enterprise Manager.

Verify Firewall Configuration

To verify the firewall configuration:

1. Open Database Ports

The database listener ports must be opened for the Oracle Enterprise Manager OMS servers. Note that Exadata Cloud Service databases will use SCAN listeners; so, ports will need to be opened for the base VM Guest, the VM Guest virtual IP, and scan listeners addresses.

For example, if an Exadata Cloud Service quarter rack has been configured with two VM Guests - `exadbnod1.example.com` and `exadbnod2.example.com` - and the listeners are using port 1521, then port 1521 will have to be opened to the Oracle Enterprise Manager Server for the following addresses:

- The VM Guest hostnames - `exadbnod1.example.com` and `exadbnod2.example.com`
- The virtual IPs for each VM Guest - `exadbnod1-vip.example.com` and `exadbnod1-vip.example.com`
- The scan listener hostname - `scan-exadatadb`

2. Open Enterprise Manager Upload Port

The Oracle Enterprise Manager agents require access to the Oracle Enterprise Manager Servers upload service, normally configured on port 4889 for HTTP uploads and 4900 for HTTPS. To verify the ports assigned, run the following command on the OMS server command line.

```
$ emctl status oms -details
```

These ports must be opened for each of the VM Guests.

3. Open Agent Ports

The OMS servers must be able to connect to the Oracle Enterprise Manager agent HTTP/HTTPS port on each VM Guest. The agent port defaults to 3872. If port 3872 is not available, the next available port starting from port 1830 is used.

To identify the port used:

- Run the following command on the VM Guest command line:

```
$ emctl status agent
```

- Alternatively, you can look for the value of the `EMD_URL` property in the `emd.properties` file the following directory:

```
<AGENT_HOME>/agent_inst/sysman/config
```

4. Verify SSL Certificate

Verify that the SSL certificate sent by OCI Service and the one received at the Oracle Enterprise Manager agent side are from trusted source, that is Oracle Corporation / DigiCert. Run the following command on the agent host:

```
openssl s_client -connect database.us-ashburn-1.oraclecloud.com:443
```

The example output of the above command where you can check the authenticity of the certificate provider:

```
CONNECTED(00000003)depth=2 C = US, O = DigiCert Inc, OU = www.digicert.com, CN = DigiCert Global Root
G2verify return:1depth=1 C = US, O = DigiCert Inc, CN = DigiCert Global G2 TLS RSA SHA256 2020 CA1verify
return:1depth=0 C = US, ST = California, L = Redwood City, O = Oracle Corporation, CN =
*.us-ashburn-1.oc-test.comverify return:1---Certificate chain 0 s:/C=US/ST=California/L=Redwood City/O=Oracle
Corporation/CN=*.us-ashburn-1.oc-test.com i:/C=US/O=DigiCert Inc/CN=DigiCert Global G2 TLS RSA SHA256 2020
CA1 1 s:/C=US/O=DigiCert Inc/CN=DigiCert Global G2 TLS RSA SHA256 2020 CA1 i:/C=US/O=DigiCert Inc/OU
```

If the certificate is not from a trusted or known source, then see [Troubleshooting](#).

The hosts where the agent selected for discovery and monitoring of the Exadata Cloud Services resides must use `https` on port 443 to communicate with the OCI APIs.

Ensure that `tcp` is open on the database listener port on the Exadata Cloud VM Guest hosts for the OMS to communicate with the databases.

Table 1-1 Firewall Ports

Component	Notes
OMS	Upload http/https port - usually 4903
Agent	The OMS servers will need to be able to connect to the Oracle Enterprise Manager Agent HTTP/HTTPS port on each VM Guest. The agent port defaults to 3872. If port 3872 is not available, the next available port starting from port 1830 is used.

Create Required Named Credentials

The user for whom the named credentials are created, must have **Monitoring** privileges.

- [Verify ExaCLI Credentials](#)
- [Create IAM Policies to Support Discovery and Monitoring](#)
- [Create Named Credentials for OCI Monitoring](#)
- [Create SSH Based Credentials](#)
- [Create Storage Server Credentials](#)

Verify ExaCLI Credentials

Run the following command on *VM Guest* to verify the `exacli` credentials:

```
exacli --xml -c <user>@<ip_address> -e 'list cell attributes all'
```

Collect the storage server *IP address* from `/etc/oracle/cell/network-config/cellip.ora` on the *VM Guest*.

Create IAM Policies to Support Discovery and Monitoring

In Oracle Cloud Infrastructure (OCI), access is granted to resources in an identity domain by creating policies. To successfully discover and monitor Exadata Database Service and Autonomous Database on Oracle Exadata Database Service on Cloud@Customer, Oracle Exadata Database Service on Dedicated Infrastructure, and Oracle Exadata Database Service on Exascale Infrastructure in OCI for Oracle Enterprise Manager, you must manually grant and maintain IAM policies to provide Oracle Enterprise Manager with necessary privileges for Exadata Database Service resources in OCI.

There are two approaches to maintain these policies:

- The simplest approach involves the definition of policies using predefined OCI aggregate resource types, which provide access to resources that are typically used together.

- For organizations with more fine-grained security requirements, a more detailed approach involves the definition of policies created specifically on each individual applicable resource, further supporting the principle of least privilege. There are separate service-specific individual resource types for deployments on Oracle Exadata Database Service on Cloud@Customer, Oracle Exadata Database Service on Dedicated Infrastructure, and Oracle Exadata Database Service on Exascale Infrastructure in OCI.

Each of these approaches can be implemented tenancy-wide or at a compartment level for finer control. Furthermore, the policies can be created for individual users or for a group of users. The examples in this document reference groups. For details on creating groups, see [Managing Groups](#) in the *Oracle Cloud Infrastructure Documentation*.

The following sections provide the required policy statements for each approach including the policies to be granted at aggregate resource level or individual resource level, at the compartment level or tenancy level. In case of granting at the individual resource level, the information is also provided by service type and deployment. Select the approach required by your organization's security requirements and follow the corresponding section.

Note that all the policy statements include the identify domain and are granted at the group level. Substitute the actual identity domain name for *<Domain>* and a suitable group name for *<Group Name>* in the policy statement. If the identity domain is Default, then substitute *<Domain>* with Default. If required, all statements can be modified to grant to an individual user instead of a group. Also, insert the right compartment name for *<Compartment Name>* in the policy statements.

For more information on managing access to OCI resources, see [Managing Access to Resources](#) in *Oracle Cloud Infrastructure Documentation*.

Create Policies for Aggregate Resources

Create and maintain policies at the compartment or tenancy level to support the discovery and monitoring using aggregate resource types. The same aggregate resource types support Oracle Exadata Database Service on Cloud@Customer, Oracle Exadata Database Service on Dedicated Infrastructure, and Oracle Exadata Database Service on Exascale Infrastructure deployments.

- Creating at compartment level:**

Allow group *<Domain>/<Group Name>* to inspect database-family in compartment *<Compartment Name>*
 Allow group *<Domain>/<Group Name>* to inspect autonomous-database-family in compartment *<Compartment Name>*

- Creating at tenancy level:**

Allow group *<Domain>/<Group Name>* to inspect database-family in tenancy
 Allow group *<Domain>/<Group Name>* to inspect autonomous-database-family in tenancy

Create Policies for Individual Resources

Create and maintain policies at the compartment or tenancy level to support the discovery and monitoring using individual resource types. There are separate individual resource types for Oracle Exadata Database Service on Cloud@Customer and Oracle Exadata Database Service on Dedicated Infrastructure deployments.

Oracle Exadata Database Service on Cloud@Customer

- **Creating at compartment level:**

Allow group *<Domain>/<Group Name>* to inspect exadata-infrastructures in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect vmclusters in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect db-nodes in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect db-homes in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect databases in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect pluggable-databases in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect autonomous-vmclusters in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect autonomous-container-databases in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect autonomous-databases in compartment *<Compartment Name>*

- **Creating at tenancy level:**

Allow group *<Domain>/<Group Name>* to inspect exadata-infrastructures in tenancy

Allow group *<Domain>/<Group Name>* to inspect vmclusters in tenancy

Allow group *<Domain>/<Group Name>* to inspect db-nodes in tenancy

Allow group *<Domain>/<Group Name>* to inspect db-homes in tenancy

Allow group *<Domain>/<Group Name>* to inspect databases in tenancy

Allow group *<Domain>/<Group Name>* to inspect pluggable-databases in tenancy

Allow group *<Domain>/<Group Name>* to inspect autonomous-vmclusters in tenancy

Allow group *<Domain>/<Group Name>* to inspect autonomous-container-databases in tenancy

Allow group *<Domain>/<Group Name>* to inspect autonomous-databases in tenancy

Oracle Exadata Database Service on Dedicated Infrastructure

- **Creating at compartment level:**

Allow group *<Domain>/<Group Name>* to inspect cloud-exadata-infrastructures in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect cloud-vmclusters in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect db-nodes in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect db-homes in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect databases in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect pluggable-databases in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect cloud-autonomous-vmclusters in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect autonomous-container-databases in compartment *<Compartment Name>*

Allow group *<Domain>/<Group Name>* to inspect autonomous-databases in compartment *<Compartment Name>*

- **Creating at tenancy level:**

Allow group *<Domain>/<Group Name>* to inspect cloud-exadata-infrastructures in tenancy

Allow group *<Domain>/<Group Name>* to inspect cloud-vmclusters in tenancy

Allow group *<Domain>/<Group Name>* to inspect db-nodes in tenancy

Allow group *<Domain>/<Group Name>* to inspect db-homes in tenancy

Allow group <Domain>/<Group Name> to inspect databases in tenancy
Allow group <Domain>/<Group Name> to inspect pluggable-databases in tenancy
Allow group <Domain>/<Group Name> to inspect cloud-autonomous-vmclusters in tenancy
Allow group <Domain>/<Group Name> to inspect autonomous-container-databases in tenancy
Allow group <Domain>/<Group Name> to inspect autonomous-databases in tenancy

Create Named Credentials for OCI Monitoring


To enable access to the Oracle Cloud resources, Enterprise Manager requires API key credentials associated with an Oracle Cloud Infrastructure IAM user. This IAM user must be assigned the required permissions using policies.

During the discovery flow, the list of available resources to discover as targets will be controlled by the policies and permissions of this IAM user. If the IAM user has limited visibility to the Exadata Cloud resources in the tenancy, the resources available during discovery will be restricted to this list. However, if the IAM user has permissions to view all the resources in the tenancy, then the list of Exadata Cloud resources returned during the discovery flow will include them all.

For details on creating the required policies, see [Create IAM Policies to Support Discovery and Monitoring](#).

Below are the steps for creating the named credential required for monitoring the Exadata Infrastructure and VM Clusters. This uses the API Signing Key for the Oracle IAM user granted the proper policies above. This is an RSA key pair **in PEM format** (minimum 2048 bits). See [How to Generate an API Signing Key](#).

Additionally, you will need the fingerprint of the public key. See [How to Get the Key's Fingerprint](#).

1. Click the setup icon , click **Security**, and select **Named Credentials**. The **Named Credentials** page is displayed.
2. Click **Create**. The **Create Credential** page opens.
- 3.

Security

[Named Credentials](#) > Create Credential

General Properties

* **Credential name**

Credential description

* **Authenticating Target Type**

* **Credential type**

Scope ☐ Target ☒ Global

Credential Properties

* **Tenancy OCID**

* **User OCID**

* **Public Key Fingerprint**

* **Private Key**

Private Key Passphrase

Provide the following information:

- **Credential Name:** Provide a suitable name to the credential. For example, OCI_CREDS.
- **Credential Description:** Describe the purpose of the credential and the intended use.
- **Authenticating Target Type:** Specify the target type for which this credential set will be used for authentication. Select Oracle Cloud Infrastructure from the menu.
- **Credential Type:** Specify the type of the credential that you're creating. Select Oracle Cloud Infrastructure Credential from the menu.
- **Scope:** Select the visibility of the credential in Enterprise Manager. Select Global.
- **Tenancy OCID:** The OCID of your OCI tenancy.
- **User OCID:** The OCID of the user to access OCI.
- **Public Key Fingerprint:** Provide the fingerprint of the public key of your API signing key pair.
- **Private Key:** Provide the private key of your API signing key pair.
- **Private Key Passphrase:** Specify the passphrase of your private key, if any.

Click **Save**.

You can test these credential for proper functioning by following the steps in [Test the Named Credentials Created for OCI Monitoring](#).

Create SSH Based Credentials

The *VM Guest* hosts are accessed through SSH and therefore the agents must be installed using SSH based named credentials. It is recommended that you use one of the two SSH credential methods described below.

- Create SSH based named credentials for the *opc* user that utilizes `sudo` to **oracle** and SSH based named credentials for the *opc* user that utilizes `sudo` to **root**. Provide the private and public SSH keys for the *opc* user.
- Create a new SSH key and add it to the *authorized_keys* file for the *oracle* user. Create an SSH based named credential for the *oracle* user using this newly created private and public key.

For steps to create SSH based named credentials and SSH keys, see *Setting Up SSH Key-based Host Authentication* in Host Authentication Features.


Create Storage Server Credentials

To associate the storage servers of the grid infrastructure with the cloud target, create a named credential for storing the ExaCLI user name and password used for connecting to the Exadata Storage Server.

For the relevant documentation about the user name and password, see:

- Exadata Database on Cloud@Customer: [Monitoring and Managing Exadata Storage Servers with ExaCLI](#) in *Exadata Database Service on Cloud@Customer Administrator's Guide*
- Exadata Database on Cloud@Customer: [Monitoring and Managing Exadata Storage Servers with ExaCLI](#) in *Exadata Database Service on Dedicated Infrastructure*

During the discovery of the cloud target, you will point to this named credential set for discovering the storage servers. This must be a named credential of the type *ExaCLI* or *RESTful API*.

1. Click the setup icon , click **Security**, and select **Named Credentials**. The **Named Credentials** page is displayed.
2. Click **Create**. The **Create Credential** page opens.
3. Provide the following information:
 - **Credential Name:** Provide a suitable name to the credential. For example, EXADATA_CRED.
 - **Credential Description:** Describe the purpose of the credential and the intended use.
 - **Authenticating Target Type:** Specify the target type for which this credential set will be used for authentication. Select Oracle Exadata Storage Server from the menu.
 - **Credential Type:** Specify the type of the credential that you're creating. Select Credential for ExaCLI or RESTful API from the menu.
 - **Scope:** Select the visibility of the credential in Enterprise Manager. Select Global.

- **Username and Password:** Provide the user name and password to access the storage cells. This is the exacli user that was generated as part of the Exadata Cloud Service setup.
- **Run Privilege:** You can select the level of privilege provided to the user for access.
- **Access Control** section: You can add grants, and select the grantee in this section.

Click **Save**.

You can now view the new credential created in the **Named Credentials** page.

Agent Installation

- [Overview of the Required Agents](#)
- [Install Agents](#)
- [Set Up a Proxy Connection](#)

Overview of the Required Agents

The Enterprise Manager Agent will be installed on the guest hosts of each VM cluster for the purpose of monitoring the resources that reside in that VM Cluster like host, database, listener, ASM, etc

In addition to these agents, one EM Agent outside of the Exadata Cloud is required for discovery and Exadata Cloud target resource monitoring. This agent is installed in a location that is able to communicate with the OCI API rest endpoints as well as the Enterprise Manager OMS servers. This agent can be used to discover and monitor multiple Exadata Cloud targets. Below are the recommendations for this agent:

- The agent should be placed on a host within the customer's data center.
- The chosen agent should not be one of the central agents on the Oracle Management Servers (OMS) nor should it be an agent that is installed on one of the VM guest hosts of the Exadata Cloud VM Cluster.
- The host should meet the requirements for installing an Enterprise Manager Agent. For details on the host requirements for installing an EM Agent, see Meeting the Generic Prerequisites for Installing Standalone Management Agents Using Add Host Targets Wizard or EM CLI.
- For HA/DR considerations, there should be both primary and backup agents selected but care should be taken to consider the proper location for the primary and backup agents.
- If Enterprise Manager is configured for HA/DR, agents must be located in sites that are accessible from both the primary and DR Enterprise Manager sites. That is, the agent selected for OCI monitoring should be on a host in the same data center or site where the primary EM is running and the backup agent should be located on a host in the data center or site where the EM DR environment is located.
- Agents must be able to communicate to the OCI API Rest endpoints, over public network, for example, by using a proxy or VPN. For more information about the endpoints, see [Internal Server Error in the Exadata Cloud discovery wizard](#).

An Enterprise Manager agent should also be installed on each of the guest VM hosts on the Exadata Cloud VM Cluster for the purpose of monitoring each of the resources that reside in that VM Cluster like host, database, ASM, etc.

Install Agents

Install an agent on the hosts selected for discovery and monitoring of the Exadata Cloud Services.

Install an agent onto each of the VM Cluster guest hosts using the SSH based named credential created earlier. Manually deploy the Exadata Plug-in to these agents.

For steps to install agents, see Installing Oracle Management Agents in *Enterprise Manager Basic Installation Guide*.

Set Up a Proxy Connection

To configure the EM agent to connect to OCI through a proxy server, use the following steps:

1. Log in to your agent host.
2. Edit the file \$EMDROOT/work/agentStateDir/sysman/config/s_jvm_options.opt. Add the following line at end of the file and save the file:

```
-Djdk.http.auth.tunneling.disabledSchemes=""
```

3. Edit the file \$EMDROOT/work/agentStateDir/sysman/config/emd.properties. Include the below properties in the file emd.properties:

```
oci_http_proxy_host= <proxy host>  
oci_http_proxy_port= <proxy port>  
oci_proxy_username= <proxy user>  
oci_proxy_password= <proxy password>
```

4. Restart the agent:

```
emctl stop agent ; emctl start agent ;
```

Test the Named Credentials Created for OCI Monitoring

You can test if the named credentials that you created for OCI monitoring are functioning in establishing communication with OCI. Before testing the credentials, ensure that the following two tasks are completed:

- [Create Named Credentials for OCI Monitoring](#)
- [Agent Installation](#)

For steps to test the credentials, see Test OCI Credentials in *Enterprise Manager Security Guide*.

2

Discover Exadata Cloud Targets

Topics:

- [Discover the Exadata Cloud Target Using Console](#)

Discover the Exadata Cloud Target Using Console

- Starting with Oracle Enterprise Manager *24ai Release 1 Update 4 (24.1.0.4)*, support is available for monitoring mixed configuration of Exadata VM Clusters where different VM clusters in Exadata Infrastructure support different storage management like ASM and Exascale.
- Starting with Oracle Enterprise Manager *24ai Release 1 Update 3 (24.1.0.3)*, support is available for Exascale storage management technology in Exadata Cloud targets.
- Starting with Oracle Enterprise Manager *24ai Release 1 Update 3 (24.1.0.3)*, the discovery flow supports Multicloud@Google Cloud and Multicloud@AWS targets.
- Starting with Oracle Enterprise Manager *24.1 Release Update 01*, the discovery of *Oracle Exadata Database Service on Exascale Infrastructure* target is available in Oracle Enterprise Manager.
- Starting with Oracle Enterprise Manager *24.1 Base release*, the discovery flow supports Multicloud@Azure target.
- Starting with Oracle Enterprise Manager *13.5 Release Update 21*, the discovery flow supports single-node VM cluster.
- Starting with Oracle Enterprise Manager *13.5 Release Update 14*, the discovery of *Oracle Exadata Database Service on Dedicated Infrastructure* target is available in Oracle Enterprise Manager. This discovery is integrated with the discovery of Oracle Autonomous VM Clusters and Exadata VM Clusters.
- Starting with Oracle Enterprise Manager *13.5 Release Update 12*, the discovery of *Oracle Exadata Database Service on Cloud@Customer* target is integrated with the discovery of Oracle Autonomous VM Clusters and Oracle Storage Servers.
- Starting with Oracle Enterprise Manager *13.5 Release Update 10*, the discovery of *Oracle Exadata Database Service on Cloud@Customer* target supports the discovery of Exadata VM Clusters.

The Exadata VM cluster targets of the type *X8M* and *X9M* are multi-VM clusters.

Before you begin, ensure that you completed the prerequisite tasks.

Note

The *Oracle Exadata Database Service on Exascale Infrastructure* target is available as Exadata VM Cluster top-level resource in the console. It does not have an associated Exadata Infrastructure resource in the console. It has its own associated Exascale Database Storage Vault resource and does not have an associated Exadata Storage Server. You can scale the Exadata VM Cluster ECPUs and node count.

Here are the steps to discover the Exadata Cloud target:

1. Open the discovery wizard:

- a. From Oracle Enterprise Manager menu, select **Setup, Add Target**, and then **Add Targets Manually**.
- b. In the **Add Targets Manually** page, click **Add Using Guided Process**.
- c. In the **Add Using Guided Process** dialog box, under the **Guided Discovery** column, select Exadata Cloud and click **Add**.

The Exadata Cloud option includes targets such as *Autonomous Exadata VM Cluster*, *Exadata VM Cluster*, *Oracle Exadata Storage Server*, *Oracle Exadata Storage Server Grid*, and *Exadata Infrastructure*.

Exadata Infrastructure covers the cloud targets *Oracle Exadata Database Service on Dedicated Infrastructure* and *Oracle Exadata Database Service on Cloud@Customer*.

Exadata VM Cluster includes the cloud target *Oracle Exadata Database Service on Exascale Infrastructure*.

The **Exadata Cloud Discovery** page opens.

2. Provide details of the monitoring agents:

- a. Under **Monitoring Agent**, select the agent that has access to OCI REST API endpoints. See [Agent Installation](#).
- b. Optionally, you can also specify an agent as **Backup Agent**. If the agent designated as Monitoring Agent is down, the Backup Agent will be auto-assigned to monitor the target.

3. Provide Oracle Cloud Infrastructure (OCI) details:

When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy in one region. This is the home region. It contains the account information and the identity resources. The subscribed regions are those regions to which you might have subscribed later for accessing resources in a different geographical location.

Note that the guided discovery process performs discovery of Exadata Cloud resources in a single subscribed region at a time.

- a. Select the **Home Region** in which your tenancy is provisioned. If your region is not listed in the menu, then select **Other**.

When you select **Other** as your home region, then you can specify the region identifier in the adjacent text box. Select from the **Region Identifier** column corresponding to your region listed at [Oracle Cloud Infrastructure Documentation: Regions and Availability Domains](#).

- b. Under **OCI Credentials**, from the menu, select the named credentials associated with your OCI account that you had identified as part of the prerequisite tasks. The list of

cloud resources returned is determined by the permissions of the user whose API key is provided.

- c. Select the **Subscribed Region** where you want to perform the discovery. The options available in this menu are updated based on your selection of home region, OCI named credentials, and monitoring agent.

Click **Next**.




4. **Select the Exadata Cloud targets to discover:**

Under **Discovered Targets**, select the Exadata Cloud infrastructure host target to discover. The columns in the table like **Deployment Type**, and **Target Already Managed** can help you to select the right targets.

For the *Oracle Exadata Database Service on Cloud@Customer* target, the deployment type is Cloud@Customer. For the *Oracle Exadata Database Service on Dedicated Infrastructure* and *Oracle Exadata Database Service on Exascale Infrastructure* targets, the deployment type is Oracle Cloud Infrastructure.

For the *Oracle Exadata Database Service on Exascale Infrastructure* target, the infrastructure type is Shared. The infrastructure type for other Exadata Infrastructure targets is Dedicated.

The various Exadata Cloud targets can be identified using the following icons:

- Exadata Database on Exascale Infrastructure VM Cluster: 
- Autonomous Exadata VM Cluster: 
- Exadata Infrastructure: 

5. **Provide the storage server details** (Optional):

This option is not available for *Oracle Exadata Database Service on Exascale Infrastructure* target.

If you have installed the monitoring agent on one of the VM cluster hosts before starting the discovery, then the button **Set Storage Server Credentials** is displayed in the column **Storage Server Credentials**. Click the button. The **Set Storage Server Credentials** dialog box opens.

Under **Storage Server Credentials**, from the menu, select the named credentials associated with the storage server that you had identified as part of the prerequisite tasks.

Under **OCI VM Cluster**, select the VM cluster which has the agent to monitor the storage server.

Under **Monitoring Mechanism**, select from ExaCli or RESTful API. If you selected RESTful API, then specify the **HTTPS Protocol**, and enable **Trust Self-Signed Certificate** if applicable.

Click **Save**.

At this point, based on your selection of Exadata Cloud infrastructure host target for discovery, all the underlying VM Clusters, Autonomous VM Clusters, and storage servers will automatically get discovered.

6. At a later point, if there are more child targets added to the infrastructure target, then you can re-run the discovery process to include them.

Click **Next**.

7. **Verify your selection of targets:**

Based on your selection of the infrastructure target in the prior step, a summary is displayed. If you provided the storage server details, then the Exadata Grid targets are also included in the summary.

Click **Submit** to complete the discovery.

After the discovery process is complete, a summary of the discovery is displayed. All the targets and their discovery status can be noted. The target types listed are Oracle Exadata Storage Server Grid, Exadata Infrastructure, Oracle Exadata Storage Server, Exadata VM Cluster, and Autonomous Exadata VM Cluster.

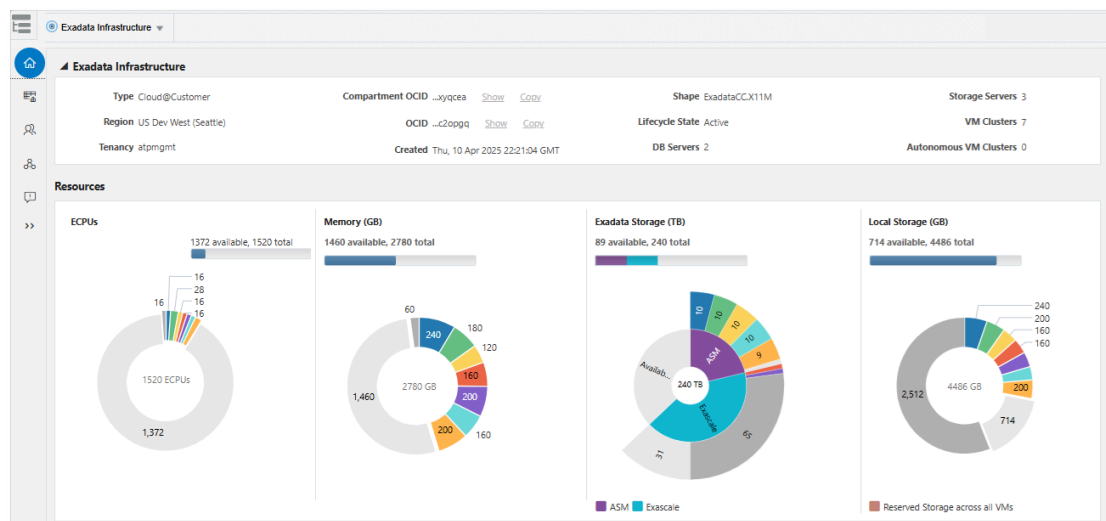
To view the target home pages, click **Show All Exadata Systems**. The page **Oracle Exadata Database Machines and Cloud Services** is displayed.

Explore the Exadata Cloud Target Home Pages

Topics:

- # Exadata Infrastructure Home Page

Home page of *Oracle Exadata Database Service on Dedicated Infrastructure* target with multi-VM Clusters:



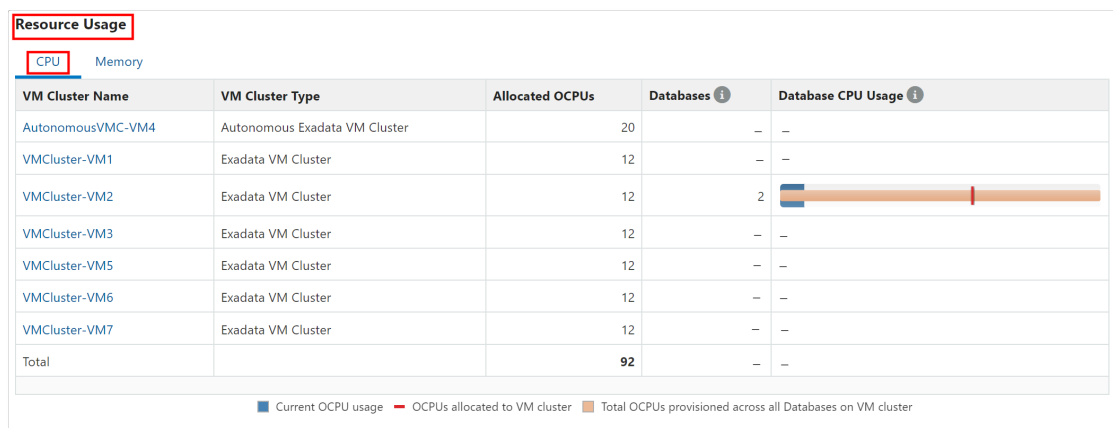
In the **Overview** tab, the **Exadata Infrastructure** section on this page provides the following details for this target:

- **Type:** Type of Exadata target, Cloud@Customer or Dedicated Infrastructure.
- **Region:** OCI region, for example, US East (Ashburn)
- **Tenancy:** The name of your OCI tenancy where your Exadata Cloud target is located
- **Compartment OCID:** The OCID of the compartment inside which the target resource is located
- **OCID:** OCID of the Exadata Infrastructure target
- **Created:** The time details of when the target was created
- **Shape:** The shape used for creating the target
- **Lifecycle State:** The state of the target to learn if it is active or in a different state.
- **DB Servers:** Number of underlying database servers
- **Storage Servers:** Number of underlying storage servers
- **VM Clusters:** Number of underlying VM clusters
- **Autonomous VM Clusters:** Number of underlying autonomous VM clusters

In the **Overview** tab, the **Resources** section on the page displays the availability and usage data of OCPUs (ECPUs in case of Exadata VM Cluster X11 based configuration), Memory (GB), Exadata Storage (TB), and Local Storage (GB) for the VM clusters. In case of single VM clusters of the target type *Oracle Cloud Infrastructure*, the chart for Local Storage (GB) is not available. The Exadata Storage resource displays the individual storage distribution in ASM and Exascale.

Exadata Storage section in the overview page displays the hierarchical storage distribution for each VM cluster. The inner ring shows the distribution of the total storage allocated to Exadata Infrastructure based on the storage type (ASM and Exascale). The outer ring corresponds to the storage allocation of each VM Cluster/Autonomous VM Cluster for each storage type. Hover your cursor over the chart to view the allocated storage for the corresponding slice. The value in the pie center corresponds to the total allocated storage for the infrastructure.

Resource Usage tab:



<

The Networks tab displays the network information for each Exadata VM cluster such as network name, network OCID, and its state. It also provides the control plane server IP address (es), the IPv4 mask used, and the gateway used.


The Networks tab is available only for the target type *Cloud@Customer*.

Incidents tab:

Incidents							
View	2	0	0	0	Type	All Incidents	
Severity	Summary	Target Name	Target	Status	Escalation Level	Age	Last Update Date
●	Failed to connect to ASM instance. Failed to connect: java...	+ASM2_scaqan10...	Automatic Storage Man...	New	No	2 days 9 hours	2 days 16 hours
●	Oracle Cloud VM Cluster VMCluster-scaqan10...-V...	VMCluster-scaqan10...	Oracle Cloud VM Cluster	New	No	2 days 17 hours	1 day 7 hours
Total: 2							


In the Incidents tab, for each member target, you can view the alerts and their details such as the severity of the issue, summary of the incident, target name, target type, status of the incident, its escalation level, the age of the incident, and the date when it was last updated.

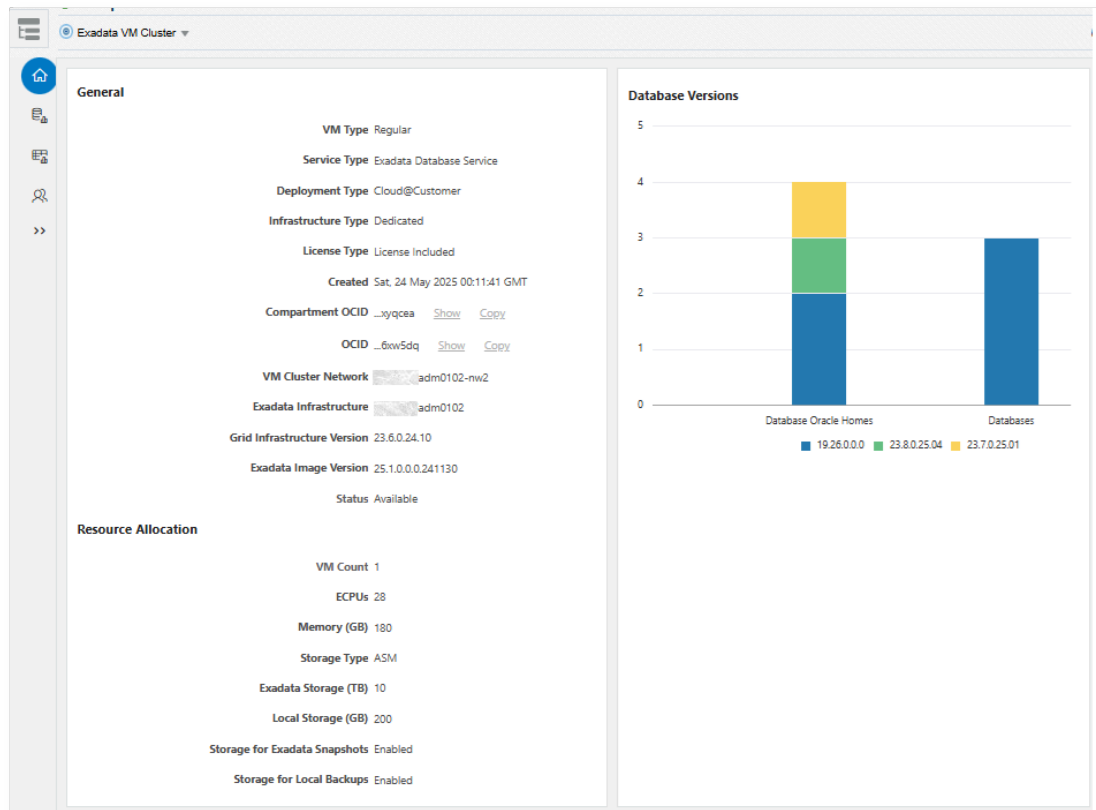
Target Navigation Tree

Click the **Target Navigation** icon  on the top left corner of the screen of the Exadata Infrastructure target home page to view the navigation tree.

The navigation tree lists all the member targets like Automatic Storage Management (ASM), Exadata VM Clusters, Autonomous Exadata VM Clusters, Databases, and Exadata Grid. Expand each member target to explore further. You can click the member target link to visit the member target home page to view its details.

Exadata VM Cluster Home Page

From Exadata Infrastructure home page, click the target Navigation icon , navigate the tree, and click the Exadata VM Cluster target link to open the home page.



Overview tab:

- General** section: The General section provides the following details for this Exadata VM cluster:
 - VM Type:** Type of the VM. This information is collected from OCI. For example, *Regular* or *Developer*.
 - Service Type:** Type of the service associated with the VM cluster. For example, *Exadata Database Service* or *Autonomous Database*.
 - Deployment Type:** Type of deployment of the Exadata VM cluster. For example, *Cloud@Customer*, *Oracle Cloud Infrastructure*, *Multicloud@Azure*, *Multicloud@Google Cloud*, or *Multicloud@AWS*.
 - Infrastructure Type:** The type of underlying infrastructure. For example, *shared* or *dedicated*.
 - License Type:** For example, *Bring Your Own License*.
 - Created:** The time details of when the Exadata VM cluster target was created
 - Compartment OCID:** The OCID of the compartment inside which the Exadata VM cluster target resource is located
 - OCID:** OCID of the Exadata VM cluster target
 - Vault OCID:** OCID of the Vault associated with the Exadata VM cluster target in case of *Oracle Exadata Database Service on Exascale Infrastructure*
 - VM Cluster Network:** The network associated with the Exadata VM cluster. This information is available only for the target type *Cloud@Customer*.

- **Exadata Infrastructure:** The parent target in case of clusters belonging to Exadata Infrastructure
- **Grid Infrastructure Version:** The version of the storage grid associated with the Exadata VM cluster
- **Exadata Image Version:** The version of the Exadata image used for creating the Exadata cloud target
- **Status:** Status of the Exadata VM cluster
- **Resource Allocation** section: This section provides the details for the Exadata VM cluster like VM Count, OCPUs (ECPUs in case of Exadata VM Cluster X11 based configuration), Memory (GB), Storage Type (ASM or Exascale), Exadata Storage (GB), Local Storage (GB), Storage for Exadata Snapshots and Storage for Local Backups. In case of *Oracle Exadata Database Service on Exascale Infrastructure*, the details available are VM Count, Reserved ECPUs, Enabled ECPUs, Memory (GB), Exascale database Storage, Storage Capacity (Free/Total, GB), and Total VM Storage (GB).
- **Database Versions** section: This section provides information about the databases available on the Exadata VM cluster.

Resource Usage tab:




Similar to Exadata Infrastructure home page, the Resource Usage tab displays the CPU usage information such as allocated OCPUs (ECPUs in case of Exadata VM Cluster X11 based configuration), number of databases, the database CPU usage and if the CPU is over-provisioned for the databases in the Exadata VM cluster. You can also view the CPU allocation by each database. Hover your cursor on the graph representing each database to view the CPU allocated to the VM and the current CPU usage of the database on the VM.

Members tab:

Target Name	Target Type	Status	0	1	2	EM Lifecycle Status	Version
Cluster	Cluster	Up	0	0	0	-	19.9.0.0.0
Host	Host	Up	0	0	0	-	7.9.0.0.0
Host	Host	Up	0	0	0	-	7.9.0.0.0
+ASM	Cluster ASM	Up	0	0	0	-	19.0.0.0.0
+ASM1	Automatic Storage Management	Up	0	0	0	-	19.0.0.0.0
+ASM2	Automatic Storage Management	Down	1	0	0	-	Not Available
+APX1	Asm Proxy	Up	0	0	0	-	19.0.0.0.0
DB19DEV	Cluster Database	Up	0	0	0	-	19.9.0.0.0
DB19DEV1	Database Instance	Up	0	0	0	-	19.9.0.0.0
DB19DEV2	Database Instance	Up	0	0	0	-	19.9.0.0.0
CDBROOT	Pluggable Database	Up	0	0	0	-	19.9.0.0.0
PDB19	Pluggable Database	Up	0	0	0	-	19.9.0.0.0
DB223003	Cluster Database	Up	0	0	0	-	19.9.0.0.0

The **Members** tab on the page displays a list of the members of the Exadata VM Cluster and its child targets, target type, status, and versions.

Oracle Autonomous Exadata VM Cluster Home Page

From Exadata Infrastructure home page, click the target Navigation icon , navigate the tree, and click the Autonomous Exadata VM Cluster target link to open the home page.

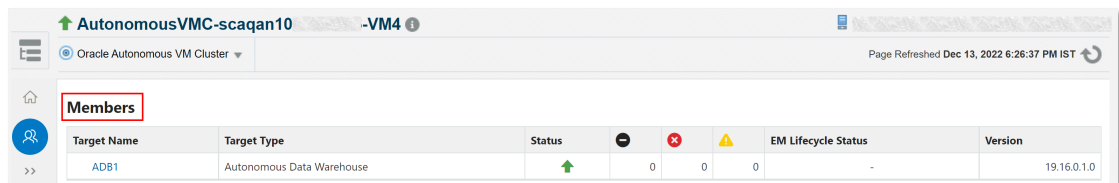
General	Resource Allocation
Service Type Autonomous Database	Node Count 2
Deployment Type Cloud@Customer	Total OCPUs 20
Infrastructure Type Dedicated	Available OCPUs 12
License Type License Included	Reclaimable OCPUs 0
Created Tue, 27 Sep 2022 21:32:04 GMT	Memory (GB) 200
Compartment OCID ...w5wu5q Show Copy	Autonomous Container 0 available, 1 total
OCID ...3npneq Show Copy	Database
VM Cluster Network VMClusterNetwork-adm0506	Autonomous Database 1.72 available, 2 total
Exadata Infrastructure adm0506	Storage (TB)
Status Available	Database Memory per OCPU (GB) 6
	Exadata Storage (TB) 2.866
	Local Storage (GB) 620
	Storage for Local Backups Not Enabled

Overview tab:

- **General** section: The General section provides the following details for this Autonomous Exadata VM cluster:
 - **Service Type:** Type of Autonomous Exadata VM cluster target

- **Deployment Type:** Type of deployment of the Autonomous Exadata VM cluster
- **Infrastructure Type:** Type of infrastructure the Autonomous Exadata VM cluster is deployed on
- **License Type:** For example, License Included
- **Created:** The time details of when the Autonomous Exadata VM cluster target was created
- **Compartment OCID:** The OCID of the compartment inside which the Autonomous Exadata VM cluster target resource is located
- **OCID:** OCID of the Autonomous Exadata VM cluster target
- **VM Cluster Network:** The network associated with the Autonomous Exadata VM cluster. This information is available only for the target type *Cloud@Customer*.
- **Exadata Infrastructure:** The parent target
- **Status:** Status of the Autonomous Exadata VM cluster
- **Resource Allocation** section: This section provides the details for the Autonomous Exadata VM cluster like Node Count, Total OCPUs, Available OCPUs, Reclaimable OCPUs, Memory (GB), Autonomous Container Database, Autonomous Database Storage (TB), Database Memory per OCPU (GB), Exadata Storage (GB), Local Storage (GB), and Storage for Local Backups (This information is available only for the target type *Cloud@Customer*).

Members tab:




Target Name	Target Type	Status	0	0	0	EM Lifecycle Status	Version
ADB1	Autonomous Data Warehouse	Up	0	0	0	-	19.16.0.1.0

The **Members** tab on the page displays a list of the members of the Oracle Autonomous Exadata VM Cluster and its child targets, target type, status, and versions.

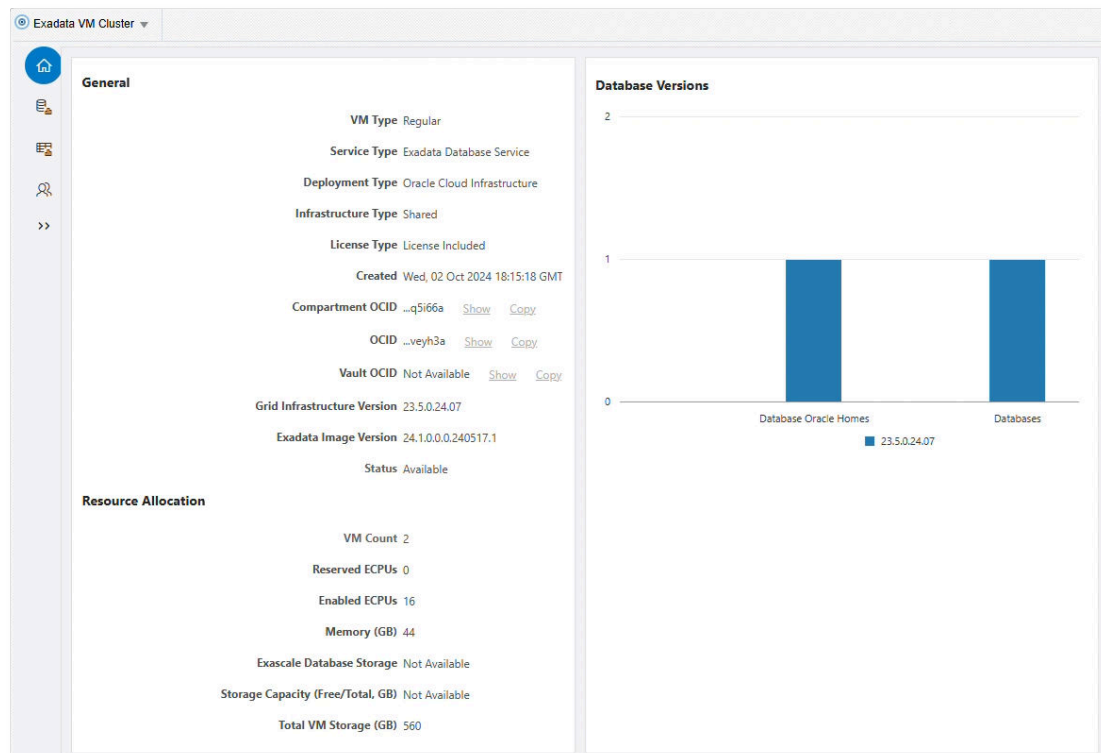
Oracle Exadata Database Service on Exascale Infrastructure Home Page

Go to Enterprise Manager menu, click **Targets**, and click **Exadata**. This action navigates to a page that lists all the Exadata systems monitored by Enterprise Manager, both on-premises and in the cloud. Click the link on the Exadata VM Cluster that corresponds to the *Oracle Exadata Database Service on Exascale Infrastructure* target name to open the home page.

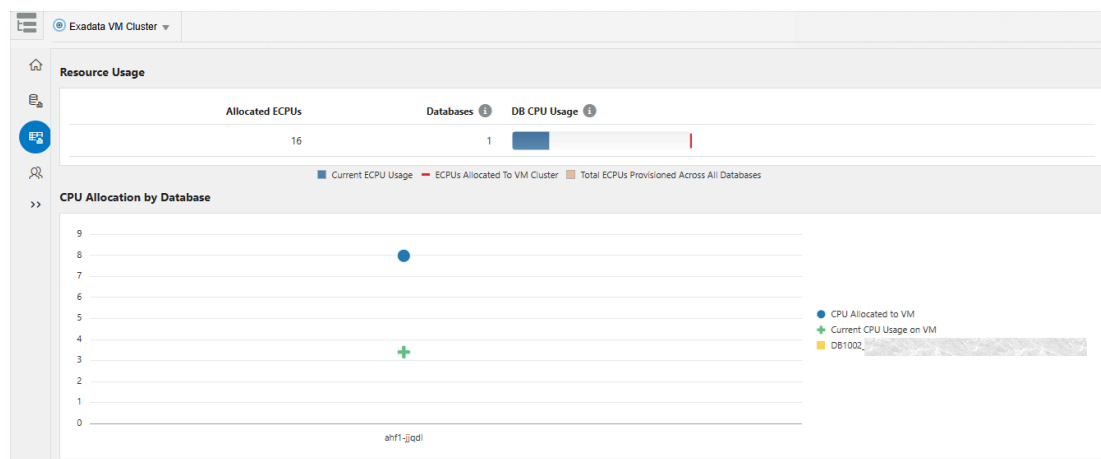
Click the **Target Navigation** icon  on the top left corner of the screen of the Exadata VM Cluster target home page to view the navigation tree.

The navigation tree lists all the member targets like hosts, Automatic Storage Management (ASM), and databases. Expand each member target to explore further. You can click the member target link to visit the member target home page to view its details.

Overview tab in the home page:



Resource Usage tab in the home page:



For details about the various parameters displayed in the Overview tab and Resource Usage tab, see [Exadata VM Cluster Home Page](#).

Part III

Exadata Cloud Lifecycle Management

Perform database lifecycle management activities like create, patch, clone, relocate, strengthen security posture with compliance management, and establish standardization of configuration on your pluggable database across your Exadata Cloud hybrid environment. You can also set up PDB-as-a-Service (PDBaaS) and adopt a DevOps model of database deployment on your Exadata Cloud environment.

The Database Lifecycle Management and Cloud Management Pack features are available:

- Starting with Oracle Enterprise Manager 13.5 Release Update 10 on *Oracle Exadata Database Service on Cloud@Customer* (ExaCC).
- Starting with Oracle Enterprise Manager 13.5 Release Update 21 on *Oracle Exadata Database Service on Dedicated Infrastructure* (ExaCS).

Before You Begin

Before you begin Database Lifecycle Management activities for your Exadata Cloud database targets, you must ensure that they are up and discovered in Oracle Enterprise Manager. See [Discovery of Exadata Cloud](#).

Oracle Database 19.6 or higher is required to perform Exadata Cloud Lifecycle Management operations.

In order to discover assets first you need to install agents on Exadata Cloud resources, for more information see: Deploying Agents on OCI Resources in *Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Once the agents have been deployed discover the targets in Enterprise Manager, for more information see [Discovery of Exadata Cloud](#).

Topics:

- [Database Lifecycle Management in Exadata Cloud](#)
- [Hybrid Multicloud Management - Exadata Cloud](#)

4

Database Lifecycle Management in Exadata Cloud

You can perform the following database lifecycle management tasks on your Exadata Cloud environment by adopting *Database Lifecycle Management Pack*:

For more information on the pack, see Database Lifecycle Management Pack for Oracle Database in *Oracle Enterprise Manager Licensing Information User Manual*.

Topics:

- [Provision a Pluggable Database \(PDB\)](#)
- [Update and Upgrade PDBs using Fleet Maintenance](#)
- [Relocate a PDB](#)
- [Unplug and Plug a PDB](#)
- [Manage Configuration of PDB](#)
- [Manage Compliance of a PDB](#)
- [Clone a PDB](#)
- [Delete a PDB](#)

Before You Begin

Before you begin with Lifecycle Management operations on Exadata Cloud, make sure that the following prerequisites are met:

- Oracle Enterprise Manager 13.5 Release Update 10 (13.5.0.10) or higher is deployed for *Oracle Exadata Database Service on Cloud@Customer*, and Oracle Enterprise Manager 13.5 Release Update 21 (13.5.0.21) or higher is deployed for *Oracle Exadata Database Service on Dedicated Infrastructure*.
- The Exadata Cloud targets are discovered in Oracle Enterprise Manager 13.5 or higher. For more information on discovering Exadata Cloud Targets in Oracle Enterprise Manager, see [Discovery of Exadata Cloud](#).

Note

Exadata Cloud PDBs use TDE (Transparent Data Encryption). In line with Oracle's best practices, Oracle Enterprise Manager only supports TDE configuration when it is configured using `init.ora`. To review additional information and actions on configuring and managing TDE see: [Using Transparent Data Encryption](#) in *Oracle Database Advanced Security Guide*, and Managing Oracle Transparent Data Encryption in *Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Provision a Pluggable Database (PDB)

In order to provision a PDB follow these steps:

To provision and create a PDB follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching** then select **Database Provisioning**. In the Database Provisioning page, under **Related Links** of the left menu pane, click **Provision Pluggable Databases**.
2. In the Provision Pluggable Database Console, under the Container Database section, select the CDB where you want to create new PDBs.
3. In the PDB Operations section, select **Create New Pluggable Databases**.
4. Click **Launch**.
5. In the **Creation Options** page under **Pluggable Database Creation Options** select **Create a New PDB**.
6. Within **Container Database Host Credentials** select or specify the target CDB Oracle home owner host credentials. Select **Preferred**, **Named** or **New**. For Exadata Cloud RAC use the OPC user for Normal Host credentials.
7. Click **Next**.
8. In the Identification page, enter a unique name for the PDB you are creating. Enter the TDE Wallet Password, and the Wallet Location where the Wallet will be stored. For RAC databases, the wallet location should be a shared location accessible from all nodes.
9. In the PDB Administrator section enter the credentials for the admin user account.
10. Click **Next**.
11. In the Storage page for **PDB Datafile Locations** select the type of location where the datafiles will be stored.
12. In the Temporary Working Directory section, enter a location where the temporary files generated during the PDB creation process can be stored.

Note

For RAC deployments the location should be shared and accessible to all RAC nodes.

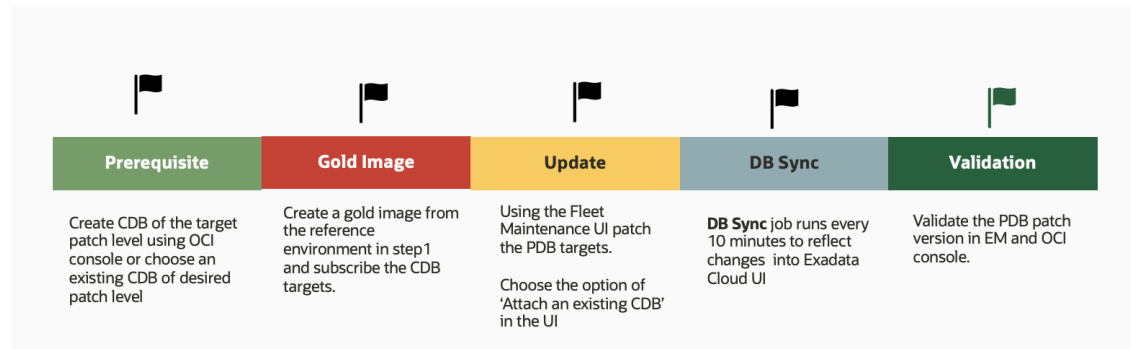
13. In the Post-Creation Scripts section, select a custom SQL script you want to run as part of this procedure, once the PDB is created.
14. Click **Next**.
15. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment.
16. Click **Next**.
17. In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click **Submit**.
18. In the Procedure Activity page, view the status of the procedure. From the **Procedure Actions** menu, you can select **Debug** to set the logging level to Debug, and select **Stop** to stop the procedure execution.

For more information on provisioning PDBs see: *Creating a New Pluggable Database in Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Update and Upgrade PDBs using Fleet Maintenance

Oracle Enterprise Manager Fleet Maintenance offers a complete and intuitive way to patch and upgrade databases, follow these steps to patch and upgrade PDBs:

Overview



Steps to Patch and Upgrade

Before you begin patching, ensure that:

- There is an existing CDB at the target patch level or create a new CDB through the OCI console or cloud tooling with the target patch version to successfully patch or upgrade a PDB.
 - The Gold image of the target CDB home is created using EMCLI or REST API. Optionally, you can also use the Fleet Maintenance Hub to create gold image using UI. See Fleet Maintenance Hub.
 - The CDB-level targets are subscribed using EMCLI or REST API commands, which in turn will subscribe the targets at PDB level.
1. To begin performing Fleet Maintenance operations, go to **Targets** and select the **Databases**
 2. Select **Administration** then **Fleet Maintenance**. In the **Fleet Maintenance** home page,
 - a. Select the Gold image from the **Image** drop down menu.
 - b. Select the target type **Pluggable Database**.
 - c. In the **Operation** drop down menu, select Update or Upgrade based on your requirement.
 - d. Select the PDB targets that you want to patch or upgrade and click **Next**.
 3. Within the **Operations** screen check mark the operations to be performed under **Maintenance Tasks**:
 - a. Under Maintenance Tasks, select **Attach an existing CDB** as Destination CDB.
 - b. In the **Attach an Existing CDB** section, select your Destination CDB.
 - c. **Update PDB**: Select a start time or set to start automatically after the previous step.
 4. Enter the **Credentials** for Normal Host, Privileged Host and Database Credentials. You can select whether to use a Named or Preferred credentials. For Exadata Cloud RAC use the OPC user for Normal Host credentials.

5. Under **Options**, you can set **Work Directory**, and optionally upload **Custom Scripts** which can be executed either before or after update. Click **Next**.
6. In the **Review** screen, verify all options entered in the previous screen. Click on **Validate** to perform a **Quick** or **Full** validation of your Fleet Maintenance Operation before submitting.
7. To start your Fleet Maintenance operation, click **Submit**.

For more information on Fleet Maintenance activities see Fleet Maintenance Hub in *Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Relocate a PDB

To relocate a PDB follow these steps:

1. From the **Targets** menu, select **Databases**.
2. Under **View** select **Search List**. From the **View** menu select **Expand All**.
3. Look for the source PDB and right-click the name of the PDB that you want to relocate.
4. From the **Oracle Database** menu, and select **Relocate**.
5. On the Source and Destination: Relocate Pluggable Database page specify the SYSDBA and host credentials for the source CDB. For Exadata Cloud RAC use the OPC user for Normal Host credentials. In the Container Database Destination section, select the destination CDB and specify a name, and a display name for the PDB.
If you want to use an existing database link for a PDB relocation, click on **Select an existing database link** option and select the database link to be used. If you do not select one, then the system will create a database link to connect to the source PDB or database.

Provide the target CDB TDE wallet password.

Select either of the following Relocation modes:

- Maximum Availability ensures there is no need to update the Connect Descriptor soon after the PDB relocation. You may do it at a later time, if required.
 - Normal Availability requires you to update the Connect Descriptor after the PDB relocation.
6. Click **Relocate**.

For more information on PDB relocation see: Relocating Pluggable Databases in *Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Unplug and Plug a PDB

To plug and unplug PDBs follow these steps:

Unplug a PDB

1. From the **Enterprise** menu, select **Provisioning and Patching** select **Database Provisioning**. In the Database Provisioning page under Related Links, click **Provision Pluggable Databases**.
2. In the Provision Pluggable Database Console, select the CDB from which you want to unplug the PDBs.
3. In the PDB Operations section, select **Unplug Pluggable Database**.
4. Click **Launch**.

5. In the Select PDB page of the Unplug Pluggable Database Wizard, in the Select Pluggable Database section, select the PDB you want to unplug. Note that the PDB once unplugged will be stopped and dropped.
6. In the CDB Host Credentials section, select or specify the target CDB Oracle Home owner host credentials. Select **Preferred**, **Named** or **New** credentials. For Exadata Cloud RAC use the OPC user for Normal Host credentials.
7. In the Destination page, select the type of PDB template you want to generate for unplugging the PDB, and the location where you want to store it. The PDB template consists of all datafiles as well as the metadata XML file.
 - Select **Target Host File System** to store the PDB template on your CDB host (where you are unplugging the PDB)
 - Select **Generate PDB File Set** to generate an RMAN file set. This creates a DFB file with all the data files and a separate metadata XML file.

Note

Oracle recommends you to select this option if the source and target CDBs are using ASM for storage.

- Select **Generate PDB Metadata File** to generate only a metadata XML file, and select a location where the metadata XML file will be created. In this option the PDB datafiles remain on the source PDB location after unplugging.
- **Generate PDB Archive** a TAR file with the datafiles and the metadata XML file included in it is not a valid option for Exadata Cloud TDE enabled databases 19c and above. Oracle recommends that you select this option if the source and target CDBs are using file system for storage. This option is not supported for PDBs using ASM as storage.
- Select **Software Library** to store the PDB template in the Oracle Software Library.
 - Select **Generate PDB File Set** to generate an archive file set. This creates a DFB file with all the data files and a separate metadata XML file. This is not a valid option for Exadata Cloud TDE enabled databases 19c and above.
 - Select **Generate PDB Metadata File** to generate only a metadata XML file, and select a location where the metadata XML file will be created.
 - **Generate PDB Archive** a TAR file with the datafiles and the metadata XML file included in it is not a valid option for Exadata Cloud. Oracle recommends that you select this option if the source and target CDBs are using file system for storage. This option is not supported for PDBs using ASM as storage.
- Enter a unique PDB template name. The template is created in the default location: Database Configuration/<db_release>/<platform>/Database Templates
- Enter a temporary location where the archive can be created by Enterprise Manager before it is uploaded to the Software Library.

If the CDB is TDE enabled, enter the TDE wallet password. Provide a new wallet transport secret and wallet location where the wallet will be stored after unplug operation. In case of RAC database, ensure that the TDE wallet location is accessible from all nodes.

8. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment.
9. Click **Next**.

10. In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details click **Submit**.
11. In the Procedure Activity page, view the status of the procedure. From the **Procedure Actions** menu select **Debug** to set the logging level debug. Select **Stop** to stop the procedure execution.

For more information on unplugging a PDB see Unplugging and Dropping a Pluggable Database in *Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Plug a PDB

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**. In the Database Provisioning page, in the Related Links section of the left menu pane, click **Provision Pluggable Databases**.
2. In the Provision Pluggable Database Console, under the Container Database section, select the CDB to which you want to add the unplugged PDBs.
3. In the PDB Operations section, select **Create New Pluggable Databases**.
4. Click **Launch**.
5. In the Creation Options page of the Create Pluggable Database Wizard, in the Pluggable Database Creation Options section, select **Plug an Unplugged PDB**.
6. In the Container Database Host Credentials section, select or specify the target CDB Oracle home owner host credentials, you can select **Preferred**, **Named**, or **New**.
7. Click **Next**.
8. In the Identification page, enter a unique name for the PDB you are plugging in. Select **Create As Clone** to ensure that Oracle Database generates a unique PDB DBID, GUID, and other identifiers expected for the new PDB.

For a TDE enabled PDB enter the Wallet Password and Wallet Transport for the PDB plug operation.
9. If you want to create a brand new admin user account for the PDB you are plugging in select **Create PDB Administrator** and enter the desired credentials. If not deselect **Create PDB Administrator**.
10. In the PDB Template Location section, select the location where the source PDB's template is available, and then select the type of PDB template. Plugging In an Unplugged Pluggable Database displays the PDB Template Location section of the Identification page.
11. For TDE enabled CDB enter the wallet file, destination CDB wallet password, and Wallet Transport Secret passed during the unplug operation.
12. Click **Next**.
13. In the Storage page, select **Use Oracle Managed Files (OMF)** or **Use Common Location for PDB datafiles** depending on your desired configuration.
14. In the Temporary Working Directory section, enter a location where the temporary files generated during the PDB creation process can be stored.
15. In the Post-Creation Scripts section, select a custom SQL script you want to run as part of this procedure, once the PDB is plugged in.
16. Click **Next**.
17. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment.

18. Click **Next**.

19. In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click **Submit**.

For more information on plugging a PDB see: *Plugging In an Unplugged Pluggable Database in Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Manage Configuration of PDB

Oracle Enterprise Manager offers a centralized platform for managing database configurations across multiple instances. It allows you to:

- Using Oracle Enterprise Manager, you can create and maintain configuration templates to enforce standardization and best practices.
- Compare configurations across various instances to identify inconsistencies.
- Track configuration changes and generate alerts to detect unauthorized modifications.

Oracle Enterprise Manager Configuration Management enables you to view, save, track, compare, search, and customize collected configuration information for all managed targets known to Oracle Enterprise Manager. Additionally, the Configuration Topology Viewer provides a visual layout of a target's relationships with other targets; for example, you can determine a system's structure by viewing the members of a system and their interrelationships.

With Configuration Management you review the following information:

- Comparisons
- Configuration Searches
- Configuration Browser
- Configuration History
- Configuration Extensions and Collections
- Configuration Topology Viewer
- Parsers
- Relationships

For more information see: *Manage Configuration Information in Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Manage Compliance of a PDB

With Oracle Enterprise Manager, you can ensure that your PDBs are within established compliance benchmarks.

Compliance Management provides the ability to evaluate the compliance of targets and systems as they relate to business best practices for configuration, security, and storage. This is accomplished by defining, customizing, and managing compliance frameworks, compliance standards, and compliance standard rules. In addition, Compliance Management provides advice of how to change configuration to bring your targets and systems into compliance.

There are three main areas for Compliance Management:

- Compliance Frameworks
- Compliance Standards

- Compliance Standard Rules

For more information see: *Manage Compliance Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

AHF Exadata Infrastructure

Oracle AHF EXAchk is a lightweight and non-intrusive health check framework for stack of software and hardware components in Exadata. Oracle Enterprise Manager provides a set of compliance standards and associated controls for overall health monitoring, automated risk identifications and proactive notification of issues for each Exadata System component and database instances.

Oracle Enterprise Manager 13c Release 5 Update 12 (13.5.0.12) integrates Oracle Autonomous Health Framework (AHF) EXAchk for *Oracle Exadata Database Service on Cloud@Customer* (ExaCC) targets.

AHF EXAchk standards are available out-of-box. When you update AHF, the associated EXAchk compliance standards are automatically updated to the corresponding version.

For more information, see:

- Exadata System Target
- Prerequisites for EXAchk Standards
- Associate EXAchk Standards

Clone a PDB

Oracle Enterprise Manager has the ability to clone PDBs, follow these steps:

Clone PDB Prerequisites

- When cloning a PDB from one CDB to another CDB make sure both the source and target CDB have the same database domain name.

Note

This is not required if cloning on the same CDB.

- When cloning a PDB from one CDB to another CDB inside Exadata Cloud, make sure that both source and target CDB have the same database domain and the global names have the correct domain suffix.

Steps to Clone a PDB

1. From the **Targets** menu, select **Databases**.
2. For **View**, select **Search List**. From the **View** menu, select **Expand All**.
3. Look for the source CDB in the list, and right-click the name of the PDB that you want to clone and select **Oracle Database**.
4. From the **Oracle Database** menu, select **Cloning** and then **Create Full Clone**.
5. On the Create Full Clone Pluggable Database page specify the following:
 - SYSDBA credentials
 - Name and display name for the PDB Clone

- Select Administrator and Host credentials (For Exadata Cloud RAC use the OPC user for Normal Host credentials).
 - Select the Refreshable clone check box to create a manually Refreshable PDB.
 - Enter the destination CDB TDE wallet password.
6. If you want to clone a PDB to a different CDB click the **Destination CDB** check box and select the destination CDB. Enter the PDB Admin user name and password for the cloned PDB.
 7. Specify additional configuration details like mask data, as well as schedule the cloning process, by selecting the **Advanced** option.
 8. Click **Clone**.

For additional information and detailed steps see: *Creating a Full Clone Pluggable Database in Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Delete a PDB

To delete a PDB follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**. In the Database Provisioning page, in the Related Links section of the left menu pane, click **Provision Pluggable Databases**.
2. In the Provision Pluggable Database Console, in the CDB section, select the CDB from which you want to delete the PDBs.
3. In the PDB Operations section, select **Delete Pluggable Databases**.
4. Click **Launch**.
5. In the Select PDBs page of the Delete Pluggable Databases Wizard, click **Add** to select the PDBs that you want to delete, then click **Select**.
6. In the CDB Host Credentials section, select or specify the target CDB Oracle Home owner host credentials. If you have already registered the credentials with Enterprise Manager, you can select **Preferred** or **Named**. Otherwise, you can select **New** and enter the credentials. For Exadata Cloud RAC use the *opc* or *su - oracle* user for Normal Host credentials.
7. In the Schedule page, enter a unique deployment procedure instance name and a schedule for the deployment.
You can run the procedure **Immediately** or **Later** and provide time zone, start date, and start time details.

You can optionally set a grace period for this schedule. A grace period is a period of time that defines the maximum permissible delay when attempting to run a scheduled procedure. If the procedure does not start within the grace period you have set, then the procedure skips running. To set a grace period, select **Grace Period**, and set the permissible delay time.
8. Click **Next**.
9. In the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, click **Submit**.
10. In the Procedure Activity page, view the status of the procedure. From the **Procedure Actions** menu, you can select **Debug** to set the logging level to Debug, and select **Stop** to stop the procedure execution.

For more information on provisioning PDBs see: Deleting Pluggable Databases in *Oracle Enterprise Manager Database Lifecycle Management Administrator's Guide*.

Hybrid Multicloud Management - Exadata Cloud

Using Oracle Enterprise Manager, you can take advantage of the Hybrid Cloud Enablement feature. This feature leverages and consolidates resources across your environment, which can range from on-premises, Exadata Cloud, Oracle Cloud, and other public clouds, to form a dedicated private cloud known as a *Hybrid Multicloud*.

The DevOps model of database management, known as *Database-as-a-Service*, helps break down these silos by creating a streamlined and efficient model for database management. Similarly, you can enable DevOps users to deploy pluggable database on demand on your Hybrid Cloud environment which is referred as *Pluggable Database-as-a-Service* (PDBaaS). Currently, only PDBaaS is supported on Exadata Cloud.

You can deploy PDBaaS on top of your *Hybrid Multicloud*, resulting in an agile, efficient, and cost-effective way of managing databases which is referred to as *Hybrid Database-as-a-Service*.

As a DevOps user, you also get a dedicated self-service portal to provision and manage your databases on your own. The dedicated REST APIs and CLIs can be integrated with your orchestration engines like Ansible and Terraform to automate repetitive tasks.

Pluggable Database-as-a-Service (PDBaaS) is an on-demand platform that allows the consolidation of Exadata Cloud resources to form a private cloud and provide PDB management as a cloud offering.

Self Service users can provision and manage PDBs in a few simple steps.

Topics:

- [Setting up Pluggable Database as a Service](#)
- [Create an Empty PDB](#)
- [Create a PDB With Data](#)
- [Unplug and Plug a PDB](#)
- [Dehydrate and Rehydrate a PDB](#)
- [Resize PDB](#)
- [Delete a PDB](#)

Setting up Pluggable Database as a Service

Before you start with the Management operations of Oracle Exadata Database Service on Cloud@Customer, PDBaaS setup is required.

1. To enable PDBaaS, follow the steps in the Getting Started section. See Getting Started with Database as a Service .
2. Create one or more resource providers.
 - Create one or more PaaS Infrastructure Zones.

- Create one or more database pools.
See [Creating Resource Providers](#).
- 3. Configure the request settings. See [Configuring Request Settings](#).
- 4. Define quotas for each self service user. See [Defining Quotas](#).
- 5. Create a database provisioning profile. This step is optional and is not required if you are creating an empty pluggable database. See [Creating a Database Provisioning Profile Using Export Schema Objects](#).
- 6. Create a service template. A service template can contain:
 - Empty Pluggable Database: In this case, the service template is created with an empty schema.
 - Pluggable Database from a Profile: In this case, you can import schemas from a database provisioning profile. You can select this option for applications with data such as eBusiness applications.See [Creating a Service Template for PDB Based on Profile](#).
- 7. Configure the Chargeback Service. See [Configuring Chargeback](#).
- 8. While deploying a database, select the service template that you have created. See [Requesting a Pluggable Database](#).

For more details, watch the video [How to set up Pluggable Database-as-a-Service on Hybrid Multicloud](#).

Create an Empty PDB

Use the Enterprise Manager console to create a new empty PDB

To create an empty PDB, first follow the Getting Started with DBaaS steps for Provisioning a PDB outlined in Getting Started in *Oracle Enterprise Manager Cloud Administration Guide*.

Once you have completed the getting started steps and created a service template for an Empty Pluggable Database, follow these steps to request a PDB:

1. Log in to Enterprise Manager as a user with EM_SSA_USER role.
2. Expand **Enterprise**, select **Cloud** and select **Self Service Portal**.
3. Expand the right hand action menu for Database Cloud Services and select **Open Service Console**.
4. In the Database Cloud Services page click on **Create Instance**, a pop up will appear. Click **Select** on the correct Service Offerings (Service Template) for your system.
5. The Database Cloud Self Service Portal Create Pluggable Database page is displayed. Enter the following information:
 - Pluggable Database Configuration
 - Service Template (this field is auto-populated with the selection from step 4)
 - PDB Name
 - Database Service Name
 - Size
 - Pluggable Database Administrator Account
 - Administrator Name

- Password
 - Confirm Password
- Tablespaces
 - Tablespace Name
- Instance Details
 - Request Name
 - Zone
- Properties (Optional)
 - Comment
 - Line of Business
 - Location
 - Downtime Contact
 - Department
 - LifeCycle Status
 - Contact
- Duration
 - Start
 - Duration

6. Click **Submit** to deploy the database to the selected zone.

For detailed information, see Requesting a Pluggable Database in *Oracle Enterprise Manager Cloud Administration Guide*.

Also, watch the video [How to provision a pluggable database](#).

Create a PDB With Data

You can create a PDB with data by following these instructions:

Pluggable databases with data can be created using service templates. To create a service template for PDB as a service, follow these steps:

1. Log into Enterprise Manager as a user with the EM_SSA_ADMINISTRATOR role.
2. From the **Enterprise** menu, select **Cloud**, then select **Cloud Home**. The Cloud Home page appears. From the **Oracle Cloud** menu, select **Setup**, then select **Getting Started**.
3. Click the **Setup** link next to the Databases option in the left panel. Select **Pluggable Database** from the drop down menu.
4. Click the **Service Templates** link. The Service Templates page appears. Click **Create**. The Create Service Template: General page appears.
5. Enter a name and description for the service template. The description must be unique and provide information on the type of service template being created.
6. In the Pluggable Database region, select the **Create Pluggable Databases from Profile** option and select a database provisioning profile that is to be used to create the pluggable database.

7. Click the **Search** icon next to the Profile field and select a profile from the list. If a profile does not exist, you must create one.
8. In the Pools and Zones region, click **Add** to select a PaaS Infrastructure Zone into which the pluggable database instance is to be provisioned. Click **Assign Pool** to add a Database Pool (for Pluggable Databases) for the zone.

Note

A database pool (for PDBs) must be present in the PaaS Infrastructure Zone that you select here.

9. After the database pool has been assigned, the container database into which the pluggable database is to be added is displayed in the Reference Container Database field. All the validations for the service template will be performed against this database.
10. Select **Create Pluggable Databases from Profile** in the Pluggable Database region, in the Shared Location field specify the Path or the shared location on all hosts in which the schema dumps are present. The Dump File Location must be accessible by all the databases.
11. In the Placement region, the pluggable database can be provisioned on the container database using one of the following options: Selected by placement algorithm or Selected by user during request.
12. In the Identification region, in the Pluggable Database Prefix Name field, specify the prefix that is to be used to generate a unique Pluggable Database name at the time of database creation. The prefix helps to identify the PDBs that are created using this service template.
13. Click **Next**. The Create New Service Template: Configurations page appears.
14. Click **Next**. The Create Service Template: Initialization Parameters page appears. In this page, you can configure the values of various initialization parameters that affect the operation of the database instance.
15. Click **Next**. The Create Service Template: Customization page appears. Specify the custom scripts that need to be executed before and after the database is created.
16. In the Target Properties region, you can specify global target properties such as Location, Department, Line of Business, and so on. You can mark a property as mandatory or lock a property.
17. Click **Next**. The Create Service Template: Roles page appears. Click **Add** to select the EM_SSA_USER roles to which this service template will be available. All users belonging to the selected role can use this service template.
18. Click **Next**. The Create Service Template: Review page appears where you can review the details entered so far.
19. Click **Submit**. The newly created service template will appear in the Database Cloud Self Service Portal Setup: Service Templates page. You can click on the Service Template Name link to view additional details.

For more information and a complete detailed set of steps and options see: Creating a Service Template for PDB Based on Profile.

Unplug and Plug a PDB

Plugging and Unplugging cloud PDBs can be performed in Enterprise Manager using the same steps as a non-cloud PDB.

To unplug or plug a cloud based PDB follow the steps located in [Unplug and Plug a PDB](#).

Dehydrate and Rehydrate a PDB

Dehydrate a PDB

Enterprise Management provides an option for a Self Service (SSA) user to create a data profile while deleting the PDB. By enabling this option, you can preserve the backup of the PDB instance when an SSA user attempts to delete the PDB. The backup is stored in the Software library by default, and the location of the profile can be customized too.

This functionality is referred to as *Dehydrate a PDB*.

To dehydrate a PDB, see *Creating a Service Template for PDB Based on Profile* in *Oracle Enterprise Manager Cloud Administration Guide*.

Rehydrate a PDB

The Enterprise Manager self-service portal enables you to create a new PDB using a preexisting Data profile created by the SSA user during Dehydrate. This template permits the SSA user to plug the PDB back into the CDB pool and this functionality is referred to as *Rehydrate a PDB by an SSA user*.

To rehydrate a PDB, see *Pluggable Database creation using SSA User Data Profile* in *Oracle Enterprise Manager Cloud Administration Guide*.

Resize PDB

You can resize PDB instances by initiating a resize operation to a workload size defined in the corresponding service template. This resize operation ensures sufficient availability and effective utilization of resources.

Prerequisites

Before you begin resizing a PDB make sure you complete all prerequisites. For more information see: *Prerequisites PDB Resize* in *Oracle Enterprise Manager Cloud Administration Guide*.

PDB Resize: Self-service Portal

You can use the Resize button to initiate a resize operation. The Resize button is enabled only if the service template for PDB has more than one workload.

After the Resize button is clicked, you can select a size from the drop-down list.

Limitations

- If the PDB's actual storage consumption is greater than the storage value in the new workload, PDB scale down operation is not possible.
- If sufficient resources are not available on the target host for PDB resizing, PDB Relocation must be used.

For more information on PDB Resizing see: *PDB Instance Resize* in *Oracle Enterprise Manager Cloud Administration Guide*.

Part IV

Administration and Monitoring

For the various administrative tasks and the monitoring information available for Exadata Cloud, see below.

Topics:

- [Administer Exadata Cloud Service Target](#)
- [Monitor Exadata Cloud Service](#)

6

Administer Exadata Cloud Service Target

After the discovery of Exadata Cloud Service target, you can perform the following administration tasks:

Topics:

- [Navigate to Individual Targets](#)
- [View Cloud Target Alerts](#)
- [View Metrics of the Individual targets](#)
- [Delete a Target Using the UI](#)
- [Manage Exadata Storage Server Metrics and Alert Settings](#)
- [Accessing Oracle Analytics Server Reports](#)

Navigate to Individual Targets


You can drill down to a component target of the Exadata Cloud Service (such as clusters, a database instance, or an Exadata Storage Server).

1. From the **Targets** menu, select **Exadata**.

Enterprise Manager displays the **Oracle Exadata Database Machines and Cloud Services** page showing all the available Exadata targets.

2. From the **Oracle Exadata Database Machines and Cloud Services** page, select the Exadata Cloud Service target whose components you want to view.

Enterprise Manager displays the **Oracle Exadata Cloud Service** page showing an General, Members, and Incidents section for the selected cloud target.

3. From **Oracle Exadata Cloud Service** page, use the left navigation icon  to view the navigation panel. Expand the list of available targets that comprise the Exadata Cloud Service.
4. Click the target to which you want to navigate.

View Cloud Target Alerts

You can view alerts on Exadata Cloud Service and drill down to details about each alert. These alerts may be performance/configuration metrics or hardware faults.

1. From the **Targets** menu, select **Exadata**.

Enterprise Manager displays the **Oracle Exadata Database Machines and Cloud Services** page showing all the available Exadata targets.

2. From the **Oracle Exadata Database Machines and Cloud Services** page, select the Exadata Cloud Service target whose configuration information you want to view.

Enterprise Manager displays the **Oracle Exadata Cloud Service** home page on which you can see all alerts associated with the current cloud target under the **Incidents** section.

View Metrics of the Individual targets

You can view the performance metrics of the dependent targets of Exadata Cloud Service by following these steps:

1. Navigate to the Exadata Cloud Service target home page by selecting the cloud target from the **Targets** page of Enterprise Manager.

Enterprise Manager displays the **Oracle Exadata Cloud Service** home page for the cloud target that you selected.

2. Navigate to the component target whose metrics you want to view. See [Navigate to Individual Targets](#).
3. From the target drop-down menu, choose **Monitoring** and then **All Metrics**.

Delete a Target Using the UI

You can delete the cloud target or any of its member targets by using UI. These steps do not have any effect on the underlying Exadata Cloud Service and do not physically delete the target itself. They help you to remove the target from the Enterprise Manager repository.

If you want to delete the cloud target **and** its member targets using the UI, then first delete the member targets individually using the following steps. After deleting the member targets, delete the cloud target.

1. From the **Targets** menu, select **Exadata**.

Enterprise Manager displays the **Oracle Exadata Database Machines and Cloud Services** page showing all the available Exadata targets.

2. From the **Oracle Exadata Database Machines and Cloud Services** page, select the Exadata Cloud Service target that you want to navigate to.

Enterprise Manager displays the **Oracle Exadata Cloud Service** page with **General**, **Members**, and **Incidents** sections for the selected cloud target.

If you want to delete a member target of the cloud target, then navigate to that target using the navigation panel.

3. From the target drop-down menu, select **Target Setup**, then **Remove Target**.
4. A warning page will display to confirm the target deletion. Click **Yes** to continue.


Manage Exadata Storage Server Metrics and Alert Settings

To access the settings for Exadata Storage Server metrics/alert:

1. Navigate to the Exadata Cloud Service target home page by selecting the cloud target from the **Targets** page of Enterprise Manager.

Enterprise Manager displays the **Oracle Exadata Cloud Service** home page for the cloud target that you selected.

2. Navigate to the Exadata Storage Server target whose metrics you want to view. See [Navigate to Individual Targets](#).
3. From the Exadata Storage Server menu, click **Monitoring**, then **All Metrics** to display all editable metric alert setting.

4. To change a setting, click **Monitoring**, then **Metric and Collection Settings** from the Exadata Storage Server menu. The default View option "Metrics with thresholds" is displayed. You can modify the following parameters:
 - Warning Threshold
 - Collection Schedule - click the link to set a collection schedule.
 - Click the Edit icon  for advanced settings.
5. Click **OK** to save the changes.

Accessing Oracle Analytics Server Reports

Access the Oracle Analytics Server to run the Exadata report *Exadata Cloud Service Capacity Planning Report* to understand the health and utilization of your Exadata portfolio.

1. To access the Exadata reports, you will need a Standalone Oracle Analytics Server. See Standalone Oracle Analytics Server in *Oracle Enterprise Manager Administrator's Guide*.
2. Expand the **Exadata Reports** tree.

The link to **Exadata Cloud Service Capacity Planning Report** is displayed. This is the capacity planning report for Oracle Exadata Cloud.
3. Click the link on the report to view it. In the resulting report page, you can optionally obtain the report in *PDF, HTML, RTF, Excel or PowerPoint* formats.

Exadata Cloud Service Capacity Planning Report

The Capacity Planning Report enables you to analyze and visualize the utilization and health of Exadata resources like storage, CPU, memory, and IO. It also helps to plan for upcoming resource requirements.

Some of the benefits of monitoring your resources:

- *Optimum Resource Usage*: Helps you to fine tune the size and number of resources to ensure that you are getting optimum output
- *Cost Efficiency*: You pay for the exact amount of resources that you will use
- *Predictive Capacity Planning*: Based on the past performance of your configuration, you can plan your resources for future requirements
- *Improved IT Management*: Helps you to keep a close watch on the performance and health parameters of your resources

Configuration

This section reports the name of the Exadata Cloud Service target for which the capacity planning report is generated. It also lists the various hosts and their corresponding host types associated with the configuration. The host type can be **Database Node** or **Storage Cell**.

Global Summary

This section summarizes the key metrics of your configuration for the time ranges like *Last 24 hours*, *Last 7 Days*, *Last 31 Days*, and *Last 365 Days*. The same key metrics are analyzed in greater details in the following sections of the report.

The metrics for *Last 24 Hours* time range are calculated based on the metric collection frequency.

The metrics for the *Last 7 Days* time range are calculated based on the hourly aggregated data.

The metrics for the *Last 31 Days* and *Last 365 Days* time ranges are calculated based on the daily aggregated data.

This is an example global summary of an Exadata Cloud Service:

Time Period	Maximum CPU %	Average CPU %	Maximum Memory %	Average Memory %
Last 24 Hours	33.26	12.34	69.88	67.78
Last 7 Days	56.41	30.03	71.14	69.29
Last 31 Days	56.41	25.88	71.14	69.17
Last 365 Days	56.41	25.88	71.14	69.17

Storage Capacity Planning Details

As an administrator, you can use *Storage Capacity Planning Details* section to diagnose and resolve issues in your storage systems, thus resulting in better performance of not only the storage systems but also the applications relying on them. This section displays the current and historical storage space usage information for all the disk groups.

The current space usage is reported based on the data that was last collected. The following parameters are tabulated for each disk group:

- Redundancy
- Used (%)
- Total Size (GB) - Raw
- Free Space (GB) - Raw
- Used Space (GB) - Raw
- Total Size (GB) - Usable
- Free Space (GB) - Usable
- Used Space (GB) - Usable

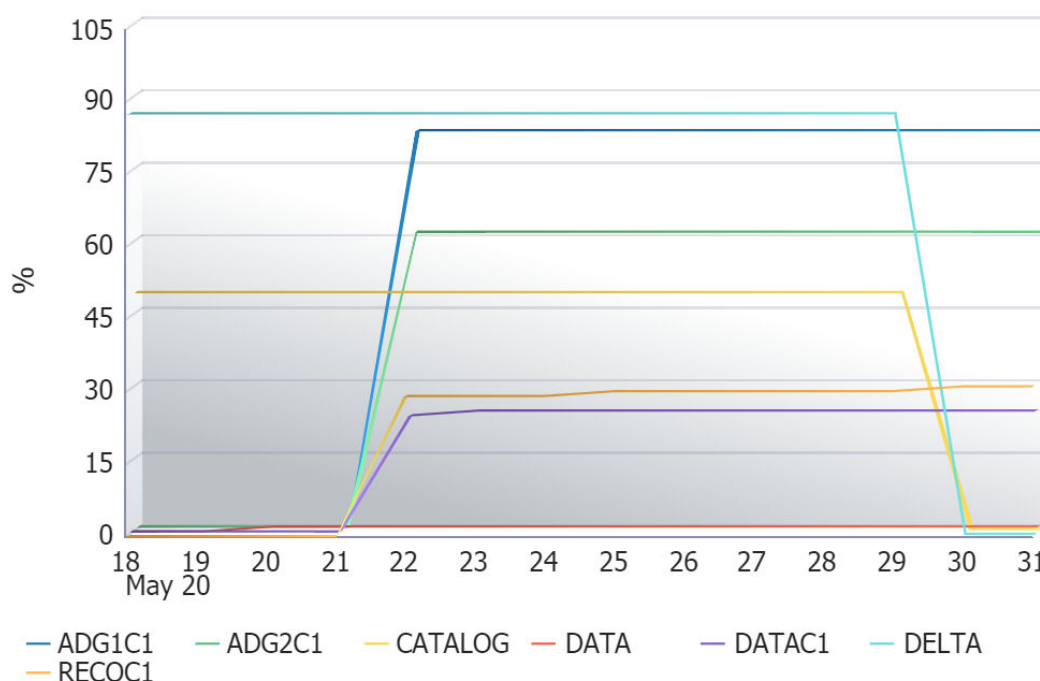
Disk Group	Redundancy	Used (%)	Total Size (GB) - Raw	Free Space (GB) - Raw	Used Space (GB) - Raw	Total Size (GB) - Usable	Free Space (GB) - Usable	Used Space (GB) - Usable
DELTA	NORMAL	87	436824	67379	369445	212345	27623	184722
DATA1	HIGH	25	7344	5573	1771	2380	1790	590
CATALOG	HIGH	49	5760	3112	2648	1813	931	882
RECOC1	HIGH	31	1800	1249	551	583	400	183
ADG1C1	NORMAL	82	1440	285	1155	708	130	578
ADG2C1	NORMAL	61	480	193	287	236	92	144
DATA	EXTERN	1	59	58	1	59	58	1

The historical report is generated based on the data available for the selected time range. The following parameters are tracked and plotted in a line chart along the X-axis against time which is on Y-axis:

- Historical Space Usage By Disk Group (GB) - Usable
- Historical Space Usage By Disk Group (%) - Usable
- Historical Total Space Usage (GB) - Raw
- Historical Total Space Usage (GB) - Usable

- Historical Space Usage by the Top 5 Serviced Databases - As seen daily

Following is an example of the graph for Historical Space Usage by Disk Group (%) - Usable:

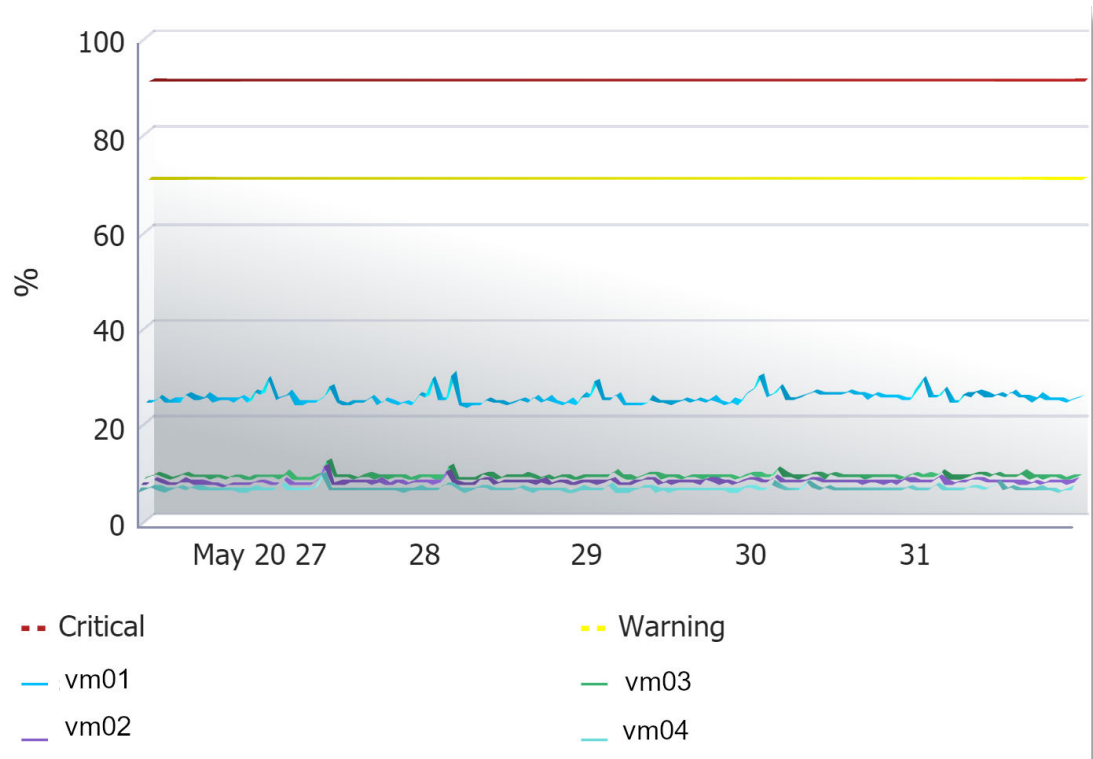


CPU and Memory Capacity Planning Details

Your CPU and Memory utilizations should not be at critical limits during normal load hours. You should determine CPU and Memory utilizations based on your application needs, including CPU and Memory cycles for peak usage. If your CPU and Memory utilizations are optimized at 100% during normal load hours, you have no **capacity** to handle a peak load. A mismatch between the capacity and the demands can result in unsatisfactory utilization of resources.

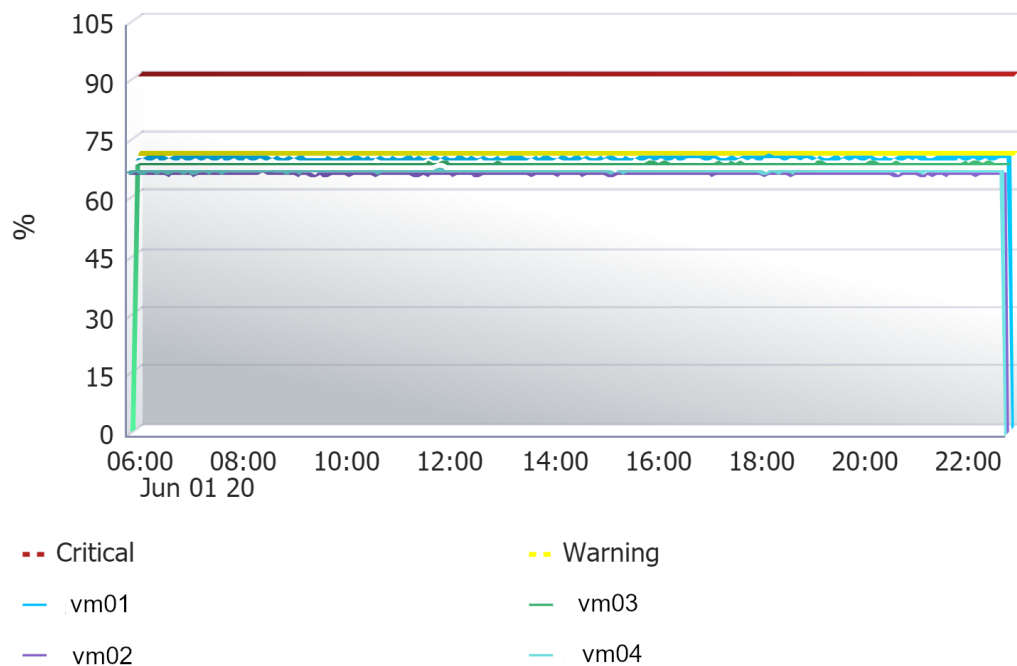
In the *CPU and Memory Capacity Planning Details* section, **CPU Utilization** and **Memory Utilization** parameters of each database node are measured and plotted in line charts for the time ranges *Last 24 Hours*, *Last 7 Days*, *Last 31 Days*, and *Last 365 Days*.

Following is an example of the CPU Utilization metric for the 7 days time range:



Following is an example of the Memory Utilization metric for the 24 hours time range:

Memory Utilization



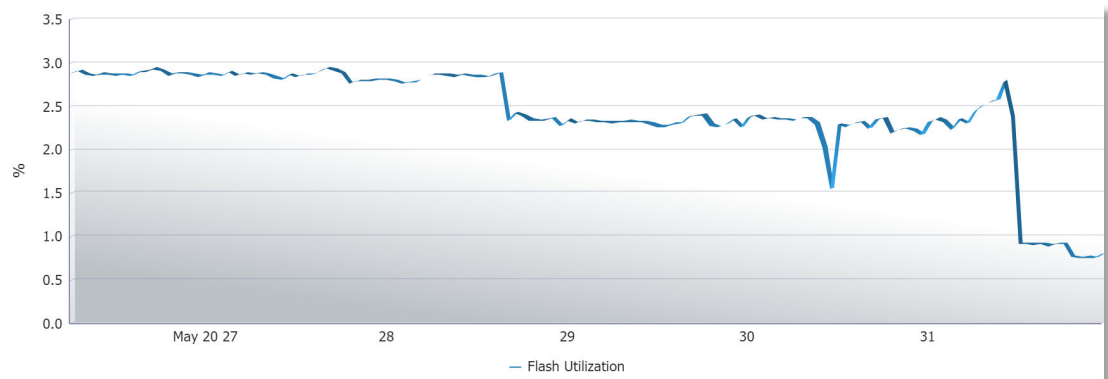
IO Capacity Planning Details

The *IO Capacity Planning Details* section displays the current and historical IO, throughput and disk utilization aggregated across all the storage servers for the selected Exadata Cloud Service target. It summarizes the amount of data being stored in the disks, the read/write throughput rate, and load handling in response to normal and peak demands. Use this report to optimize your disk resources and to plan for future demands.

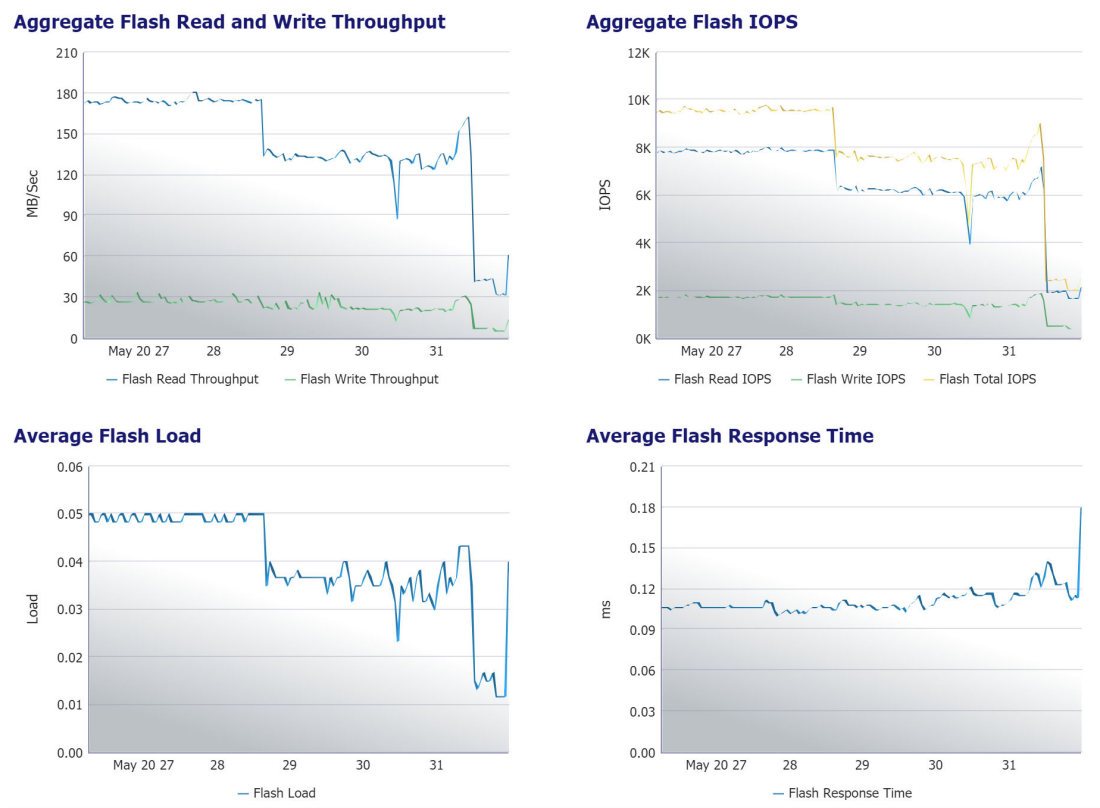
The following parameters of **Cell Disk** and **Flash** are measured and plotted in line charts for the time ranges *Last 24 Hours*, *Last 7 Days*, *Last 31 Days*, and *Last 365 Days*:

- Average Cell Disk Utilization
- Aggregate Cell Disk Read and Write throughput
 - Disk Read Throughput
 - Disk Read Throughput
- Aggregate Cell Disk IOPS
 - Read IOPS
 - Write IOPS
 - Total IOPS
- Average Cell Disk Load
- Average Cell Disk Response Time
- Average Flash Utilization
- Aggregate Flash Read and Write throughput
 - Flash Read Throughput
 - Flash Read Throughput
- Aggregate Flash IOPS
 - Read IOPS
 - Write IOPS
 - Total IOPS
- Average Flash Load
- Average Flash Response Time

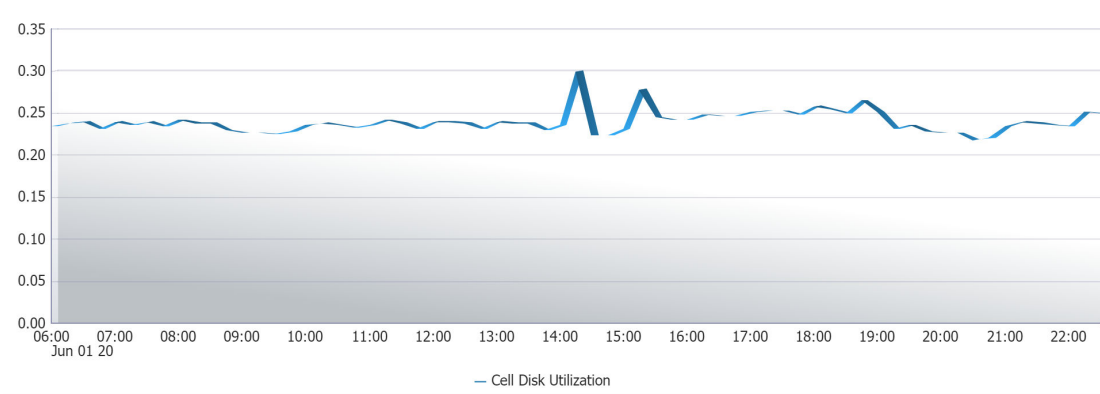
Following is an example of the average utilization of Flash over a 7-day time range:



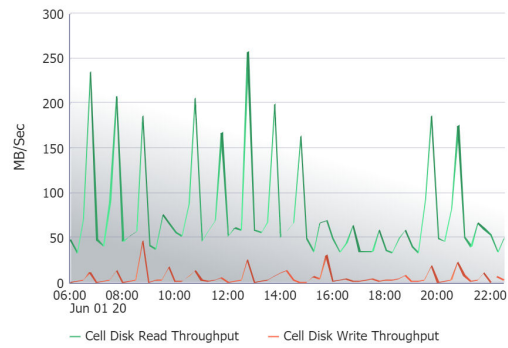
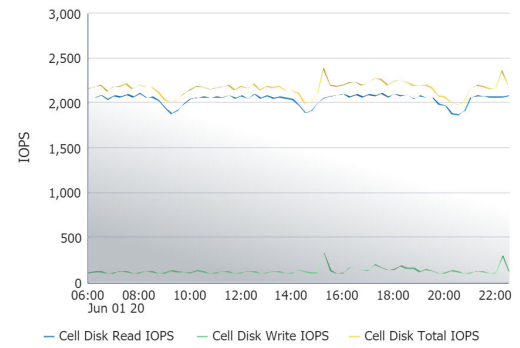
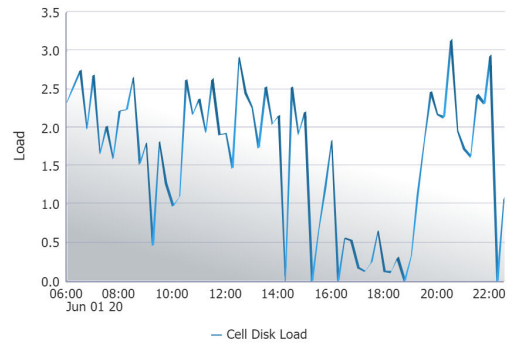
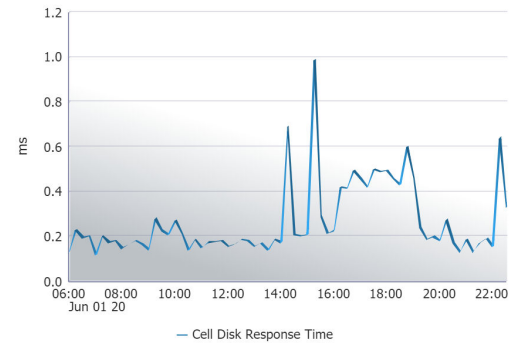
Following is an example of the various IO capacity metrics of Flash over a 7-day time range:



Following is an example of the average utilization of Cell Disk over a 24-hour time range:



Following is an example of the various IO capacity metrics of Cell Disk over a 24-hour time range:

Aggregate Cell Disk Read and Write Throughput**Aggregate Cell Disk IOPS****Average Cell Disk Load****Average Cell Disk Response Time**

7

Monitor Exadata Cloud Service

After the Exadata Cloud Service target and its member targets are discovered, you can use the Oracle Enterprise Manager environment to monitor the targets and obtain insights into their performance.

Topics:

- [Access Exadata Metrics](#)
- [Aggregated Exadata FlashDisk and HardDisk Metric Example](#)
- [Exadata Cell Metric Example](#)
- [Exadata Key Performance Indicators Metrics Examples](#)

Above performance information is currently available for *Oracle Exadata Database Service on Cloud@Customer* targets.

Access Exadata Metrics

For a complete list of metrics available for Oracle Exadata, see *Oracle Exadata in Enterprise Manager Oracle Engineered Systems Metric Reference Manual*.

To access the available Exadata Metrics:

1. Go to Enterprise Manager menu, click **Targets**, and click **Exadata**. All the Exadata targets are listed in this page. Click the link on your cloud target name to open the home page.
The **Oracle Exadata Cloud Service** target home page is displayed.
2. From the Target Navigation tree, expand the **Exadata Grid** and select an Exadata Storage Server. See [View Metrics of the Individual targets](#).
3. On the Exadata Storage Server page, click the **Exadata Storage Server** menu, select **Monitoring**, and select **All Metrics**.
4. On the All Metrics page, a variety of metrics are available. Select a metric or click to expand available metric details.

Aggregated Exadata FlashDisk and HardDisk Metric Example

This metric category contains metrics that are aggregated over either the hard disks or flash disks in a cell. Selecting this metric from the All Metrics page generates a high-level summary, as shown below:

Aggregated Exadata FlashDisk and HardDisk Metric

Collection Schedule Every 15 Minutes [Modify](#)

Upload Interval Every Collection

Last Upload Jan 21, 2015 8:08:53 PM GMT

	CellDisk Type	Average CellDisk IO Load	Average CellDisk Read IOPS	Average CellDisk Read Response Time	Average CellDisk Read Throughput	Average CellDisk Small Read Response Time
▶	FlashDisk	1	8.86	0.4	0.07	0.4
▶	HardDisk	4.33	216.55	59.73	215.64	59.73

☒ Data shown in above table is collected in real time.

Expand the **Aggregated Exadata FlashDisk and HardDisk Metric** in the All Metrics page to show a variety of metric details, such as **Average CellDisk Read Throughput**, which gives an indication of the average number of bytes read from the cell disk, or **Total CellDisk IO Load**, which gives an indication of the total input/output load to the cell disk.

View Data Last 24 Hours Auto Refresh Off

Average CellDisk Read Throughput

CellDisk Type	Average Value	Low Value	High Value	Last Known Value	Current Severity	Alert Triggered
FlashDisk	0.09	0.02	0.36	0.06	Not Applicable	-
HardDisk	190.36	92.54	368.79	130.37	Not Applicable	-

View Data

Last 24 Hours

 Auto Refresh

Off

Total CellDisk IO Load

CellDisk Type	Average Value	Low Value	High Value	Last Known Value	Current Severity	Alert Triggered
FlashDisk	16	16	16	16	Not Applicable	-
HardDisk	51.69	48	93	49	Not Applicable	-

Exadata Cell Metric Example

This metric category contains the performance metrics collected at the cell level for each cell, such as CPU utilization and memory utilization. Selecting this metric from the All Metrics page generates a high-level summary, as shown below:



Expand the **Exadata Cell Metric** in the All Metrics page to show a variety of metric details, such as **CPU Utilization**, which provides information about the CPU utilization, or **Disk I/O Objective**, which provides the optimization objective which IORM is configured to achieve, for example, *Low Latency* or *Balanced* for OLTP-oriented databases, or *High Throughput* for data warehouses.





Exadata Key Performance Indicators Metrics Examples

Access the Key Performance Indicators Metrics for Exadata Storage Servers and Exadata Storage Server Grid to view the metrics of the parameters as listed in the image below for an Exadata Storage Server Grid:

Exadata Grid Exadata Storage Server Grid Page Refreshed May 22, 2020 7:29:29 PM MDT

Oracle Exadata Storage Server Grid: Exadata Grid > All Metrics

All Metrics

Search

View ▼

Exadata Grid

- Exadata Key Performance Indicators**
- Average Flash Disk IO Load
- Average Flash Disk Response Time
- Average Hard Disk IO Load
- Average Hard Disk Response Time
- Flash Disk IO Health Exceptions
- Hard Disk IO Health Exceptions
- Total Flash Disk IOPS
- Total Flash Disk Throughput
- Total Hard Disk IOPS
- Total Hard Disk Throughput

Exadata Key Performance Indicators
Description : Key performance indicators for the Exadata Storage Server Grid.

Collection Schedule Every 5 Minutes

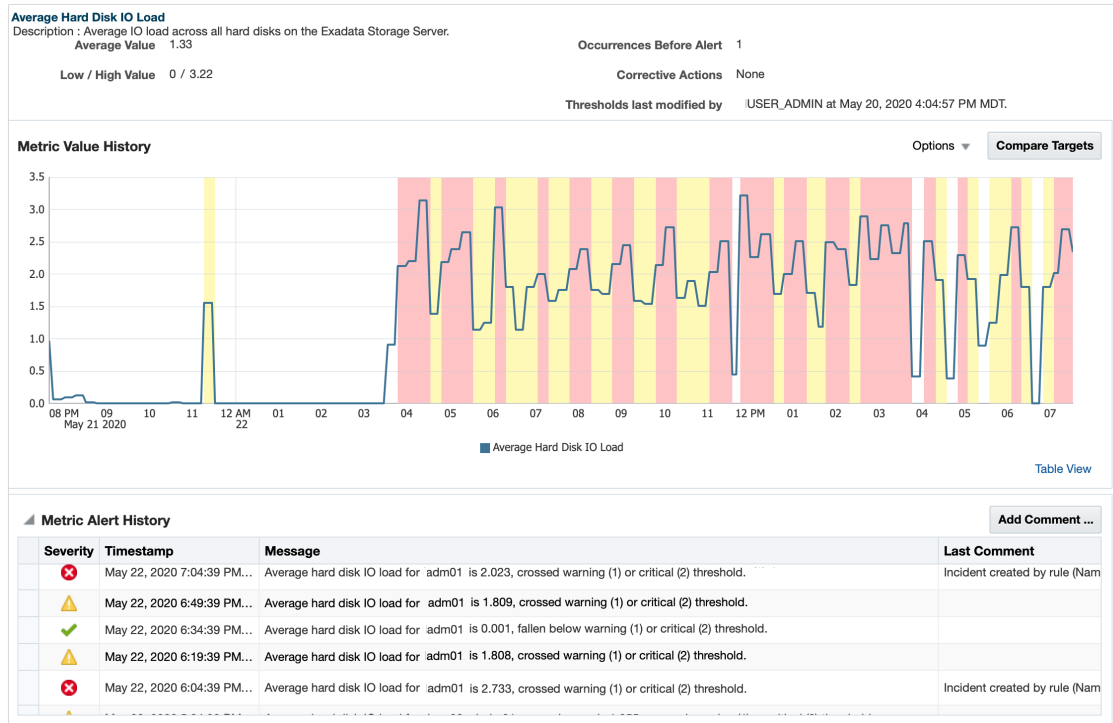
Upload Interval Every Collection

Last Upload May 22, 2020 7:28:41 PM MDT

Metric	Thresholds	Real Time Value
Average Flash Disk IO Load	Not Set	2.25
Average Flash Disk Response Time	Not Set	0.25
Average Hard Disk IO Load	Not Set	2.21
Average Hard Disk Response Time	Not Set	1.87
Flash Disk IO Health Exceptions	Set	0
Hard Disk IO Health Exceptions	Set	1
Total Flash Disk IOPS	Not Set	19.49
Total Flash Disk Throughput	Not Set	45.7
Total Hard Disk IOPS	Not Set	38.73
Total Hard Disk Throughput	Not Set	182.5

☒ Data shown in above table is collected in real time.

The following is an example of the Average Hard Disk IO Load metric of an Exadata Storage Server:



The following is an example of the Flash Disk IO Health Exceptions metric of an Exadata Storage Server:



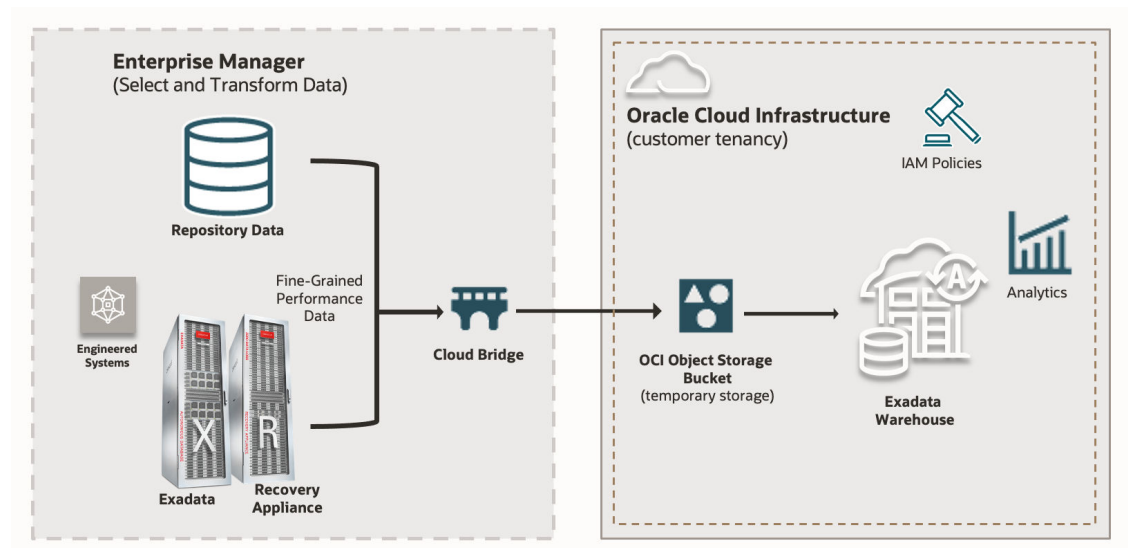
8

Exadata Warehouse

Starting with Oracle Enterprise Manager 13.5 Release Update 15, you can set up Exadata Warehouse through Oracle Enterprise Manager such that it serves as a repository for fine-grained performance metric data collected from *Oracle Exadata Database Service on Cloud@Customer* (ExaCC).

It provides a platform for long-term retention of historical and forecast data for performance and capacity analysis and planning.

The following image illustrates the data flow into the Exadata Warehouse after the configuration is complete:



The fine-grained data from Exadata components like storage servers and compute servers are uploaded to an Autonomous Database. The Exadata Warehouse generates periodic forecasts based on historical data, which are available for queries and export. Older forecast data points are retained for easy comparison.

Since the Exadata Warehouse resides in an Autonomous Database that you specify, the data uploaded to the Exadata Warehouse can be accessed even after data uploads for the Exadata systems have been paused or terminated.

Topics:

- [Setting up Exadata Warehouse and Analyzing Performance Data](#)
- [Simplifying Capacity Planning and Forecasting](#)
- [Exadata Warehouse Sample Queries](#)
- [Configuring Purge Settings for Old Metric Data](#)

Setting up Exadata Warehouse and Analyzing Performance Data

For steps to set up Exadata Warehouse through Oracle Enterprise Manager, see *Exporting and Analyzing Oracle Enterprise Manager Engineered Systems Performance Data Using Exadata Warehouse* ([Tutorial](#)).

Simplifying Capacity Planning and Forecasting

Oracle Cloud Infrastructure (OCI) Operations Insights supports Exadata Warehouse for monitoring Oracle engineered systems like Exadata Database Machine and Zero Data Loss Recovery Appliance (ZDLRA).

Capabilities of Operations Insights' Exadata Warehouse:

- Analyze resource usage for databases running on on-premises Exadata platforms across the enterprise
- Forecast future demand for resources based on historical trends

With the introduction of Exadata Warehouse, resource utilization of the Exadata systems can be analyzed to identify spare capacity for new workloads. This feature is available for Exadata systems that are managed by Oracle Enterprise Manager (EM) 13c Release 5 Update 7 (13.5.0.7) and above. The use of Exadata Warehouse features requires OCI Operations Insights service license subscription.

For more information, see [Operations Insights: Analyze Exadata Resources](#).

Key functionality of fine-grained metric ingestion and forecasting:

- Ingest metrics to Autonomous Databases in customer's tenancy.
- Store fine-grained metric data with per-minute granularity. Offers up to 1,440 data points per metric, if available.
- *Back-fill* metric data to accommodate periods where monitoring of Exadata through Enterprise Manager is interrupted.
- Automated forecasts and roll-ups, updated periodically
- Forecast data using machine learning with seasonality models.
- Forecasts can be generated for up to the next 3 months by analyzing 9 months of historical data, subject to the availability of data.

Capacity Planning:

Exadata Warehouse collects metrics from ASM Cluster, Exadata Storage Server and Exadata VM Guests and provides coverage for the areas below:

- CPU/vCPU and Memory consumption
- Storage Utilization
- IO Capacity, utilization and performance
- Network performance

Default metrics collected by Exadata warehouse:

ASM Cluster	Exadata Storage Server	Exadata VM Guests	Zero Data Loss Recovery Appliance
<ul style="list-style-type: none"> • Disk Group Physical Free (MB) • Disk Group Free (MB) • Disk Group Used % • Used % of Safely Usable • Physical Size (MB) • Size (MB) • Disk Group Usable Free (MB) • Disk Group Usable (MB) 	<ul style="list-style-type: none"> • Large Read/Write Throughput • Small Read/Write Throughput • Large Write Throughput • Small Write Throughput • I/O Load • Large Read Requests • Small Read Requests • Large Write Requests • Small Write Requests • Average Large Read Response Time • Average Small Read Response Time • Average Large Write Response Time • Average Small Write Response Time • I/O Utilization • Flash Cache Size • Flash Cache Used • Read Throughput for Scan • Read Throughput Redirected to Disk for Scan • Read Throughput Redirected to Disk • Read IOPS for Random I/O 	<ul style="list-style-type: none"> • CPU Utilization • Memory Utilization • Network Read Rate • Network Write Rate 	<ul style="list-style-type: none"> • Storage Locations Unused Space (GB) • Storage Locations Size (GB) • Storage Locations Recovery Window Space (GB) • Number of Protected Databases

Applications of Exadata Warehouse:

- Metric data from the Exadata Warehouse can be leveraged to build reports through Oracle Analytics Cloud to visualize historical and forecasted capacity utilization and resource usage.

Exadata Warehouse Sample Queries

Topics:

- [Query 1: Top 10 Exadata Targets Based on Average CPU Usage Over a Chosen Time Period](#)
- [Query 2: Exadata Targets with Under-utilized IOPS](#)

Query 1: Top 10 Exadata Targets Based on Average CPU Usage Over a Chosen Time Period

Use this query to get a list of the top 10 Oracle Exadata Database Machines, based on Average CPU Usage. The query takes the *EM_ID* and historic period in days as parameters and returns a list of up to 10 Oracle Exadata Database Machines sorted in descending order on the basis of CPU utilization over the selected historical period.

Input Parameters:

- EM_ID
- Historical time period in days, for example, 180 for 6 months

```
WITH cpu_usage_historical AS (
  SELECT
    em_id,
    target_name,
    target_type,
    mean_value v
  FROM
    XAWH_DAILY_METRIC_DATA
  WHERE
    em_id = :1
    AND target_type = 'host'
    AND metric_column = 'DS_CPUT'
    AND metric_name = 'xa_analytics'
    AND metric_time > SYSDATE - :2
    AND metric_time < SYSDATE
), targets_with_exadata_name AS (
  SELECT DISTINCT
    em_id,
    target_name,
    target_type,
    CONNECT_BY_ROOT target_name AS exadata_name
  FROM
    XAWH_TARGETS
  WHERE
    target_type = 'host'
  CONNECT BY PRIOR target_guid = parent_guid
    AND em_id = :1 START WITH ( parent_guid IS NULL )
)
SELECT
  t.exadata_name "Target",
  TRUNC(AVG(cpu_usage_historical.v),2) "Average CPU Usage"
FROM
  cpu_usage_historical,
  targets_with_exadata_name t
WHERE
  cpu_usage_historical.em_id = t.em_id
  AND cpu_usage_historical.target_name = t.target_name
  AND cpu_usage_historical.target_type = t.target_type
GROUP BY
  t.exadata_name
ORDER BY
```

AVG(cpu_usage_historical.v) DESC
FETCH FIRST 10 ROWS ONLY;

Sample Output:

Input parameters provided:

- EM_ID: 19102682514EE3FAAE6DC8AFB460EAC8
- Historical period in days: 365

Output:

Target	Average CPU Usage
DB Machine sample1.example.com	29.08
DB Machine sample2.example.com	24.48
DB Machine sample3.example.com	14.52
DB Machine sample4.example.com	7.57
DB Machine sample5.example.com	3.68

Query 2: Exadata Targets with Under-utilized IOPS

Use this query to get a list of the top 10 under-utilized Oracle Exadata Database Machines, based on IOPS utilization computed for the following metrics:

1. Number of requests to read small blocks per second from a cell disk (CD_IO_RQ_R_SM_SEC)
2. Number of requests to read large blocks per second from a cell disk (CD_IO_RQ_R_LG_SEC)
3. Number of requests to write small blocks per second to a cell disk (CD_IO_RQ_W_SM_SEC)
4. Number of requests to write large blocks per second to a cell disk (CD_IO_RQ_W_LG_SEC)

The query takes the EM_ID and historic period in days as parameters and returns a list of up to 10 Oracle Exadata Database Machines sorted in ascending order on the basis of least IOPS usage over the selected historical period.

Input Parameters:

- EM_ID
- Historical time period in days, for example, 365 for 1 year

```
WITH iops_usage_historical AS (
  SELECT
    em_id,
    target_name,
    target_type,
    mean_value v,
    CASE
      WHEN KEY_PART_1 LIKE 'CD%' THEN
        'CD'
      ELSE
        'FD'
    END disk_type
  FROM
    XAWH_DAILY_METRIC_DATA
  WHERE
    em_id = :1
```

```

AND target_type = 'oracle_exadata'
AND metric_column IN ( 'CD_IO_RQ_R_SM_SEC', 'CD_IO_RQ_R_LG_SEC', 'CD_IO_RQ_W_SM_SEC',
'CD_IO_RQ_W_LG_SEC' )
AND metric_name = 'xa_analytics'
AND metric_time > SYSDATE - :2
AND metric_time < SYSDATE
), targets_with_exadata_name AS (
SELECT DISTINCT
  em_id,
  target_name,
  target_type,
  CONNECT_BY_ROOT target_name AS exadata_name
FROM
  XAWH_TARGETS
WHERE
  target_type = 'oracle_exadata'
CONNECT BY PRIOR target_guid = parent_guid
AND em_id = :1 START WITH ( parent_guid IS NULL )
)
SELECT
  t.exadata_name "Target",
  t.target_name "Oracle Storage Server",
  disk_type "Disk Type",
  SUM(v) "Total IOPS"
FROM
  iops_usage_historical,
  targets_with_exadata_name t
WHERE
  iops_usage_historical.em_id = t.em_id
AND iops_usage_historical.target_name = t.target_name
AND iops_usage_historical.target_type = t.target_type
GROUP BY
  t.exadata_name,
  t.target_name,
  disk_type
ORDER BY
  AVG(iops_usage_historical.v) ASC
FETCH FIRST 10 ROWS ONLY;

```

Sample Output:

Input parameters provided:

- EM_ID: 19102682514EE3FAAE6DC8AFB460EAC8
- Historical period in days: 365

Output:

Target	Oracle Storage Server	Disk Type	Total IOPS
DB Machine sample4.example.com	sample4adm09.example. com	FD	14368.85
DB Machine sample4.example.com	sample4adm11.example. com	FD	14605.37
DB Machine sample4.example.com	sample4adm10.example. com	FD	15722.89

Target	Oracle Storage Server	Disk Type	Total IOPS
DB Machine sample3.example.com	sample3adm02.example.com	CD	8364.61
DB Machine sample3.example.com	sample3adm01.example.com	CD	9614.26
DB Machine sample1.example.com	sample1adm03.example.com	CD	26910.9
DB Machine sample3.example.com	sample3adm02.example.com	FD	27305.82
DB Machine sample2.example.com	sample2adm06.example.com	FD	222255.41
DB Machine sample2.example.com	sample2adm04.example.com	FD	222294.83
DB Machine sample2.example.com	sample2adm05.example.com	FD	222339.38

Configuring Purge Settings for Old Metric Data

The volume of fine-grained data and rolled-up statistics stored in the Exadata Warehouse grows over time. You can configure settings to periodically delete old data from the Exadata Warehouse.

Data retention settings:

Data Type	Default Value (days)	Minimum Value (days)
Fine-grained metric data	395 (13 months)	30 (1 month)
Hourly roll-up metric data	1460 (4 years)	395 (13 months)
Daily roll-up metric data	2555 (7 years)	395 (13 months)
Diagnostic data	395 (13 months)	60 (2 months)

Customize Data Retention Settings for Exadata Warehouse

To override the default retention periods and to view the latest retention settings for the Exadata Warehouse tables, use the EM Job type **Exadata Warehouse Auto Purge Settings**. In the job, you can specify the value for the retention period for fine-grained data, hourly roll-up data, daily roll-up data, and diagnostic data. To create a new Job of the type:

1. In Oracle Enterprise Manager, click **Enterprise**, then **Job**, then click **Activity**.
2. On the Job Activity page, select the job type **Exadata Warehouse Auto Purge Settings** from the Create Job menu and click **Go**.
3. In the **General** tab, provide a name for the job and add the individual targets or one composite target such as a Group.
4. In the **Parameters** tab, provide value for the retention period for fine-grained data, hourly roll-up data, daily roll-up data, and diagnostic data, all specified in days.
5. In the **Credentials** tab, select an appropriate option for credentials.
6. In the **Schedule** tab, schedule the job.
7. In the **Access** tab, define or modify the access you want other users to have to this job.
8. Click **Submit**.

Auto-Purge Job Details

After the auto-purge job is created to override the default purge settings, you can view the job or modify the auto-purge settings:

Job

Create 'Exadata Warehouse Auto Purge Settings' Job

CancelSubmit

GeneralParametersCredentialsScheduleAccess

Retention period for fine-grained data (days)

Default setting is 395 days (13 months) and minimum value supported is 30 days (1 month)

Retention period for hourly rollup data (days)

Default setting is 1460 days (4 years) and minimum value supported is 395 days (13 months)

Retention period for daily rollup data (days)

Default setting is 2555 days (7 years) and minimum value supported is 395 days (13 months)

Retention period for diagnostic data (days)

Default setting is 395 days (13 months) and minimum value supported is 60 days (2 months)

Target Properties

Target properties can be used in parameters
Property names are case-sensitive. To escape '%', use '%%'.

Name	Description
%emd_root%	location of Agent
%perlbin%	location of Perl binary used by Agent
%TargetName%	target name
%TargetType%	target type

CancelSubmit

If the job is run without any change to the parameters, the current retention settings are displayed in the job output.

Part V

Troubleshooting

For troubleshooting issues related to Oracle Exadata Database Machine, see Troubleshooting the Exadata Plug-in in *Oracle Exadata Database Machine Getting Started Guide*.

Topics:

- [Internal Server Error in the Exadata Cloud discovery wizard](#)
- [List of members of the cloud target does not show Database Instance](#)
- [List of members of the cloud target does not show Storage Servers](#)
- [Cell Admin command fails for RestAPI monitored cells](#)
- [Discovery of Exadata Storage Server failed when \\$HOME/.exacli/cookiejar is present](#)

Internal Server Error in the Exadata Cloud discovery wizard

This error usually occurs in the step where you select the Monitoring Agent - OCI Credential Name, and the Home Region in the wizard.

Check the agent logs *gcagent.log* and look for entries in the format `http://<service>.<region_identifier>.<realm_domain>`.

The Exadata cloud discovery flow makes use of three OCI services, *Database Service API*, *Identity and Access Management API* and *Search Service API*. For example, `https://identity.us-ashburn.oraclecloud.com`, `https://database.us-ashburn.oraclecloud.com`, and `https://query.us-ashburn.oraclecloud.com`.

1. Verify that the OCI API service host names are reachable from agent host.
If they are not reachable, then verify the correctness of the URL. The OCI regions can belong to three different realms - commercial realm, Oracle Government Cloud realms, dedicated realms. The API URL formats are different for each realm. If incorrect URL is used, then contact *Oracle Support* to get the required patch.
2. If proxy is used, then see [Set Up a Proxy Connection](#).
 - a. Ensure that the following properties are set in `$EMSTATE/config/emd.properties` on the agent:

```
oci_http_proxy_port=<proxy port>
oci_http_proxy_host=<proxy url>
```

Optionally, set the properties `oci_proxy_username` and `oci_proxy_password` if authenticated proxy is used.

If connection using proxy host fails, then try using the proxy host IP address instead to rule out DNS related issues.

- b. Edit the file `$EMDROOT/work/agentStateDir/sysman/config/s_jvm_options.opt` and add the following line at end of the file and save the file:

```
-Djdk.http.auth.tunneling.disabledSchemes=""
```

3. Verify the OCI credentials by running the following OCI CLI command on the agent host:

```
oci db vm-cluster list --compartment-id <compartment id> --exadata-infrastructure-id <exadata infrastructure id>
```

Ensure that the profile information is present in `~/.oci/config` file before running the OCI CLI command. For example:

```
[DEFAULT]
user=<user ocid>
fingerprint=<finger print>
tenancy=<tenancy ocid>
key_file=<pem key file path>
region=<region code for example: us-ashburn-1>
```

4. Verify that the SSL certificate sent by OCI Service and the one received at the EM agent side are from trusted source - Oracle Corporation / Digicert. Run the following command on the agent host:

```
[oracle@<server>]$ openssl s_client -connect database.us-ashburn-1.oraclecloud.com:443
CONNECTED(00000003)
depth=2 C = US, O = DigiCert Inc, OU = www.digicert.com, CN = DigiCert Global Root G2
verify return:1
depth=1 C = US, O = DigiCert Inc, CN = DigiCert Global G2 TLS RSA SHA256 2020 CA1
verify return:1
---
Certificate chain
0 s:/C=US/ST=California/L=Redwood City/O=Oracle Corporation/CN=*.us-ashburn-1.oc-test.com
i:/C=US/O=DigiCert Inc/CN=DigiCert Global G2 TLS RSA SHA256 2020 CA1
1 s:/C=US/O=DigiCert Inc/CN=DigiCert Global G2 TLS RSA SHA256 2020 CA1
i:/C=US/O=DigiCert Inc/OU
```

If the certificate is not from a trusted Certification Authority, for example, a self-signed certificate, and the proxy or firewall is replacing it with another certificate, then import the proxy or firewall certificate into the agent's JVM trust store (not keystore):

- a. Copy the certificate *crt* file on the agent host.
- b. Change to the right directory:

```
cd $EMDROOT/oracle_common/jdk/jre/lib/security/
```


- c. Import the certificate:

```
keytool -importcert -keystore ./cacerts -file <certificate file> -alias <alias_name> -storepass <password>
```

- d. Restart the agent:

```
emctl stop agent; emctl start agent
```

List of members of the cloud target does not show Database Instance

From cloud target home page, use the left navigation icon  to view the navigation panel. Expand the list of available targets in the navigation panel, locate your database instance, and check if it belongs to a cluster database.

If the database belongs to a cluster database, then this is an expected behavior. It is not associated to the cloud target explicitly as a system member.

List of members of the cloud target does not show Storage Servers

Discovery is usually skipped in one of the following scenarios:

- Credentials are not supplied during the discovery.
- IP's of the storage servers are typically listed in the file `/etc/oracle/cell/network-config/cellip.ora`. Discovery will be skipped if the file is not present on any of the hosts.
- Credentials are incorrect. See [Create Storage Server Credentials](#).
- `exacli` is not installed on any of the hosts.

To resolve the issue, try one of the following:

- Ensure that proper named credentials are created and provided in the property file.
- Ensure that `exacli` is installed on each of the hosts listed in the property file that was used for discovering the cloud target.

Cell Admin command fails for RestAPI monitored cells

This error occurs when the agent is running with the user that is different from the user who installed the agent. For RestAPI monitoring mechanism, the HOST PDP settings are mandatory.

To resolve the issue, set the privilege delegation for the user that is running the agent. To set the Host PDP settings, click **Setup Menu**, click **Security**, and select **Privilege Delegation**.

Discovery of Exadata Storage Server failed when `$HOME/.exacli/cookiejar` is present

The discovery of the Exadata Storage Server fails when the cookie-jar is present at `$HOME/.exacli/cookiejar`. Note that cookie-jar is not supported for Enterprise Manager monitoring of Exadata Storage Server.

To resolve the issue, move cookie-jar from `$HOME/.exacli/cookiejar` and attempt the discovery of the Exadata Storage Server.