

Oracle® Enterprise Manager

Oracle Exadata Database Machine Getting Started Guide



13c Release 5
F37174-17
November 2025



Copyright © 2015, 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
Conventions	i

1 Introduction to the Plug-in

Oracle Exadata Plug-in Features	1
Monitoring and Notification Features	1
Hardware Support Features	2
Target Discovery Features	3
Exadata Storage Server Grid Home Page and Server Home Page Features	3
Exadata Performance Page Features	4
Exadata Metrics Features	4
Oracle Exadata Database Machine Supported Hardware and Software	4
Exadata Database Machine Configuration Support	4
Hypervisor Support	5
Multi-Rack Support	5
Partitioned Support	5
IPv6 Support	5
Oracle SuperCluster Support	5
Supported Component Versions	6
Supported Operating Systems	7
Oracle Exadata Database Machine Hardware Not Supported	7

2 Prerequisites

Create a Database Server ILOM Service Processor User	1
Create an ExaCLI or RESTful API User	2
Create an Exascale Monitoring User	3
Create SNMPv3 Users	4
Create SNMPv3 Users on Compute Nodes and Storage Servers	4
Create Individual SNMPv3 Users on Exadata 19.3 and Above	5
Create All SNMPv3 Users on Exadata 19.2. and Below	5

Create SNMPv3 Users on Cisco Ethernet Switches	5
Create SNMP Community on Cisco Ethernet Switches	6
Create an SNMPv3 User in IB Switches	6
Create SNMP Community for InfiniBand Switch Targets	7
Enable SNMPv3 on PDUs	8
Create an SNMPv3 user on PDUs	8
Create SNMP Community on PDUs	8
Verify Software Versions	9
Exadata Storage Server Software	9
InfiniBand Switch	9
Cisco Ethernet Switch	10
PDU Firmware	10
Grid Infrastructure/DB Cluster	10
Verify Names Resolution	10
Verify Firewall Configuration	11
User Roles	13
Install Oracle Management Agent	13
Manually Deploy Exadata and Related Plug-ins	13
Prerequisites for Discovery Using EMCLI	14
Create Named Credentials for All the Target Types	14
Named Credentials and Commands to Create Them	15
Set Up SSH Key Authentication	25
Generate the SSH Key	25
Add SSH Public Key to Each Target	26
Create the Input File	27
Configuration Parameters	29
Parameters for Adding or Skipping Components	31
Parameters for Adding or Skipping Virtual Platform Components	32
Optional Verification of the Prerequisites	33
Prerequisites for Discovery and Monitoring Using Non-root User	34
Non-root User for Systems Infrastructure Remote Host Targets	34
Non-root User for Virtual Platform Targets	34

3 Exadata Database Machine Discovery

Discovering an Oracle Exadata Database Machine	1
Discover the Oracle Exadata Database Machine Target Using the Console	1
Fresh Discovery of Oracle Exadata Database Machine Target	2
Fresh Discovery of Exadata Database Machine on SuperCluster Systems	9
Convert Database Machine Targets with 12c Target Types to 13c Using the Console	11
Discover the Oracle Exadata Database Machine Target Using EMCLI	13
Submit the Deployment Procedure for the Discovery	13

View the Discovery Status	14
Convert 12c Type Database Machine Targets to 13c Using EMCLI	15
Discovering Grid Infrastructure and RAC	17
Switch from Using CellCLI to RESTful API	17
Update the Monitoring Configuration Using Console	17
Update the Monitoring Configuration Using EMCLI for Switching from CellCLI to RESTful API	18

4 Post-Discovery Configuration and Verification

Verify SNMP Subscription of the Oracle Management Agent to Exadata Storage Server	1
Using the ALTER CELL Command	2
Check Current SNMP Configuration	3
Configure SNMP Values Using dcli (optional)	4
Verify SSH Connectivity Using cellcli	4
Remove a Subscription	4
Verify SNMP Configuration for InfiniBand Switch Targets	5
Verify the Cisco Ethernet Switch SNMP Configuration	6
Verify the PDU SNMP Configuration	6
Convert SNMP V2 to V3 at Scale for All Hardware Components in Exadata	6

5 Oracle Exadata Database Machine Administration

Creating Roles to Manage the Plug-in	1
Update Monitoring User of an Already Discovered SI Target	2
Exadata Database Machine Topology	3
Drilling Down to Individual Targets of Database Machine	4
Viewing Virtual Exadata Targets	4
Viewing Critical Hardware Information for the Database Machine	5
Viewing Exadata Resource Utilization	6
Compute Resources	6
Storage Resources	8
Viewing Database Machine Alerts	12
Viewing Metrics	12
Adding Exadata Database Machine Components Manually	13
Removing an Exadata Database Machine Target	13
Removing a Component of a Database Machine Target	15
Updating the Exadata Database Machine Schematic Diagram	15
Drop a Component from the Exadata Database Machine Schematic Diagram	15
Add Components to the Exadata Database Machine Schematic Diagram	15
Provisioning Virtualized Exadata Clusters	16
Creating a Database Cluster	16

Scaling Up a Database Cluster	22
Scaling Down a Database Cluster	23
Deleting a Database Cluster	24
Exadata Storage Server Metrics and Alert Settings	24
Exadata Storage Server Management	25
About Exadata Storage Server	25
Viewing an Exadata Storage Server Topology	26
Performing Administration Tasks on Exadata Storage Servers	26
Managing the I/O Resource	27
Add/Update IORM Configuration	28
The Inter-Database Plan	30
Diagnosing Exadata Storage Server Alerts	31
Managing the InfiniBand Fabric	31
InfiniBand/RoCE Switch Metrics	31
Performing Administration Tasks on InfiniBand Networks	32
Setting Up Alerts	33
Discovering and Monitoring Exascale	33
Discovering Exascale	34
Monitoring Exascale	36
Flash Cache Resource Monitoring	40
Exadata Database Machine Fault Monitoring	41
Monitoring Exadata Database Machine Components	41
Exadata Storage Servers	41
Cisco/RoCE Switch (Systems Infrastructure Switch)	41
ILOM Targets (Systems Infrastructure Server)	42
Monitoring Systems Infrastructure ILOM Server Target Using Non-root User	42
Power Distribution Units (PDUs)	42
InfiniBand Switch (Systems Infrastructure Switch)	42
Automatic Indexing for Database Nodes	43
Evaluating Compliance of Exadata Targets	44
Reporting	44
Accessing Oracle Support Workbench for Exadata Storage Server	45
Accessing Oracle Analytics Server Reports	45
Exadata Capacity Planning Details	46
Exadata vCPU Consumption Report	50
Oracle Exadata Database Machine Dashboard Creation	53
Exadata Database Machine Dashboard Creation	53
Find All Database Machine Reports	54
How to Make the Report "Public"	54

6 Exadata Warehouse

Setting Up Exadata Warehouse and Analyzing Performance Data	1
Simplifying Capacity Planning and Forecasting	2
Exadata Warehouse Sample Queries	3
Configuring Purge Settings for Old Metric Data	7

7 Exadata Metrics

Accessing Exadata Metrics	1
Example Metrics	1
Aggregated Exadata FlashDisk and HardDisk Metric	1
Exadata Storage Server Metric	2
Exadata Key Performance Indicators Metrics	4

8 Troubleshooting

Discovery Troubleshooting	1
Hardware Availability	2
Discovery Failure Diagnosis	2
Exadata Storage Server is not Discovered	2
Compute Node or InfiniBand Switch is not Discovered	3
Extra or Missing Components in the Newly Discovered Exadata Database Machine	3
ILOM, PDU, or Cisco Switch is not Discovered	3
Test Connection Fails in PDU	4
Target Does not Appear in the Components Step of Guided Discovery	4
Target Status is Down or Metric Collection Error After Discovery	4
ILOM Credential Validation Fails During Discovery	5
Discovery Process Hangs	5
SNMP Configuration Missing	6
Set Up SNMP Subscription for Exadata Storage Servers	6
Set Up SNMP Subscription for Cisco Ethernet Switch Targets	7
Set Up SNMP for Power Distribution Unit (PDU) Targets	8
Time Out Error During SNMP Subscription	8
Review Logs	9
EMCLI Discovery Issues	9
System Infrastructure Targets in Pending Status	10
Discovery Status Summary States "Completed with errors"	10
Target Discovered but Association Not Created	10
Component not Available or Accessible Before Discovery	11
Trouble Using ExaCLI during Discovery	11
Input File Errors	11

Unclear Message Shown on the Deployment Procedure Output	12
Troubleshooting the Exadata Database Machine Schematic File	12
Exadata Database Machine Management Troubleshooting	13
Oracle Database Storage Server System Target Status Pending	13
Target Status Issues	13
Metric Collection Issues	13
Status: Pending Issues	14
Cells sys Targets	14
Database Machine Target or Any Associated Components	14
Monitoring Agent Not Deployed for IPv6 Environments	14
Monitoring: Error Fetching Data for Grid and CellSys Targets IORM Page Charts	15
Not able to receive SNMP traps on Exadata Storage Servers using IPv6	15
Exascale Discovery and Monitoring Issues	16

Index

Preface

Oracle Enterprise Manager provides the tools to effectively and efficiently manage your Oracle Exadata Database Machine. With a consolidated view of the Exadata Database Machine, Oracle Enterprise Manager provides a consolidated view of all the hardware components and their physical location with indications of status. Oracle Enterprise Manager also provides a software view of the databases residing on it and their resource consumption on compute nodes and Exadata Storage Servers.

This document provides the installation and administration instructions to set up Oracle Enterprise Manager to monitor and manage your Oracle Exadata Database Machine.

Audience

This system monitoring installation guide is for users who perform administrative and problem-resolution tasks on Oracle Exadata Database Machine.

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 that 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.

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.
<code>monospace</code>	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Introduction to the Plug-in

This chapter provides a general overview of the Oracle Exadata plug-in, including supported hardware and software. The following topics are discussed:

- [Oracle Exadata Plug-in Features](#)
- [Oracle Exadata Database Machine Supported Hardware and Software](#)

Oracle Exadata Plug-in Features

Highlights of the Oracle Exadata plug-in include the following features:

- [Monitoring and Notification Features](#)
- [Hardware Support Features](#)
- [Target Discovery Features](#)
- [Exadata Storage Server Grid Home Page and Server Home Page Features](#)
- [Exadata Performance Page Features](#)
- [Exadata Metrics Features](#)

For advanced management and monitoring capabilities, enable the Exadata management Pack. See Introduction to the Engineered Systems Packs in *Advanced Management and Monitoring for Engineered Systems*.

Monitoring and Notification Features

With the Oracle Exadata plug-in, and related Systems Infrastructure and Virtual Infrastructure plug-ins, you can monitor Exadata Database Machine targets through Enterprise Manager Cloud Control 13c. These plug-ins provide seamless integration with supported Exadata hardware and software so that you can receive notifications on any Exadata Database Machine target. Features include:

- Monitoring of the Exadata Database Machine as an Enterprise Manager target.
- Monitoring of the Exadata Database Machine target, including the Exadata Storage Server's I/O Resource Management feature within Enterprise Manager.
- Support SNMP notification for Exadata Database Machine components.
- Support dashboard report creation from Enterprise Manager Cloud Control, including a simplified configuration of the service dashboard.
- Support of client network host names for compute nodes.
- Enhanced InfiniBand network fault detection and InfiniBand schematic port state reporting.
- Modification of Enterprise Manager monitoring agents as needed for all Exadata Database Machine components.
- IORM for multi-tenancy database (CDB/PDB) environment:
 - CDB-level I/O Workload Summary with PDB-level details breakdown.

- I/O Resource Management for Oracle Database 12c and above.
- Exadata Database Machine-level physical visualization of I/O Utilization for CDB and PDB on each Exadata Storage Server.
- Integration link to Database Resource Management UI.
- Support discovery of locked down Exadata Storage Servers.
- Since the release of EM 13c, the Exadata Plug-in is able to leverage the enhanced hardware monitoring features offered by the newly introduced *Oracle Systems Infrastructure Plug-in*. Users can view a photo-realistic schematic of the various hardware (including rack) and monitor the faults on individual hardware components.

Hardware Support Features

You can use the Oracle Exadata plug-in to optimize the performance of a wide variety of Exadata Database Machine targets, including:

- Oracle SuperCluster, including:
 - Configurations:
 - * LDOM: Control domain, IO/guest domain
 - * Zone: Global, non-global
 - Discover, monitor, and manage Exadata Database Machine-related components residing on SuperCluster Engineered System
 - See [Oracle SuperCluster Support](#) for more details.
- Multi-Rack support:
 - Supports discovery use cases: Initial discovery, add a rack
 - Side-by-side rack schematic
- Support for Storage Expansion Rack hardware.
- Full partition support:
 - Logical splitting of an Exadata Database Machine Rack into multiple Database Machines.
 - Each partition is defined through a single OneCommand deployment.
 - Compute nodes are not shared between partitions.
 - Multiple partitions connected through the same InfiniBand network.
 - Compute nodes in same partition share the same Cluster.
 - Ability to specify a customized DBM name during discovery of the target.
 - User can confirm and select individual components for each DBM.
 - Flexibility to select none, some, or all of the InfiniBand switch as part of monitored network, including the ability to add switches post discovery.
 - Flexibility to select some or all of the Exadata Storage Servers to be shared among Exadata Database Machines.
- Support for Exadata Extended (XT) Storage: Since XT storage servers do not have flash drives, Oracle Enterprise Manager Cloud Control will automatically suppress all the flash-related monitoring functionalities, and present only those IO Resource Management settings that are appropriate for XT storage. Exadata X8 systems configured with XT storage are automatically recognized by Oracle Enterprise Manager Cloud Control during

discovery, and the corresponding Exadata Storage Server targets automatically adapt to the underlying storage server type.

- Support for the increasing types of Exadata Database Machine targets. See [Oracle Exadata Database Machine Supported Hardware and Software](#) for a complete list of supported hardware.
- InfiniBand Switch Sensor fault detection, including power supply unit sensors and fan presence sensors.

Target Discovery Features

The target discovery process is streamlined and simplified with the Oracle Exadata plug-in. Features include:

- Automatically push the Exadata plug-in to agent during discovery.
- Discovery prerequisite checks updates, including:
 - Check for critical configuration requirements.
 - Check to ensure either `databasemachine.xml` or `catalog.xml` files exist and are readable.
 - Prevent discovered targets from being rediscovered.
- Credential validation and named credential support.
- Ability to apply a custom name to the Exadata Database Machine target.
- Support discovery using the client access network.
- Automate SNMP notification setup for Exadata Database Machine components.
- Support discovery of compute nodes with client network host names.
- Enterprise Manager Cloud Control Exadata Discovery Wizard lets you discover Exadata Database Machine targets using 13c.
- An existing Exadata Database Machine target with 12c target types can be converted to 13c target types. For more information, see [Convert Database Machine Targets with 12c Target Types to 13c Using the Console](#) for more information.

 **Note**

Exadata Database Machine targets are configured with OOB default thresholds for the metrics. No additional template is provided by Oracle.

Exadata Storage Server Grid Home Page and Server Home Page Features

The Exadata Storage Server Grid home page and Server home page provides the following features:

- Provides a fine-grained performance summary for flash and hard disk.
- Provides usage statistics to highlight flash cache and Smart Scan efficiency.
- Provides a detailed summary of flash space usage.
- Provides metrics for:
 - I/O time distribution by flash and hard disk.

- IORM wait per database.

Exadata Performance Page Features

The Performance page provides the following features:

- Side-by-side comparison of flash and hard drive performance.
- Performance comparison between multiple Exadata Storage Servers.
- Performance utilization for flash and hard drive to identify workload reaching hardware limit.
- Provides Exadata Storage Server performance charts to help with diagnosing performance issues when I/O reaching hardware limits.

Exadata Metrics Features

Metrics reports are critical to manage your Oracle Exadata Database Machine effectively. With the metrics, you can determine where additional resources are needed, when peak usage times occur, and so forth.

- Enhanced metric scalability in large environment to reduce time out by reducing calls.
- Reduce metric collection error for the Exadata HCA metric. Improvements to combine the HCA port data collection in a single call to reduce chances of time out.

Oracle Exadata Database Machine Supported Hardware and Software

The following sections describe the supported hardware and software by the Oracle Exadata plug-in:

- [Exadata Database Machine Configuration Support](#)
- [Oracle SuperCluster Support](#)
- [Supported Component Versions](#)
- [Supported Operating Systems](#)
- [Oracle Exadata Database Machine Hardware Not Supported](#)

Exadata Database Machine Configuration Support

Enterprise Manager Cloud Control 13c is supported on the following Exadata Database Machine configurations:

- [Hypervisor Support](#)
- [Multi-Rack Support](#)
- [Partitioned Support](#)
- [IPv6 Support](#)

Hypervisor Support

Exadata Database Machine Monitoring and Management supports both Xen and KVM Hypervisors when used in certified Exadata configurations.

Multi-Rack Support

Enterprise Manager supports managing multiple connected racks of Oracle Database Machine of the supported machine types listed above ([Hypervisor Support](#)). Also, the following two racks can be monitored in a multi-rack as these cannot exist as a standalone single rack:

- Storage Expansion Rack
- Compute Node Expansion Rack

Partitioned Support

The following partitioned configurations are supported:

- Partitioned Exadata Database Machine - the logical splitting of a Database Machine Rack into multiple Database Machines. The partitioned Exadata Database Machine configuration must meet the following conditions to be fully supported by Enterprise Manager Cloud Control 13c:
 - Each partition is defined through a single OneCommand deployment.
 - Exadata Storage Servers and compute nodes are not shared between partitions.
 - Multiple partitions are connected through the same InfiniBand network.
 - Compute nodes in same partition share the same Cluster.

The expected behavior of a partitioned Exadata Database Machine includes:

- The target names for the Exadata Database Machine, Exadata Grid, and InfiniBand Network will be generated automatically during discovery (for example, Database Machine dbm1.example.com, Database Machine dbm1.example.com_2, Database Machine dbm1.example.com_3, etc.). **However, users can change these target names at the last step of discovery.**
- All Infiniband Switches need to be selected as part of the Exadata Database Machine targets for every partition. Infiniband Switches will not be added automatically to subsequent Exadata Database Machine targets of other partitions. The KVM, PDU, and Cisco switches can be individually selected for the Database Machine target of each partition.
- User can confirm and select individual components for each Database Machine.

IPv6 Support

Enterprise Manager supports monitoring and managing Exadata environments configured with IPv6. For more details, see IPv6 support status on Exadata Database Machine ([Doc ID. 2056895.1](#)), available in My Oracle Support.

Oracle SuperCluster Support

Only Oracle SuperCluster with software Version 1.1 with DB Domain on Control LDOM-only environments are supported. Earlier versions of Oracle SuperCluster can be made compatible

if you update to the October 2012 QMU release. You can confirm this requirement by looking at the version of the `compmon pkg` installed on the system (using either `pkg info compmon` or `pkg list compmon` commands to check). You must have the following minimum version of `compmon` installed:

```
pkg://exa-family/system/platform/exadata/compmon@0.5.11,5.11-0.1.0.11:20120726T024158Z
```

The following **hardware configurations** are supported:

- Oracle SuperCluster:
 - T4-4
 - T5-8
 - M6-32
 - M7
 - M8
 - M12-2S

Note: Enterprise Manager Systems Infrastructure plug-in doesn't support firmware monitoring through eXtended System Control Facility (XSCF).

The following **software configurations** are supported:

- LDOM
 - Control Domain
 - IO/Guest Domain
- Zone
 - Global
 - Non-Global

The following **software versions** are supported:

- Oracle SuperCluster V1.1 and above
- Oracle SuperCluster V1.0.1 + October QMU

Supported Component Versions

For information on supported component versions of Exadata, refer to the following notes on My Oracle Support:

- [Exadata System Software and Hardware Versions Supported by Oracle Enterprise Manager Plug-ins \(Doc ID 1626579.1\)](#). This note provides information on the required Oracle Enterprise Manager and plug-in versions to support the specific versions of Exadata Software and generations of Exadata Hardware.
- [Required Enterprise Manager versions for Exadata Cisco Catalyst switches \(Doc ID 2449150.1\)](#)
- [Exadata Database Machine and Exadata Storage Server Supported Versions \(Doc ID 888828.1\)](#)

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)

Oracle Exadata Database Machine Hardware Not Supported

The following Oracle Exadata Database Machine hardware configurations are **not** supported for Enterprise Manager Cloud Control Exadata plug-in 13.x:

- V1 hardware
- V2 hardware

 **Note**

V2 machines discovered in Enterprise Manager Cloud Control 12c are still supported in 13c. However, discovery of V2 machines in Enterprise Manager Cloud Control 13c is not supported.

Prerequisites

This chapter provides the prerequisites for Exadata Database Machine discovery. The following topics are discussed:

- [Create a Database Server ILOM Service Processor User](#)
- [Create an ExaCLI or RESTful API User](#)
- [Create an Exascale Monitoring User](#)
- [Create SNMPv3 Users](#)
- [Verify Software Versions](#)
- [Verify Names Resolution](#)
- [Verify Firewall Configuration](#)
- [User Roles](#)
- [Install Oracle Management Agent](#)
- [Manually Deploy Exadata and Related Plug-ins](#)
- [Prerequisites for Discovery Using EMCLI](#)
- [Prerequisites for Discovery and Monitoring Using Non-root User](#)

Create a Database Server ILOM Service Processor User

For the Enterprise Manager Agent to communicate with an ILOM service processor, there must be a specific user ID established on the ILOM service processor.

 **Note**

Adding the specific user ID requires administrator level privilege on the ILOM service processor.

The specific ILOM user ID can be added in the ILOM service processor web interface, ILOM CLI, or with the `ipmitool` command. This example uses ILOM CLI.

For security reasons, the password to the ILOM service processor root user ID does not appear in the ILOM CLI commands in this example.

1. Log in to the Service Processor as root:

```
# ssh root@[Service Processor IP]
Password:
```

2. Change to the `users` directory:

```
# cd /SP/users
```

3. Create the `oemuser` user and password:

```
# create oemuser

Creating user...
Enter new password: ****
Enter new password again: ****

Created /SP/users/oemuser
```

4. Change to the new user's directory and set the role:

```
# cd oemuser
/SP/users/oemuser

set role='cro'
Set 'role' to 'cro'
```

5. Test the ILOM user ID created in step 3 by listing the last 10 system events:

- For Exadata X2 through X4:

```
# ipmitool -I lan -H <ilom_hostname> -U oemuser -P <oemuser password> -L USER
sel list last
10\
```

- For Exadata X5 and above (requires the -I lanplus command option):

```
# ipmitool -I lanplus -H <ilom_hostname> -U oemuser -P <oemuser password> -L
USER sel list last 10
```

6. Repeat steps 1 through 5 for the rest of the compute node ILOM service processors in your Oracle Database Machine.

Create an ExaCLI or RESTful API User

Enterprise Manager can monitor Exadata Storage Servers using `cellcli`, `exacli`, or RESTful API. In order to monitor and/or manage Exadata Storage Servers through `exacli` or the RESTful API, a user must be created and provided for Exadata Database Machine discovery. For more information on `exacli` or RESTful API user creation, see the Cell node sections in both Example 9-9 and Example 9-10 in [Creating Users for Use with ExaCLI](#) in *Oracle Exadata Database Machine Maintenance Guide*. These examples provide the steps to create both the `celladministrator` user and the `cellmon` user. Review the following best practice guidelines/requirements to determine if the `celladministrator` user is required:

- As a best practice of using the least privileged user for an operation, always create and use the `cellmon` user for the monitoring credentials.
- Create the `celladministrator` user only if one or both of these activities are required:
 - Set up of SNMP subscriptions for the Exadata Storage Servers
 - IORM administration
- During discovery, for the Exadata Storage Server credentials, the following are the best practice recommendations/requirements:
 - Supply the `cellmon` user credentials if selecting `exacli` or RESTful API for monitoring
 - If setting up SNMP subscriptions, supply the `celladministrator` credentials for the SNMP setup (the `celladministrator` credential is required)

Note

If SNMP subscriptions are not configured, then Enterprise Manager will not monitor key events such as hardware failures.

- While making IORM change through Enterprise Manager, use the `celladministrator` credential.

Create an Exascale Monitoring User

Starting with Exadata System Software version 24.1, Enterprise Manager offers the capability to monitor Oracle Exadata Exascale. To monitor Exascale, create a user on the cluster with appropriate privileges. The user authentication uses public-private key pairs.

For the purpose of this example and throughout this document, the monitoring user will be `esmonitor`. See [Oracle Exadata User-Specific Administration](#) in *Exascale User's Guide*.

For use with Exascale, the public and private keys must be in PEM format. To create a key pair, use standard utilities such as `openssl` or the `ESCLI` utility on the Exadata compute node:

1. Create a directory for your monitoring key pair:

```
mkdir /home/oracle/eskey
```

2. Protect this directory from access by other users:

```
chmod 700 /home/oracle/eskey
```

3. Generate your key using `ESCLI`:

```
/opt/oracle/dbserver/dbms/bin/escli
```

```
mkkey --private-key-file /home/oracle/eskey/esmonitor-priv.pem --public-key-file /home/oracle/eskey/esmonitor-pub.pem
```

4. Protect your key from unauthorized access:

```
chmod 600 /home/oracle/eskey/esmonitor-priv.pem /home/oracle/eskey/esmonitor-pub.pem
```

5. Provide your public key to the Exadata Exascale administrator to provision your monitoring user.

6. As the Exadata Exascale administrator, create the monitoring user, grant monitoring privileges and configure with the user's public key:

```
escli --wallet <admin wallet location> ----ctrl <ExaCTRL server"port>
```

```
mkuser esmonitor --id esmonitor
```

7. Grant privileges to monitor the cluster and all vaults in the cluster:

```
chuser esmonitor --privilege cl_monitor
```

```
chuser esmonitor --privilege +vlt_read
```

8. Associate the public key with the user:

```
chuser esmonitor --public-key-file1 esmonitor-pub.pem
```

Create SNMPv3 Users

If SNMPv3 will be used for Oracle Exadata Database Machine monitoring, ensure that the necessary SNMPv3 users are created on the components prior to discovering the Exadata to fully leverage monitoring through Enterprise Manager.

- [Create SNMPv3 Users on Compute Nodes and Storage Servers](#)
- [Create SNMPv3 Users on Cisco Ethernet Switches](#)
- [Create an SNMPv3 User in IB Switches](#)
- [Enable SNMPv3 on PDUs](#)
- [Create an SNMPv3 user on PDUs](#)

 **Note**

Oracle does not recommend to use MD5 as the AuthProtocol while creating SNMPv3 user for security reasons.

Create SNMPv3 Users on Compute Nodes and Storage Servers

The commands required to create SNMPv3 users on compute nodes and storage servers are similar, but make use of different command line interfaces and object names. In the following examples, run `cellcli` to get to the interactive prompt if on a storage server, and run `dbmcli` to get to the interactive prompt on a compute node. Specify the appropriate object name, `cell` and `dbserver` for compute node. The instructions differ between Exadata System Software releases. Please see the details for the respective Exadata System Software version in the sections below.

 **Note**

For additional information on these commands, please see the following references:

- For information on CellCLI and the `alter cell` command, see [Using the CellCLI Utility](#) in the *Oracle Exadata System Software User's Guide*.
- For information on DBMCLI and the `alter dbserver` command, see [Using the DBMCLI Utility](#) in the *Oracle Exadata Database Machine Maintenance Guide*.

Create Individual SNMPv3 Users on Exadata 19.3 and Above

Exadata System Software version 19.3 and above supports maintaining SNMPv3 users individually. Use the following command to create an SNMPv3 user.

```
CLI> alter <cell|dbserver> snmpuser.<username>
=(authprotocol=SHA,authpassword=<password>,privprotocol=AES,privpassword=<password>)
```

In the above command,

- **name:** The user name
- * is the only allowed value for the password
- **authProtocol:** The authentication protocol. The supported protocol is **SHA**. The system prompts for the authentication password. The authentication password must have 8 to 12 alphanumeric characters.
- **privProtocol:** Encryption protocol. The options supported are **none**, **AES**, and **DES**. The default value is **none** when **privProtocol** attribute is not specified. The system prompts for an encryption password if the encryption protocol is specified. The password must be exactly 8 alphanumeric characters long, and is case sensitive.

To change the password for an existing user:

```
CLI> alter <cell|dbserver> snmpuser.<username>
=(authprotocol=SHA,authpassword=<newpassword>)
```

To delete a user:

```
CLI> alter <cell|dbserver> snmpuser.<username> = ''
```

Create All SNMPv3 Users on Exadata 19.2. and Below

Exadata System Software versions 19.2 and below require maintenance of SNMPv3 users as a complete set. Use the following command to create an SNMPv3 user.

Note

Be sure to include the details of all SNMPv3 users while creating a new user as the set will be replaced with this command.

```
CLI> alter <cell|dbserver> snmpUser=( (name=<username>, authProtocol=SHA,
authPassword=<password>, privProtocol=DES, privPassword=<password>)
[,<repeat_with_details_as_necessary_for_additional_users>])
```

Create SNMPv3 Users on Cisco Ethernet Switches

The below commands configure an SNMP user with authentication and privacy parameters on the Cisco Ethernet switches running NX-OS including the admin switch, and if this is a Exadata X8M, the RoCE switches.

Starting with Enterprise Manager 13.5 Release Update 13, while configuring SNMP, Enterprise Manager does not support AUTH protocol DES and PRIV protocol MD5 in Cisco Ethernet switches running NX-OS.

The passphrase can be any case-sensitive, alphanumeric string up to 64 characters.

```
switch# configure terminal
switch(config)#
switch(config)# snmp-server user <username> auth sha <passphrase> priv
aes-128 passphrase
```

The below command displays information about one or more SNMP users.

```
switch(config)# show snmp user
```

Once all the configuration changes are done, the below command will save the configuration in persistent memory .

```
switch(config)# copy running-config startup-config
```

Create SNMP Community on Cisco Ethernet Switches

The below commands configure an SNMP community on the Cisco Ethernet switches running NX-OS including the admin switch, and if this is a Exadata X8M, the RoCE switches.

```
switch# configure terminal
switch(config)#
switch(config)# snmp-server community <community> ro
```

The below command displays information about one or more SNMP communities.

```
switch(config)# show snmp community
```

Ensure that the ACL is not assigned to the community string.

Once all the configuration changes are done, the below command will save the configuration in persistent memory .

```
switch(config)# copy running-config startup-config
```

Create an SNMPv3 User in IB Switches

If the Exadata to be discovered is an IB Exadata, follow the below commands to create an SNMPv3 user on the IB switches.

Log in to the ILOM CLI of the switch as *ilom-admin* user.

Create the SNMPv3 user account by using the following `create` command:

```
> create /SP/services/snmp/users/<v3user> privacyprotocol=<privacv3_protocol>
privacypassword=<privacv3_password>
```

```
authenticationprotocol=<authentication_protocol>
authenticationpassword=<authentication_password>
```

For example,

```
> create /SP/services/snmp/users/v3adminuser privacyprotocol=AES
privacypassword=mypassword1 authenticationprotocol=SHA
authenticationpassword=mypassword2
```

The user is created as part of the discovery flow of the switch. This is required only if you want to explicitly add an SNMPv3 user.

To edit the SNMPv3 user account, run the following `set` command:

```
> set /SP/services/snmp/users/<v3user> privacyprotocol=<privacv_protocol>
privacypassword=<privacy_password>
authenticationprotocol=<authentication_protocol>
authenticationpassword=<authentication_password>
```

Run the following command to display information about one or more SNMP users:

```
> show -d properties /SP/services/snmp/users/<username>
```

Create SNMP Community for InfiniBand Switch Targets

 **Note**

To configure the SNMP community for an InfiniBand Switch:

1. Log in to the InfiniBand Switch ILOM web interface using the URL `https://<ib_switch_hostname>` as root.

 **Note**

Try using Internet Explorer if the console does not display all fields/values in your browser of choice.

2. Click **ILOM Administration**, then **Management Access**, and finally **SNMP**.

3. Ensure the following values are set:

```
State=Enabled
Port=161
Protocols=v1,v2c,v3
```

If you need to make changes, make sure you click **Save**.

4. Click **ILOM Administration**, then **Management Access**, then **SNMP**, and finally the **Communities** link. The page scrolls to the SNMP Communities table. In the SNMP Communities table, click **Add**.

A new window opens.

5. Type the community name into the **Community Name** field. Select the permissions from the Permission drop-down menu. Click **Save**.

The community is created, and the window closes.

6. Click **Alert Management**.

7. If not already listed, for each Agent that monitors the InfiniBand Switch target, select an empty alert (one that has the Destination Summary 0.0.0.0, snmp v1, community 'public') and click **Edit**. Provide the following values:

```
Level = Minor
Type = SNMP Trap
Address = [agent compute node hostname]
Destination Port = [agent port]
SNMP Version = v1
Community Name = public
```

Click **Save**.

Enable SNMPv3 on PDUs

Follow the below steps to enable SNMPv3 on the PDUs.

1. Access the PDU metering unit from a system on the network.
2. Click on the **Net Configuration** link and log in as an admin user.
3. Select the **SNMP-Access** tab.
4. Click the **SNMP v3 Enable** checkbox to enable SNMP v3.
5. Click **Submit**.

Create an SNMPv3 user on PDUs

Follow the below steps to create an SNMPv3 user on the PDUs.

1. Access the PDU metering unit from a system on the network.
2. Click on the **Net Configuration** link and log in as an admin user.
3. Select the **SNMP-Access** tab.
4. In the SNMPv3 table, perform the following
 - a. Enter the SNMPv3 UserName.
 - b. Select the **Security Level** auth / priv.
 - c. Select **SHA** as the Auth Algorithm.
 - d. Enter the **Auth Password**.
 - e. Select **AES** as the **Privacy Algorithm**.
 - f. Enter the **Privacy Password**.
 - g. Select **Enable**.
5. Click **Submit**.

Create SNMP Community on PDUs

Follow the below steps to create an SNMP community on the PDUs.

1. Access the PDU metering unit from a system on the network.
2. Click on the **Net Configuration** link and log in as an admin user.
3. Select the **SNMP-Access** tab.
4. In the NMS (SNMPv1/v2) table, perform the following
 - a. Enter the SNMP IP address or host name of the first monitoring agent.
 - b. Select the privacy level for read-write community as **private** or **public**.
 - c. Select the privacy level for read-only community as **private** or **public**.
 - d. Select the type of **access** from the menu.
 - e. Select **Enable**.
5. Click **Submit**.

Verify Software Versions

Verify the following software versions:

- [Exadata Storage Server Software](#)
- [InfiniBand Switch](#)
- [Cisco Ethernet Switch](#)
- [PDU Firmware](#)
- [Grid Infrastructure/DB Cluster](#)

Exadata Storage Server Software

See [Oracle Exadata Database Machine Supported Hardware and Software](#) for specific supported Exadata Software releases. To verify the software version on the Exadata Storage Server, ssh to the Exadata Storage Server as the `root`, `celladmin`, or `cellmonitor` user. Run:

```
# cellcli -e 'list cell detail'
```

Look for `releaseVersion` in the output.

InfiniBand Switch

To verify the version of the InfiniBand switch firmware in your environment:

1. Log on to the management interface for the InfiniBand Switch (using SSH).
2. Run the following command:

```
# nm2version
```

The output should be similar to this:

```
# nm2version
Sun DCS 36p version: 2.2.13-2
```

This example shows a supported configuration for deploying the plug-in to monitor.

3. If the `nm2version` command returns output similar to this:

```
# nm2version
NM2-36p version: 1.0.1-1
```

Then you must upgrade your InfiniBand switch firmware. Follow the instructions listed in [My Oracle Support \(MOS\) Document 888828.1](#)

Cisco Ethernet Switch

Use the `show version` command to list the BIOS version, NXOS version, and the model details.

Log on to the management interface for the Cisco Switch (using SSH). Run the following command:

```
#show version
```

Here's an example output to the command:

```
Software
BIOS: version 05.38
NXOS: version 7.0(3)I7(8)
BIOS compile time: 06/12/2019
NXOS image file is: bootflash:///nxos.7.0.3.I7.8.bin
NXOS compile time: 3/3/2020 20:00:00 [03/04/2020 04:49:49]

Hardware
cisco Nexus9000 C9336C-FX2 Chassis
Intel(R) Xeon(R) CPU D-1526 @ 1.80GHz with 24571608 kB of memory.
Processor Board ID FDO230701AK
```

PDU Firmware

The PDU firmware version must be 2.10 or later. The current version can be obtained by logging into the web interface of the PDU. On the left side of the screen, click **Module Info** to view the PDU firmware version.

Software updates for the PDU are available at:

https://updates.oracle.com/Orion/PatchDetails/process_form?patch_num=12871297

Grid Infrastructure/DB Cluster

Grid Infrastructure/DB Cluster is required to be up and running before discovery.

Verify Names Resolution

The Enterprise Manager OMS server(s) require direct network access to each of the compute nodes. If the names of the compute nodes are not registered in the OMS nodes' DNS, then they will have to be manually entered in the `/etc/hosts` file for each OMS.

Each compute node should be verified to be able to resolve the hostnames of the ILOM servers, PDU's, Exadata Storage Servers, and InfiniBand and Cisco switches. Again, if the names of those components are not registered in DNS, then entries can be added to the `/etc/hosts` file of each compute node.

To manage the Exadata Database Machine components from Enterprise Manager Cloud Control 13c, it is necessary for your local machine to be able to resolve the host name of Cloud Control 13c.

To access any of the Exadata Database Machine components directly from your local machine, it is also necessary for your local machine to be able to resolve the names of those components.

Verify Firewall Configuration

To verify the firewall configuration:

1. Enable ping

In many secure network environments, it is normal for the `ping` service to be disabled. Enterprise Manager uses `ping` to establish the basic availability and status of the Exadata Database Machine components.

- The compute nodes need to have the `ping` service and port enabled from the OMS Server(s).
- All other Exadata Database Machine components (ILOM servers, PDU's, Exadata Storage Servers, InfiniBand switches, and Cisco switch) need to have the `ping` service and port enabled from the compute nodes (where the agents are running).

ⓘ Note

The `ping` traffic overhead is minimal. The agent pings the targets every five minutes.

2. Open Database Ports

The database listener ports must be opened for the Enterprise Manager OMS server(s). Note that Exadata Database Machine databases will use SCAN listeners; so, ports will need to be opened for the base compute node, the compute node virtual IP, and scan listeners addresses.

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

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

3. Open Enterprise Manager Upload Port

The Enterprise Manager Cloud Control 13c agents require access to the 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 will need to be opened for each of the compute nodes.

4. Open Agent Ports

The OMS server(s) will need to be able to connect to the Enterprise Manager Cloud Control 13c Agent HTTP/HTTPS port on each compute node. 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 compute node 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`

5. Open SSH Ports (port 22)

The Enterprise Manager Cloud Control 13c Agents require `ssh` access to the Exadata Database Machine components they monitor. As the Agents will run on the compute nodes the `ssh` ports, 22, on each of the Exadata Storage Servers, ILOMs, PDUs, InfiniBand switches, and Cisco switch will need to be opened for each of the compute nodes.

6. Allow UDP Traffic (SNMP Ports) (Port 162)

All Exadata Database Machine components need to be able to send SNMP traps to the Agents running on the compute nodes. SNMP uses the UDP protocol so the Agent port and port 162 need to be opened for UDP traffic between the Exadata Storage Servers, ILOMs, InfiniBand Switches, Cisco Switch, and the Agent.

Table 2-1 Firewall Ports

Component	Ping service and port	SNMP*	SSH (port 22)	Notes
PDU	From remote agent	Yes	Yes	
Compute node ILOM	From remote agent	Yes	Yes	<p>1. Remote agent needs to be able to SSH to dom0.</p> <p>2. Need SNMP port open on dom0.</p>
dom0	From remote agent	Yes	Yes	
Exadata Storage Server	From remote agent	Yes	Yes	
InfiniBand Switch	From remote agent	Yes	Yes	
Cisco Switch	From remote agent	Yes	Yes	
OMS			Yes	Upload http/https port - usually 3872
Agent				The OMS server(s) will need to be able to connect to the Enterprise Manager Cloud Control Agent HTTP/HTTPS port on each compute node. The Agent port defaults to 3872. If port 3872 is not available, the next available port starting from port 1830 is used.

User Roles

To manage the Exadata Database Machine in Enterprise Manager, you should consider creating Enterprise Manager roles and assigning them to the appropriate Enterprise Manager users. Creating these roles restricts the privileges that each user has, for example in deleting the plug-in or accessing reports. See [Oracle Exadata Database Machine Administration](#).

Install Oracle Management Agent

Enterprise Manager Exadata discovery supports the use of either management network hostname or client network hostname for the compute nodes. When installing the Oracle Management Agent on the compute nodes, you should use the same hostname as used in Oracle Clusterware.

You can identify the hostname of the nodes in the cluster by running the `olsnodes` command on one of the compute nodes. It is recommended that a fully qualified hostname, including the domain name, be used when specifying an Oracle Management Agent hostname.

Oracle Management Agents need to be installed on each compute node and must not be installed on any other Exadata Database Machine components. For physical Exadata, the agents should be installed on each compute node. For virtual Exadata, the agents should be installed on each domU (virtual machine), and not on the dom0 (hypervisor).

For information on installing agents, refer to [Installing Oracle Management Agents in Cloud Control Basic Installation Guide](#).

Manually Deploy Exadata and Related Plug-ins

Exadata and Systems Infrastructure plug-ins are the default plug-ins and are available for use. The Virtual Infrastructure plug-in must be manually deployed in OMS as well as agent before starting the discovery process for Virtual Exadata.

You may need to manually deploy the Exadata and related plug-ins to the agents on each of the compute nodes when upgrading an existing agent plug-in installation. Deploy the Exadata, Systems Infrastructure, and for virtual Exadata the Virtual Infrastructure plug-ins manually if an older version of the plug-in(s) has been deployed to the agent already and you would like to upgrade to the latest version of the plug-in(s) deployed on the OMS.

To determine if the Exadata, Systems Infrastructure, and Virtual Infrastructure plug-ins are deployed on each compute node and what versions they are, you have two options:

- From a terminal window, run the following command:

```
emctl listplugins agent
```

Note

The `emctl listplugins agent` command must be run on the compute node using the `emctl` in the agent installation directory.

- From Enterprise Manager Cloud Control, click the **Setup** menu (upper right corner), **Extensibility**, and then **Plug-ins**.

To manually deploy the Exadata and related plug-ins:

1. From the Enterprise Manager home page, click the **Setup** menu (upper right corner), **Extensibility**, and then **Plug-ins**.
2. On the *Plug-ins* page, select the plug-in of choice from the Name list.

① Note

Check *Exadata Storage Software Versions Supported by the Oracle Enterprise Manager Exadata Plug-in* (Doc ID 1626579.1) on My Oracle Support for the latest supported plug-ins. Oracle recommends that you deploy the latest version of the Exadata and related plug-ins to the agent.

3. To deploy the plug-in on Management Servers first, follow the instructions in [Downloading, Deploying, and Upgrading Plug-Ins](#) in *Cloud Control Administrator's Guide*.
4. To deploy the plug-in on the Agent, click **Deploy On**, then **Management Agent**.
5. On the *Deploy Plug-in on Management Agent* pop-up, click **Add**. A search pop-up will appear to allow you to search for targets to add. In the Target Type drop-down, select **Agent**, then click **Search**.
Select a Target name from the results list and click **Select**. Repeat for each agent target.
6. After you have added the agents, click **Next** on the *Deploy Plug-in on Management Agent* screen to review and verify the agent information.
7. Click **Deploy** to deploy the plug-in to the agents.
8. Once the plug-in has been deployed to all agents, a confirmation screen will appear. Click **OK** to dismiss the pop-up or **Show Status** to display the status of the agent in the Enterprise Manager *Deployment Activities* screen.

Prerequisites for Discovery Using EMCLI

The EMCLI based database machine discovery is performed through a deployment procedure which consumes an input file that contains all the required information like monitoring agents, configuration, components, and credentials. OMS and the monitoring agents should be *Enterprise Manager 13.4 Release Update 4* or above.

Complete the following prerequisites tasks before you can discover the Exadata Database Machine targets using EMCLI:

- [Create Named Credentials for All the Target Types](#)
- [Set Up SSH Key Authentication](#)
- [Create the Input File](#)
- [Optional Verification of the Prerequisites](#)

Create Named Credentials for All the Target Types

The value for the credential parameter in the input file must be a named credential available in Enterprise Manager. The named credentials can be created using either EMCLI or by using the Enterprise Manager UI.

To create a named credential from the UI, navigate to **Setup > Security > Named Credential**.

To create a named credential using EMCLI, issue the command specified in the table in [Named Credentials and Commands to Create Them](#). Named credential in the input file should

follow the notation <CREDENTIAL_OWNER>:<NAMED_CREDENTIAL>. For example, credMap.schematicFileHostCred=SYSMAN:SCHEMATIC_HOST_CRED.

Named Credentials and Commands to Create Them

The following table lists the named credentials, and provides the sample command that you can run to create it:

Property Name and Description	Sample Command
credMap.schematicFileHostCred Credential of host where schematic file is present. This will be used to read schematic file and credential user should have read permission for schematic file. It's mandatory to place the schematic file in the discovery agent host. Target Type: host	Credential Type: HostCreds (username, password) <pre>emcli create_named_credential -cred_name='SCHEMATIC_HOST_CRED' -auth_target_type='host' -cred_type='HostCreds' -cred_scope='global' -target_type='host' - attributes="HostUserName:sample;HostPassword: XXXXXXXXX"</pre>
	Credential Type: HostSSHCreds (public, private key) <pre>emcli create_named_credential -cred_name='SCHEMATIC_HOST_SSH_CRED' -cred_scope=global -cred_type=HostSSHCreds - auth_target_type=host - attributes="USERNAME:XX;SSH_PVT_KEY:pvt;SSH_P UB_KEY:pub" -input_file="pvt:<PATH_TO_PRV_ FILE_EMCLI_INSTALL_BOX>; pub:<PATH_TO_PUB_FILE_EMCLI_INSTALL_BOX >"</pre>

The private and public key file path is on the host where EMCLI is executed.

Property Name and Description	Sample Command
credMap.agentCred	Credential Type: HostCreds (username, password)
These were the credential used to install the agents. These credentials are required to set up passwordless SSH between the agent and Exadata Storage Server.	<pre>emcli create_named_credential -cred_name='AGENT_HOST_CRED' -auth_target_type='host' -cred_type='HostCreds' -cred_scope='global' -target_type='host' - attributes="HostUserName:sample;HostPassword: XXXXXXXXX"</pre>
Target Type: host / oracle_emd	Credential Type: HostSSHCreds (public, private key)
	<pre>emcli create_named_credential -cred_name='AGENT_HOST_SSHCRED' -cred_scope=global -cred_type=HostSSHCreds -auth_target_type=host - attributes="USERNAME:XXXX;SSH_PVT_KEY:pvt;SSH _PUB_KEY:pub" -input_file="pvt:<PATH_TO_PRV_ FILE_EMCLI_INSTALL_BOX>; pub:<PATH_TO_PUB_FILE_EMCLI_INSTALL_BOX >"</pre>
	The private and public key file path is on the host where EMCLI is executed.
credMap.computenodellomCred	Credential Type: ilom_creds (username, password)
Compute node ILOM credential is required to promote ILOM access point.	<pre>emcli create_named_credential -cred_name='CN_ILOM_CRED' -auth_target_type='oracle_si_server_map' -cred_type='ilom_creds' -cred_scope='global' -attributes="username:XXXXX;password:XXXXXXXX"</pre>
Target Type: oracle_si_server_map	Credential Type: ilom_key_auth_creds (private key file, passphrase)
	<pre>emcli create_named_credential -cred_name='CN_ILOM_SSH_CRED' -auth_target_type='oracle_si_server_map' -cred_type='ilom_key_auth_creds' -cred_scope='global' - attributes="username:XXXXX;privatekeyfile:<PR IVATE_KEY_PATH_ON_AGENT_BOX>"</pre>
	The private key file path is on the agent box.

Property Name and Description	Sample Command
credMap.computenodeAdminCred Compute node admin credential is required to setup SNMP subscription on compute node. Target Type: host	Credential Type: HostCreds (username, password) <pre>emcli create_named_credential -cred_name='HOST_ADMIN_CRED' -auth_target_type='host' -cred_type='HostCreds' -cred_scope='global' -target_type='host' - attributes="HostUserName:sample;HostPassword: XXXXXXXXX"</pre>
	Credential Type: HostSSHCreds (public, private key) <pre>emcli create_named_credential -cred_name='HOST_ADMIN_SSH_CRED' -cred_scope=global -cred_type=HostSSHCreds -auth_target_type=host - attributes="USERNAME:X;SSH_PVT_KEY:pvtkey;SSH _PUB_KEY:pubkey" -input_file="pvt:<PATH_TO_PRV_ FILE_EMCLI_INSTALL_BOX>; pub:<PATH_TO_PUB_FILE_EMCLI_INSTALL_BOX >"</pre>
credMap.computenodeSnmpCred Compute node SNMP credential is required for alerts. Target Type: oracle_si_server_map	Credential Type: SNMPV3Creds <pre>emcli create_named_credential -cred_name='CN_SNMP_V3_CRED' -auth_target_type='oracle_si_server_map' -cred_type='SNMPV3Creds' -cred_scope='global' -attributes="authUser:XXXXX;authPwd:YYYY; privProtocol: AES;privPwd: ZZZZZ;authProtocol:S HA;authPwd:WWWWW"</pre>
	Credential Type: SNMPV1Creds (SNMP V1/V2) <pre>emcli create_named_credential -cred_name='CN_SNMP_CRED' -auth_target_type='oracle_si_server_map' -cred_type='SNMPV1Creds' -cred_scope='global' -attributes="COMMUNITY:<community>"</pre>

Property Name and Description	Sample Command
credMap.computenodeliomRestCred Compute node REST credential to promote ILOM REST access point. Target Type: oracle_si_server_map	Credential Type: AliasCredential (Alias, password) <pre>emcli create_named_credential -cred_name=ILOM_HTTP_CREDS -cred_scope=global -cred_type=AliasCredential -auth_target_type=oracle_si_server_map -attributes="Alias:root;Password:XXXXXXXX"</pre>
credMap.ibILOMcred IB Switch ILOM admin credential is required to promote ILOM access point. Target Type: oracle_si_netswitch	Credential Type: ilom_creds (username, password) <pre>emcli create_named_credential -cred_name='IB_ILOM_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='ilom_creds' -cred_scope='global' -attributes="username:XXXXXX;password:XXXXXXXX"</pre>
	Credential Type: ilom_key_auth_creds (private key file, passphrase) <pre>emcli create_named_credential -cred_name='IB_ILOM_SSH_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='ilom_key_auth_creds' -cred_scope='global' -attributes="username:XXXXXX; privatekeyfile:<PRIVATE_KEY_PATH_ON_AGENT_BOX >"</pre> <p>The private key file path is on the agent box.</p>

Property Name and Description	Sample Command
credMap.ibSnmpCred IB Switch SNMP credential is required to promote SNMP access point.	Credential Type: SNMPV1Creds
Target Type: oracle_si_netswitch	<pre>emcli create_named_credential -cred_name='IB_SNMP_V1_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='SNMPV1Creds' -cred_scope='global' -attributes="COMMUNITY:public"</pre>
	Credential Type: SNMPV3Creds
	<pre>emcli create_named_credential -cred_name='IB_SNMP_V3_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='SNMPV3Creds' -cred_scope='global' -attributes="authUser:XXXXX;authPwd:YYYY; privProtocol: AES;privPwd: ZZZZZ;authProtocol:SHA;authPwd:WWWWW"</pre>

Property Name and Description	Sample Command
credMap.cellMonitorCred	Credential Type: ExaCreds (CellCLI based monitoring)
Storage Server monitored using CellCLI – root user credential required to setup password less SSH between agent and Storage Server to monitor Storage Server.	emcli create_named_credential -cred_name='CELL_ROOT_CRED' -auth_target_type='oracle_exadata' -cred_type='ExaCreds' -cred_scope='global' -target_type='oracle_exadata' - attributes="ExaUserName:root;ExaPassword:XXXX XXX"
Or Storage Server monitored using ExaCLI/ RESTAPI - monitoring credential required to monitor Storage Server.	
Target Type: oracle_exadata	
	Credential Type: ExaSSHCreds (CellCLI based monitoring)
emcli create_named_credential -cred_name=CELL_SSH_CREDS -cred_type=ExaSSHCreds -auth_target_type=oracle_exadata - attributes="USERNAME:root;SSH_PVT_KEY:pvtkey" - input_file="pvtkey:<KEY_AVAILABLE_AT_EMCLI_IN_STALL_BOX>"	
The private and public key file path is on the host where EMCLI is executed.	Credential Type: ExaCLICred (ExaCLI/RESTAPI based monitoring)
emcli create_named_credential -cred_name='CELL_EXACLI_CRED' -auth_target_type='oracle_exadata' -cred_type='ExaCLICred' -cred_scope='global' -target_type='oracle_exadata' - attributes="ExaCLIUserName:celladmin;ExaCLIPassword:XXXXXXXX"	

Property Name and Description	Sample Command
credMap.cellSnmpCred Storage Server SNMP credential is required to receive alert. Target Type: oracle_exadata	Credential Type: SNMPV1Creds
	<pre>emcli create_named_credential -cred_name='CELL_SNMP_V1_CRED' -auth_target_type='oracle_exadata' -cred_type='SNMPV1Creds' -cred_scope='global' -attributes="COMMUNITY:<community>"</pre>
	Credential Type: SNMPV3Creds
	<pre>emcli create_named_credential -cred_name='CELL_SNMP_V3_CRED' -auth_target_type='oracle_exadata' -cred_type='SNMPV3Creds' -cred_scope='global' -attributes="authUser:XXXXX;authPwd:YYYY; privProtocol: AES128;privPwd: ZZZZZ;authProtocol: SHA;authPwd: WWWW"</pre>
	To specify the privacy protocol name, use AES128 or DES.

Property Name and Description	Sample Command
credMap.cellSntpSetupCred Storage Server admin credential is required to setup SNMP subscription. Target Type: oracle_exadata	Credential Type: ExaCreds (CellCLI based monitoring)
	<pre>emcli create_named_credential -cred_name='CELL_ADMIN_CRED' -auth_target_type='oracle_exadata' -cred_type='ExaCreds' -cred_scope='global' -target_type='oracle_exadata' - attributes="ExaUserName:YYYYYY;ExaPassword:XX XXXXX"</pre>
	Credential Type: ExaSSHCreds (CellCLI based monitoring)
	<pre>emcli create_named_credential -cred_name=CELL_ADMIN_SSH_CREDS -cred_type=ExaSSHCreds -auth_target_type=oracle_exadata - attributes="USERNAME:YYYYYY;SSH_PVT_KEY:pvtkey " - input_file="pvtkey:<KEY_AVAILABLE_AT_EMCLI_IN STALL_BOX>"</pre>
	Credential Type: ExaCLICred (ExaCLI/RESTAPI based monitoring)
	<pre>emcli create_named_credential -cred_name='CELL_ADMIN_EXACLI_CRED' -auth_target_type='oracle_exadata' -cred_type='ExaCLICred' -cred_scope='global' -target_type='oracle_exadata' - attributes="ExaCLIUserName:celladmin;ExaCLIPa ssword:XXXXXXXX"</pre>

Property Name and Description	Sample Command
credMap.ciscoIosCred Cisco Switch IOS credential is required to promote Cisco IOS access point	Credential Type: cisco_creds (username, password, privilege password)
Target Type: oracle_si_netswitch	<pre>emcli create_named_credential -cred_name='CISCO_IOS_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='cisco_creds' -cred_scope='global' - attributes="username:XXXX;userpass:YYY;privpass:ZZZZ"</pre>
	Credential Type: cisco_key_auth_creds (username, private key file, passphrase, privilege password)
	<pre>emcli create_named_credential -cred_name='CISCO_IOS_KEY_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='cisco_key_auth_creds' -cred_scope='global' -attributes="username:XXXX; privatekeyfile:<PRV_KEY_FILE_LOCATION_ON_AGENT_HOST>; privatekeypassphrase:YYYYYY;privpass:ZZZZ"</pre>
	The private key file path is on the agent host.
credMap.ciscoSnmpCred Cisco Switch SNMP credential is required to promote Cisco SNMP access point.	Credential Type: SNMPV1Creds
Target Type: oracle_si_netswitch	<pre>emcli create_named_credential -cred_name='CISCO_SNMP_V1_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='SNMPV1Creds' -cred_scope='global' -attributes="COMMUNITY:public"</pre>
	Credential Type: SNMPV3Creds
	<pre>emcli create_named_credential -cred_name='CISCO_SNMP_V3_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='SNMPV3Creds' -cred_scope='global' -attributes="authUser:XXXXX;authPwd:YYYY; privProtocol: AES;privPwd: ZZZZZ;authProtocol: SHA;authPwd: WWWW"</pre>

Property Name and Description	Sample Command
credMap.roceSwCred RoCE Switch credential is required to promote RoCE IOS access point	Credential Type: cisco_creds (username, password, privilege password)
Target Type: oracle_si_netswitch	<pre>emcli create_named_credential -cred_name='ROCE_IOS_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='cisco_creds' -cred_scope='global' - attributes="username:XXXX;userpass:YYY;privpass:ZZZZ"</pre>
	Credential Type: cisco_key_auth_creds (username, private key file, passphrase, privilege password)
	<pre>emcli create_named_credential -cred_name='ROCE_IOS_KEY_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='cisco_key_auth_creds' -cred_scope='global' -attributes="username:XXXX; privatekeyfile:<PRV_KEY_FILE_LOCATION_ON_AGENT_BOX>; privatekeypassphrase:YYYYYY;privpass:ZZZZ"</pre>
	The private key file path is on the agent box.
credMap.roceSnmpCred This credential is required to promote ROCE SNMP access point.	Credential Type: SNMPV1Creds
Target Type: oracle_si_netswitch	<pre>emcli create_named_credential -cred_name='ROCE_SNMP_V1_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='SNMPV1Creds' -cred_scope='global' -attributes="COMMUNITY:public"</pre>
	Credential Type: SNMPV3Creds
	<pre>emcli create_named_credential -cred_name='ROCE_SNMP_V3_CRED' -auth_target_type='oracle_si_netswitch' -cred_type='SNMPV3Creds' -cred_scope='global' -attributes="authUser:XXXXX;authPwd:YYYY; privProtocol:DES;privPwd:ZZZZZ;authProtocol:SHA;authPwd:WWWWW"</pre>

Property Name and Description	Sample Command
credMap.pduHttpCred PDU HTTP credential is required to promote PDU. Target Type: oracle_si_pdu	Credential Type: http (username, password) <pre data-bbox="812 329 1336 413">emcli create_named_credential -cred_name='PDU_HTTP_CRED' -auth_target_type='oracle_si_pdu' -cred_type='http' -cred_scope='global' -attributes="username:YYYYY;password:XXXXX"</pre>
credMap.pduSnmpCred PDU SNMP credential is required setup SNMP subscription. Target Type: oracle_si_pdu	Credential Type: SNMPV1Creds <pre data-bbox="812 413 1336 604">emcli create_named_credential -cred_name='PDU_SNMP_CRED' -auth_target_type='oracle_si_pdu' -cred_type='SNMPV1Creds' -cred_scope='global' -attributes="COMMUNITY:public"</pre>

Set Up SSH Key Authentication

To use SSH Keys for authentication during discovery, generate and install the public key on each target that must be discovered.

- [Generate the SSH Key](#)
- [Add SSH Public Key to Each Target](#)

Generate the SSH Key

Use `ssh-keygen` tool to generate the new authentication key pairs for SSH. Command to generate key pair (Public key and Private Key) based on RSA authentication algorithm:

```
ssh-keygen -f <FILE_NAME> -t rsa -N ""
```

In the above command:

- `-f` option to indicate file name
- `-t` option to indicate algorithm (rsa/dsa)

- `-N` option to indicate passphrase

The same key pair can be used on all the monitoring agent hosts by copying them. Alternatively, you can generate different key pair on each monitoring agent host.

Add SSH Public Key to Each Target

Follow these steps to install the SSH public key on the various targets that you want to discover:

- [Add SSH Public Key to Exadata Storage Server and Exadata Compute Node](#)
- [Add SSH Public Key to ILOM](#)
- [Add SSH Public Key to Cisco Ethernet / RoCE Switch](#)

Add SSH Public Key to Exadata Storage Server and Exadata Compute Node

Copy the public key to the Exadata Storage Server, Compute Node, and the monitoring Agent host using the following command:

```
ssh-copy-id -i <RSA_PUBLIC_KEY_LOCATION> oracle@host
```

Add SSH Public Key to ILOM

You can use the following steps for InfiniBand ILOM as well as Compute Node ILOM:

1. Log in to the ILOM as a user that has administration capabilities that will be used as part of Database Machine discovery.
2. Navigate to the directory location of the user's SSH key.
3. Add SSH public key to the user's account:

```
set load_uri=<transfer_method>://  
<username>:<password>@<fully_qualified_hostname>/  
<path_to_public_key_filename>
```

For example:

```
set load_uri=scp://adminuser:userpswd@192.0.2.1/keys/sshkey_1.pub
```

Add SSH Public Key to Cisco Ethernet / RoCE Switch

1. Log in to the switch.
2. Enter the global configuration mode on the switch using the command:

```
configure terminal
```

3. Specify the public key for the user by pasting the entire content of the key in the specified location in the command:

```
<username> admin sshkey <paste_entire_public_key_content>
```

4. Save the configuration:

```
copy running-config startup-config
```

Create the Input File

Create the input file with the details of the configuration parameters, and named credentials. You can optionally add or remove components from the configuration by using the `skipComponentList` or `addComponentList` parameters.

Similarly, for Virtual Exadata Database Machine discovery, you can provide the component details in the input file.

The following tables provide the information about all the parameters that can be used in the input file:

- [Configuration Parameters](#)
- [Named Credentials and Commands to Create Them](#)
- [Parameters for Adding or Skipping Components](#)
- [Parameters for Adding or Skipping Virtual Platform Components](#)

Following is an example input file created by using the configuration parameters, named credentials, and additional components:

```
#####
#           Discovery and Monitoring Agent
#####
configMap.discoveryAgentUrl=https://cn01vm01.example.com:3442/emd/main/
configMap.monitorAgentUrl.0=https://cn01vm01.example.com:3442/emd/main/
configMap.monitorAgentUrl.1=https://cn02vm01.example.com:3442/emd/main/
#####
#           Database Machine Configuration and required configuration
#####

# 0 For No, 1 for Yes
configMap.isRemoteAgent=0
configMap.isIPV6=0

# Identify if the network is roce or ib
configMap.networkType=roce
# Uncomment/update either the configMap.ibSwitchName or
configMap.roceSwitchName parameter depending on what type of switch you have
#configMap.ibSwitchName=sw-iba0.example.com
configMap.roceSwitchName=sw-rocea0.example.com

#directory must exist for file creation
configMap.outputFileLoc=/home/oracle/exadata_discovery/output.txt

# Either CellCLI or ExaCLI or RESTAPI
configMap.cellMetricSrc=CellCLI

#These properties would be considered if Storage Server is monitored using
RESTAPI
configMap.HttpsProtocol=TLSv1.2
configMap.TrustSelfSignedCert=1
```

```
# 0 (no) or 1 (yes) to configure SNMP
configMap.computenodeSnmpSub=1
configMap.cellSnmpSub=1

#Location of the schematic file on discovery agent
configMap.schematicFileHost=cn01vm01.example.com
configMap.schematicFilePath=/opt/oracle.SupportTools/onecommand/
databasemachine.xml
#####
#           TARGET CREDENTIALS
#####
#Please provide named credentials for all target types of Database Machine.

# Remote Agent or Compute Node Host Credential
credMap.agentCred=SYSMAN:AGENT_HOST_CRED
credMap.schematicFileHostCred=SYSMAN:AGENT_HOST_CRED

# CELL (Storage Server) Credentials
credMap.cellMonitorCred=SYSMAN:CELL_ROOT_CRED
credMap.cellSnmpCred=SYSMAN:CELL_SNMP_CRED
credMap.cellSnmpSetupCred=SYSMAN:CELL_ROOT_CRED

# Compute Node Credentials
credMap.computenodeIломCred=SYSMAN:CN_ILOM_CRED
credMap.computenodeAdminCred=SYSMAN:CN_ADMIN_CRED
credMap.computenodeSnmpCred=SYSMAN:CN_SNMP_CRED

# PDU Credentials
credMap.pduSnmpCred=SYSMAN:PDU_SNMP_CRED
credMap.pduHttpCred=SYSMAN:PDU_HTTP_CRED

# CISCO Switch (Admin) Credentials
credMap.ciscoIosCred=SYSMAN:CISCO_IOS_CRED
credMap.ciscoSnmpCred=SYSMAN:CISCO_SNMP_CRED

# CISCO Switch (RoCE) Credentials
#credMap.roceSwCred=SYSMAN:ROCE_SWITCH_CRED
#credMap.roceSnmpCred=SYSMAN:ROCE_SNMP_CRED

# Infiniband Switch Credentials
credMap.ibIломCred=SYSMAN:IBILOM_ADMIN_CRED
credMap.ibSnmpCred=SYSMAN:SI_SWITCH_SNMP_CRED

#####
#           Virtual Platform Config Data
#####
#Virtual Platform Credentials
#Uncomment/complete if this is a Virtual Database Machine (Virtual targets
#are not yet discovered in EM)
vpTargetMap.vpCred=SYSMAN:VP_CRED

#Please provide virtual platform targets information
vpTargetMap.name.0=vp01.example.com
vpTargetMap.name.1=vp02.example.com
```

```
#####
#           SKIP Components
#####
#Specify any components you wish to skip
componentMap.skipComponentList.0=pdu0.example.com
```

Configuration Parameters

The following table lists the **mandatory** configuration parameters and their details that must be specified in the input file for EMCLI based discovery:

Parameter Name and Description	Allowed Value / Sample Value
configMap.discoveryAgentUrl Discovery agent URL on which schematic file is available.	Allowed Value: https://<em_agent_target_name>:<em_agent_port>/emd/main Sample Value: configMap.discoveryAgentUrl=https://agent1:1838/emd/main
configMap.monitorAgentUrl.0 configMap.monitorAgentUrl.1 Monitoring agent URL, one for primary and one for backup.	Allowed Value: https://<em_agent_target_name>:<em_agent_port>/emd/main Sample Value: configMap.monitorAgentUrl.0=https://agent01:1838/emd/main configMap.monitorAgentUrl.1=https://agent02:1838/emd/main
configMap.isRemoteAgent This property indicates monitoring agent(s) is a remote one, that is installed outside of the Database Machine. The purpose of a remote monitoring agent is to avoid consuming Database Machine resources for monitoring.	Allowed Value: 0 - Discovery is performed by a local agent 1 - Discovery is performed by remote agent Sample Value: configMap.isRemoteAgent=0
configMap.isIPV6 This property is to indicate IP Protocol used for communication.	Allowed Value: 0 - IPV4 Exadata Database Machine (<i>Default</i>) 1 – IPV6 Exadata Database Machine Sample Value: configMap.isIPV6=0
configMap.schematicFileHost This property is to specify the host name on which schematic file available. On the same host, the discovery agent should also be present.	Allowed Value: Host name Sample Value: configMap.schematicFileHost=adm01.example.com

Parameter Name and Description	Allowed Value / Sample Value
configMap.schematicFilePath	Allowed Value:
This property is to specify the schematic file full path on the host (host name specified for property configMap.schematicFileHost).	Full path of schematic file.
	Sample Value:
	configMap.schematicFilePath=/opt/oracle.SupportTools/onecommand/databasemachine.xml
configMap.ibSwitchName	Allowed Value:
This property is to indicate the InfiniBand Switch name if configMap.networkType parameter is set to ib.	Fully qualified InfiniBand Switch host name
This property is mandatory based on the network type.	Sample Value:
	configMap.ibSwitchName=sw-iba0.XXXX.YYY
configMap.roceSwitchName	Allowed Value:
This property is to indicate the RoCE Switch name if configMap.networkType parameter is set to roce.	Fully qualified ROCE Switch host name
This property is mandatory based on the network type.	Sample Value:
	configMap.roceSwitchName=sw-ROCEa0.XXXX.YYY
configMap.cellMetricSrc	Allowed Value:
This property is to indicate Exadata Storage Server monitoring mechanism.	CellCLI / ExaCLI / RESTAPI
	Sample Value:
	configMap.cellMetricSrc=RESTAPI
configMap.outputFileLoc	Allowed Value:
This property is to indicate the full path of the file on OMS host to log Deployment Procedure (DBMachineSystemCheck's) output.	Full path to file on OMS host
	Sample Value:
	configMap.outputFileLoc=/scratch/user/dp_output

The following table lists the **optional** configuration parameters and their details that you can specify in the input file for EMCLI based discovery:

Parameter Name and Description	Allowed Value / Sample Value
configMap.isStatic	Allowed Value:
Database Machine discovery by default performs dynamic discovery by invoking IB commands and other commands. But, in some cases discovery may not be able to proceed with dynamic approach like IB Switch is not accessible, Failure in detecting rack type, etc.,	0 - dynamic discovery be performed (<i>Default</i>) 1 – static discovery
In this case static discovery would be useful. Static discovery depends on schematic file for identifying components.	Sample Value:
	configMap.isStatic=0
configMap.DBMachineName	Allowed Value:
This property allows to define a custom name for Database Machine.	Any string can be specified as value for this property. If not specified discovery will automatically prepare a name for Exadata Database Machine as default.
	Sample Value:
	configMap.DBMachineName=DB Machine

Parameter Name and Description	Allowed Value / Sample Value
configMap.ExadataGridName This property allows to define a custom name for Exadata Grid.	Allowed Value: Any string can be specified as value for this property. If not specified discovery will automatically prepare a name for Exadata Grid as default. Sample Value: configMap.ExadataGridName=Exadata Grid
configMap.networkType This property is to indicate which network type is part of discovering Exadata Database Machine.	Allowed Value: roce to indicate RoCE X8M ib to indicate InfiniBand Network (<i>Default</i>) Sample Value: configMap.networkType=roce
configMap.HttpsProtocol This property is to indicate communication protocol used for Exadata Storage Server communication. This property needs to specify when Storage Server is monitored using RESTAPI. Otherwise, this property is not required.	Allowed Value: TLSv1.2 (<i>Default</i>) Sample Value: configMap.HttpsProtocol=TLSv1.2
configMap.TrustSelfSignedCert This property is to indicate whether a self-signed certificate is used for communication with Storage Server.	Allowed Value: 0 or 1 (<i>Default</i>) Sample Value: configMap.TrustSelfSignedCert=1
configMap.computenodeSnmpSub This property is to indicate whether SNMP Subscription for Compute Nodes is required.	Allowed Value: 0 - Skip SNMP Subscription (<i>Default</i>) 1 - Set up SNMP subscription Sample Value: configMap.computenodeSnmpSub=1
configMap.ilomVersionSupportRest This property is to indicate whether to promote ILOM REST Access Point using the computenodellomRestCred credential.	Allowed Value: true - Promote ILOM REST Access Point. When this property is set to true, you must also include the computenodellomRestCred property and its value in the input file. false - Skip ILOM REST Access Point Sample Value: configMap.ilomVersionSupportRest=true
configMap.cellSnmpSub This property is to indicate whether SNMP Subscription for Storage Servers is required.	Allowed Value: 0 - Skip SNMP Subscription (<i>Default</i>) 1 - Set up SNMP subscription Sample Value: configMap.cellSnmpSub=1

Parameters for Adding or Skipping Components

The following **optional** properties can be used to skip or add the components as part of EMCLI based discovery of Exadata Database Machine:

Property Name	Description	Example
componentMap.skipComponentList	Use this property to list the components that must be excluded from Exadata Database Machine discovery.	componentMap.skipComponentList.0=pdu1 componentMap.skipComponentList.1=pdu2
componentMap.addComponentList	Use this property to add new components to an existing database machine as part of Exadata Database Machine discovery.	componentMap.addComponentList.0=cel01 componentMap.addComponentList.1=cel02

Parameters for Adding or Skipping Virtual Platform Components

For Virtual Exadata Database Machine, the Virtual Platform targets can also be discovered as part of the EMCLI based Database Machine Discovery if not yet discovered in Enterprise Manager. The following input parameters can be used to specify the components:

Parameter Name and Description	Sample Value
vpTargetMap.name List of virtual platform targets to be discovered.	vpTargetMap.name.0=hostvp1 vpTargetMap.name.1=hostvp2

Parameter Name and Description	Sample Value
vpTargetMap.vpCred Target Type: oracle_si_virtual_platform	Credential Type: DMOvsBasicCreds

```

emcli create_named_credential
-cred_name='VP_CRED'
-
auth_target_type='oracle_si_virtual_platform'
-cred_type='DMOvsBasicCreds'
-cred_scope='global'
-
attributes="OVSUsername:root;OVSPassword:XXXXXXXX"

```

Credential Type: DMOvsSshKeyCreds

```

emcli create_named_credential
-cred_name='VPSSHCRED'
-
auth_target_type='oracle_si_virtual_platform'
-cred_scope='global'
-cred_type='DMOvsSshKeyCreds'
-
attributes="OVSUsername:root;OVSPassword:XXXXXXXX;
PassPhrase:test;SshPrivateKey:pvtkey;
SshPublicKey:pubkey;"
-
input_file="pvtkey:<PVT_KEY_LOC>;pubkey:<PUB_KEY_LOC> "

```

The private and public key file path is on the host where EMCLI is executed.

Optional Verification of the Prerequisites

To ensure that all the required parameters are provided in the input file, run a prerequisite-check deployment procedure before submitting the Exadata Database Machine discovery deployment procedure. Use the input file that you created earlier in the prerequisite-check deployment procedure.

This prerequisite-check deployment procedure is optional because it is also executed as part of the discovery deployment procedure.

Run this EMCLI command to perform the prerequisite-check and collect the information about the targets related to the Exadata Database Machine:

```
emcli submit_procedure -name=DBMachineSystemCheck -input_file="data://<input_file_absolute_path>"
```

For more information on the deployment procedure for EMCLI based discovery of Exadata Database Machine, see [Deployment Procedure for EMCLI](#).

Prerequisites for Discovery and Monitoring Using Non-root User

You can optionally perform discovery and monitoring operations using a non-root user.

Topics:

- [Non-root User for Systems Infrastructure Remote Host Targets](#)
- [Non-root User for Virtual Platform Targets](#)

Non-root User for Systems Infrastructure Remote Host Targets

To support the discovery through a non-root user, the list of commands used by the remote host and server OS metrics which require *sudo* privileges must be placed in the file */etc/sudoers*.

Place the following commands in the sudoers file:

```
<non-root-username> ALL=(ALL) /usr/sbin/sminfo,/usr/bin/ipmitool,/bin/cat
    /etc/iscsi/iscsid.conf,/sbin/fdisk,/sbin/vgs,/sbin/pvs,/sbin/lvs,/sbin/
    dmsetup,/usr/sbin/dbmcli,/sbin/service,/sbin/initctl,/opt/oracle.cellos/
    compmon/exadata_mon_hw_asr.pl,/usr/local/bin/imageinfo,/usr/sbin/
    exadata.img.hw,/usr/sbin/ibnetdiscover,/usr/sbin/smartctl
```

Use EMCLI to discover the remote host using a non-root user. Specify the dispatch URL as 'sudo-ssh://<IP/hostname>:22' by running the following EMCLI command:

```
emcli add_target -name=<REMOTE_HOST_NAME> -type=oracle_si_host_remote -
host=<AGENT_HOST_NAME>
-
monitoring_cred='ssh_creds_set;oracle_si_host_remote_ssh;ssh_creds;username:<user>;userpass:<password>;'
    -subseparator=properties='=' -properties='dispatch.url=sudo-ssh://
<target ip/hostname>:22'
    -access_point_name="<AP_NAME>" -
access_point_type="oracle_si_host_remote_ssh"
```

Non-root User for Virtual Platform Targets

To be able to discover the virtual platform targets using the non-root user, you must first update the */etc/sudoers* file with the list of commands used by the targets which require *sudo* privileges. Next, update the non-root user monitoring credentials with the *sudo* attributes.

Xen Virtual Platform

Content for the */etc/sudoers* file for Xen virtual platform target:

```
Cmnd_Alias OEM_MON_CMNDS = /usr/sbin/xl list*,/usr/sbin/xl info*,/usr/sbin/xl
    network-list*,/usr/sbin/xl uptime*,/usr/sbin/xl
    vcpu-list*,/usr/sbin/dmidecode,/usr/sbin/xentop,/usr/bin/xenstore-
```

```
read,/usr/bin/xenstore-list,/usr/bin/xenstore-ls,/sbin/ethtool
<non-root-user>    ALL=(ALL)      NOPASSWD:OEM_MON_CMNDS
```

KVM Virtual Platform

Content for the `/etc/sudoers` file for KVM virtual platform target:

```
non-root-user ALL=(ALL) NOPASSWD: /usr/bin/virsh list*, /usr/bin/virsh
dominfo*, /usr/bin/virsh nodememstats*, /usr/bin/virsh domstats*, /usr/bin/virsh
capabilities, /usr/bin/virsh domblklist*, /usr/bin/virsh domiflist*, /usr/bin/virsh
vcpuinfo*, /usr/bin/virsh domifstat*, /bin/virsh domblkstat*, /bin/virsh
dommemstat*, /bin/virsh nodeinfo, /bin/virsh domblkinfo*, /usr/sbin/dmidecode, /sbin/
ethtool, /usr/bin/ipmitool, /usr/sbin/imageinfo, /usr/local/bin/imageinfo,
/opt/oracle/bda/bin/imageinfo, /opt/exadata_ovm/vm_maker, /usr/sbin/
brctl
```

Update and save the monitoring credentials of the Virtual Platform target

Set the following `sudo` attributes:

1. Set **Privilege Type** to `sudo`.
2. Set **Run as** to `root`.
3. Set the **Privilege Command** to `/usr/bin/sudo -S -u %RUN_AS% %COMMAND%`. Verify the path of `sudo`.
4. Set **Is Privileged User** to `Yes`.
5. Unless the `requiretty` flag is explicitly un-set for the non-root user in the `/etc/sudoers` file, set **Enable Pseudo Terminal** to `Yes`.

Now, you can use the non-root user instead of the `sudo` username for target discovery.

Exadata Database Machine Discovery

This chapter provides instructions for discovery of the Oracle Exadata Database Machine through Enterprise Manager.

The following sections describe how to discover the Exadata Database Machine and other supported targets:

1. [Discovering an Oracle Exadata Database Machine](#)
2. [Discovering Grid Infrastructure and RAC](#)
3. [Switch from Using CellCLI to RESTful API](#)

Discovering an Oracle Exadata Database Machine

This section describes various discovery flows that are supported in Enterprise Manager.

- [Discover the Oracle Exadata Database Machine Target Using the Console](#)
- [Discover the Oracle Exadata Database Machine Target Using EMCLI](#)

Comparison between the two methods for discovery:

EMCLI	Console
<ul style="list-style-type: none"> • The Database Machine discovery using EMCLI helps to automate the discovery process through scripts where several Database Machines must be discovered offline or in background. • It is accessible through classic programming language constructs, enabling tasks to be created and run either from the command-line or programmatically. EMCLI enables you to access Enterprise Manager functionality from text-based consoles (shells and command-line windows) for a variety of operating systems. • EMCLI is fully integrated with Enterprise Manager's security and user administration functions, enabling you to carry out operations using EM CLI with the same security and confidentiality as the Enterprise Manager console. For example, you can only see and operate on targets for which you are authorized. 	<ul style="list-style-type: none"> • Use the Console for an interactive experience of the discovery. • The entire discovery process is packaged together and is available in the Enterprise Manager console. • The discovery process is self-intuitive. You must only make sure to have the required information available to enter in the dialog boxes. • This method is best suited for one-time set up.

Discover the Oracle Exadata Database Machine Target Using the Console

- [Fresh Discovery of Oracle Exadata Database Machine Target](#)
- [Fresh Discovery of Exadata Database Machine on SuperCluster Systems](#)
- [Convert Database Machine Targets with 12c Target Types to 13c Using the Console](#)

Fresh Discovery of Oracle Exadata Database Machine Target

This discovery flow shows the steps involved in discovering Database Machine targets. This flow is the recommended option for discovering Database Machine targets afresh.

An Oracle Exadata Database Machine can be deployed in a physical or virtual configuration. Enterprise Manager supports both of these configurations. There are additional steps required in order to discover a virtual Exadata Database Machine. These additional steps are addressed inline as necessary.

Discovering the Oracle Virtual Platform

If discovering a virtual Exadata, the Oracle Virtual Platform target must be discovered before discovering the Oracle Exadata Database Machine target. The steps below detail the discovery of the Oracle Virtual Platform target.

Note

If you are using EMCLI to perform Exadata Database Machine discovery, the Oracle Virtual Platform targets can be discovered as part of that discovery without a separate discovery operation. See [Discover the Oracle Exadata Database Machine Target Using EMCLI](#).

1. From the Enterprise Manager home page, select the **Setup** menu (upper right corner), **Add Target**, and then **Add Targets Manually**.
2. On the Add Targets Manually page, click **Add Targets Using Guided Process**.
3. Under the Guided Discovery column, select **Oracle Virtual Platform** and click **Add**.
4. On the Discover Oracle Virtual Platforms screen, do the following:
 - a. Select the agent installed on the first compute node *domU* (virtual machine) for the Monitoring Agent and select the second compute node *domU* (virtual machine) agent for the *Failover Monitoring Agent*.
 - b. For the *Credential Properties*, enter the user name for the root or non-root user and provide the password.
 - c. In the *Hostname and IP Address* section, click on the **Add** button and type in the fully qualified hostname for the *dom0* servers (one per line) and click **Add**.
 - d. Click **Submit**.

The screenshot shows the 'Discover Oracle Virtual Platforms' dialog. It includes sections for 'Default Monitoring Agent and Credentials' (with monitoring agents 'host3vm1.example.com' and 'host4vm1.example.com'), 'Metric Execution Timeout (seconds)' (set to 90), 'Privileged User' (set to 'Yes'), and 'Credential Type' (set to 'Host Credentials'). The 'Hostname and IP Address' section lists hosts 'host3.example.com' and 'host4.example.com' with their respective agent and monitoring agent details. A 'Global Target Properties (Optional)' section is also present.

- After the discovery of the Oracle Virtual Platform target is complete, follow the steps in the next section to discover the Oracle Exadata Database Machine target.

Discovering the Oracle Exadata Database Machine

- From the Enterprise Manager home page, select the **Setup** menu (upper right corner), **Add Target**, and then **Add Targets Manually**.
- On the Add Targets Manually page, click **Add Targets Using Guided Process**.
- From the Add Using Guided Process window, select **Oracle Exadata Database Machine** from the list and click **Add**.
- On the Oracle Exadata Database Machine Discovery page, the following two options are presented. This document details the steps for discovering a new Exadata Database machine and components:
 - Discover a new Database Machine and its hardware components as targets. The table will update to show the target types and the credentials required for discovery.
 - Discover newly added hardware components in an existing Database Machine as targets. Select a Database Machine from the drop-down menu. The table will update to show the target types and the credentials required for discovery. The subsequent dialogs in the discovery flow will be similar to those depicted in this section but will be restricted to only undiscovered components.

Select the task and click **Discover Targets**. The Exadata Discovery Wizard begins.

- On the Discovery Inputs page, enter the following information:
 - For the Discovery Agents section:
 - Agent URL:** Select an agent deployed on a physical compute node or if virtual Exadata on a domU (virtual machine) compute node. Click the search icon to select from available agents.
 - For the Schematic Files section:
 - Once you have specified the Agent URL, a new row (hostname and schematic file information) is automatically added. The default schematic file, `databasemachine.xml`, describes the hardware components of the Exadata Database Machine.
 - Click **Set Credential** to set the credentials for the host.
 - Check and modify if necessary the schematic file location.

- Select the schematic file name from drop-down menu.

 **Note**

You only need to customize the Schematic file location if one of the following is true:

- The Schematic file is not in the default location.
- The Database Machine contains a Storage Expansion Rack.
- The hardware components are specified in multiple schematic files.

Click **Next**.

6. The InfiniBand Discovery page looks different depending on whether the Oracle Exadata Database Machine has been identified as having an InfiniBand (IB) or RDMA over Converged Ethernet (RoCE) storage network. Screenshots and the details for each option are shown below.

If this is an IB Oracle Exadata, the following information is required,

- **IB Switch Host Name:** The host name for one of the IB switches for the Oracle Exadata. The IB Switch hostname is usually pre-populated.
- **InfiniBand Switch ILOM host credential:** The `root` or `ilom-admin` user name and password for the InfiniBand switch ILOM host.

Click **Next**.

If this is Oracle Exadata X8M, no credentials are required at this point in the discovery flow. The InfiniBand Discovery page will display a message that the Database Machine to be discovered contains RoCE switches and no additional inputs are needed.

Click **Next**.

 **Note**

If this is Oracle Exadata X8M and the form displayed requests InfiniBand information, then it likely means there was an issue accessing the `databaseMachine.xml` file specified in the previous step. Navigate back and correct this issue. Ensure the proper location and credentials are specified for the `databaseMachine.xml` file and that the file exists.

7. On the Prerequisite Check page, Enterprise Manager will try to discover hardware components dynamically based on your environment. If any critical issues are encountered, then you can click **Back** to resolve them. Enterprise Manager will show the problem and its severity (Info, Warning, or Critical).

Any warning issues or informational messages may also be displayed. These do not interfere with the discovery process.

Click **Next**.

8. On the Components page, the components below are pre-selected. You can deselect any unwanted component. Components include, but are not limited to, the following:

- **Compute Node:** Select the hosts that are compute nodes in the Oracle Exadata Database Machine.
- **Oracle Exadata Storage Server:** Select the Oracle Exadata Storage Servers that are part of this Oracle Exadata Database Machine target.

- **InfiniBand Switch:** Displayed if this is an IB Oracle Exadata. Select the InfiniBand Switches that are part of the Oracle Exadata Database Machine. These also will be added as managed targets.
- **Ethernet Switch:** Select the Ethernet switches that are part of the Oracle Exadata Database Machine. If this is an IB Oracle Exadata, select the admin switch. If this is Oracle Exadata X8M, select the admin and RoCE switches. The Ethernet switches will be added as managed targets.
- **Compute Node ILOM:** Select the Integrated Lights Out Managers (ILOMs) of the compute nodes that are part of this Oracle Exadata Database Machine. These ILOMs will be added as managed targets.

 **Note**

A partitioned rack can list other components that should be deselected manually (for example, compute nodes).

Oracle Exadata Database Machine Discovery

Discovery Inputs Infiniband Discovery Prerequisite Check Components Monitoring Agents Credentials Review

Database Machine Discovery: Components

Set Target Properties Back Step 4 of 7 Next Cancel

Below are the new components that are discovered by examining the network topology and configuration files. Select the different types of components that are part of Database Machine. Note that all components are pre-selected by default the first time you visit the page.

Compute Nodes/Servers

Select the servers that are Compute Nodes in the Oracle Exadata Database Machine. These servers must be added as managed targets before the Oracle Exadata Database Machine target creation can proceed. Monitoring of the selected Compute Node Servers through their ILOM will also be enabled as part of the discovery process.

Server Name
host1vm01.example.com
host2vm01.example.com
host3vm01.example.com
host4vm01.example.com

Oracle Exadata Storage Server

Select the Oracle Exadata Storage Servers that are part of this Oracle Exadata Database Machine target.

Exadata Storage Server Name	Management IP
celadm01.example.com	10.110.10.1
celadm02.example.com	10.110.10.2
celadm03.example.com	10.110.10.3
celadm04.example.com	10.110.10.4

Infiniband Switch

Select the Infiniband switches that are part of this Oracle Exadata Database Machine target. If this DB Machine is partitioned or this is a SuperCluster, you should select all Infiniband switches in this DB Machine or SuperCluster when you discover the first DB Machine target. All subsequent DB Machine targets will share the same set of Infiniband switch targets.

Infiniband Switch Name
hostsw-iba0.example.com
hostsw-ibb0.example.com
hostsw-lbs0.example.com

Ethernet Switch

Select the Ethernet switches that are part of this Oracle Exadata Database Machine target.

Click **Next**.

9. By default, the agents are automatically assigned to the targets. To add more monitoring agents:

In the **Monitoring Agents** page, click **Select Agents** > in the pop-up window, select the required agents for each target on the Exadata Database Machine. Use the **Ctrl** key to select multiple agents.

① Note

A Warning will be thrown if only one agent is used as monitoring and backup agents. Click **Select Agents** and add more agents, if available.

Database Machine Discovery: Monitoring Agents

DB Machine guided discovery has automatically assigned agents to your targets. You can choose to override the assigned monitoring agents. Click Select Agents and then select the monitoring and backup agents from the lists below. Click Reset to return all targets to automatic assignments. When you are satisfied with your agent selection, click Next to continue.

Select Agents **Reset**

Target Name	Monitoring Agent	Backup Monitoring Agent
DB Machine dbm01.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
Exadata Storage Server Name	Monitoring Agent	Backup Monitoring Agent
celadm02.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
celadm03.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
celadm04.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
Infiniband Switch Name	Monitoring Agent	Backup Monitoring Agent
hostsw-iba0.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
hostsw-ibb0.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
hostsw-lbs0.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
Ethernet Switch Name	Monitoring Agent	Backup Monitoring Agent
hostsw-adm0.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
Compute Node Server ILOM Name	Monitoring Agent	Backup Monitoring Agent

Click **Next**.

10. On the Credentials page, set the credentials for all components within the Oracle Exadata Database Machine (for example, the Storage Server, PDU, InfiniBand Switch, and so forth). Click **Select Agents** to set the credential for the component. If the users and passwords are the same for all components of the type, choose **Same for all** and enter the user and password combination.
 - Depending on the component, you may need to provide additional information. For example, in the Oracle Exadata Storage Server Credentials window, you must enter the SNMP Credentials as well as the monitoring Credentials.
 - The recommended monitoring mechanism for the Oracle Exadata Storage Servers is RESTful API. An ExaCLI user is required to use the RESTful API. For details on how to create the necessary ExaCLI user, refer to [Creating Users for Use with ExaCLI](#) in *Oracle Exadata Database Machine Maintenance Guide*.
 - SNMP subscriptions are required for EM to receive alerts for hardware and software failures from Exadata components. Each component credential screen provides fields for SNMP configuration. The recommendation for SNMP subscriptions is to use SNMP v3 for all subscriptions. For IPv6 environments, only SNMPV3 credentials are supported.
 - The SNMP community string is like a password. It is sent along with each SNMP Get-Request, and allows or denies access to the device. Please check with your network administrator to get the appropriate community string details for your environment.

- For the Compute Node Server ILOM Credentials, you can additionally specify the named credentials for monitoring using RESTful API. Click **Test Connection** to verify the proper working of the SSH and RESTful API credentials. The ILOM version for REST monitoring must be *5.0.1 or higher*. Also, for REST monitoring, ensure to use the latest Exadata plug-in release update on the agent side. You can create a Compute node REST credential to promote ILOM REST access point by navigating in the Enterprise Manager as follows: Click **Setup** in the Enterprise Manager home > Click **Security** > Click **Named Credential** > Enter the details for the REST credentials.

The table below provides sample guidance on the recommended values for the credentials.

Table 3-1 Credential Details

Component	Credential	Monitoring Mechanism	SNMP Subscription	Notes
Agent	oracle (or the user who owns the agent installation)	.	.	.
Oracle Exadata Storage Server	celladministrator	RESTful API	SNMP V3	.
Infiniband Switch	root / ilom-admin	.	SNMP V3	This component will not be present when discovering RoCE based systems
Compute Node Server ILOM	root	SSH / Enhanced monitoring using RESTful API	SNMP V3	In case of enhanced ILOM monitoring using RESTful API, both <i>root</i> and <i>admin</i> user credentials are supported.
PDU	admin	.	SNMP V3	.
Ethernet Switch	admin	.	SNMP V3	In a RoCE based system, these credentials will be used for the Management switch discovery as well as the RoCE switch discovery. If the passwords are different between the Management and RoCE switches, passwords can be specified by switch.

Oracle Exadata Database Machine Discovery

Discovery Inputs Infiniband Discovery Prerequisite Check Components Monitoring Agents Credentials Review

Database Machine Discovery: Credentials

Agent
Specify the agent host users and passwords for all agents. The agent users and passwords are needed to set up SSH user equivalence between the agents and the Exadata Storage Server targets monitored by the agents and to configure the targets. They will not be saved in Enterprise Manager unless you save them as named credentials.

Set Credential	
Agent URL	*Credential
https://host3vm01.example.com:3872/emd/main/	oracle*****
https://host4vm01.example.com:3872/emd/main/	oracle*****

Compute Node Server ILOM
Specify credentials for all the Compute Node Server Integrated Lights Out Managers (ILOM). They will be saved in Enterprise Manager as monitoring credentials.

Set Credential	
ILOM Name	*Credential
host3-ilom.example.com	root*****, root*****, v3u...
host4-ilom.example.com	root*****, root*****, v3u...

PDU
Specify the HTTP credentials and SNMP credentials for all PDUs. They will be saved in Enterprise Manager as monitoring credentials.

Set Credential	
PDU Name	*Credential
pdua0.example.com	admin*****, v3user*****
pdub0.example.com	admin*****, v3user*****

Ethernet Switch
Specify the Cisco IOS and SNMP credentials for all the Ethernet switches. They will be saved in Enterprise Manager as monitoring credentials.

Set Credential	
Ethernet Switch Name	*Credential
hostsw-adm0.example.com	admin*****, v3user*****

Infiniband Switch
Specify the Integrated Lights Out Manager (ILOM) credentials and SNMP credentials for all Infiniband switches. They will be saved in Enterprise Manager as monitoring credentials and preferred credentials for administration tasks.

Set Credential	
Infiniband Switch Name	*Credential
hostsw-iba0.example.com	ilom-admin*****, v3user..."
hostsw-ibb0.example.com	ilom-admin*****, v3user..."
hostsw-ibs0.example.com	ilom-admin*****, v3user..."

Click Next.

At this point, if you want to optionally discover Exascale, for details see [Discovering Exascale](#).

11. On the Review page, verify each section is correct. The following figure shows an example of an accurate review.

Oracle Exadata Database Machine Discovery

Discovery Inputs Infiniband Discovery Prerequisite Check Components Monitoring Agents Credentials Review

Database Machine Discovery: Review

Verify the following information. You can click Back to revise the inputs or click Submit to complete the discovery process.

System Target

The following system targets will be added in Enterprise Manager. You can keep the pre-filled system target names or specify your own system target names. If you choose to use your own system target names, make sure that 1) the system target names are unique in Enterprise Manager and 2) you use similar target names for the systems so that you can easily identify systems that are in the same DB Machine.

Target Name	Target Type	Monitoring Agent	Backup Monitoring Agent
DB Machine dbm01.example.com	Oracle Exadata Database Machine	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
Exadata Grid dbm01.example.com	Oracle Exadata Storage Server Grid	Agent not needed for repository target	Agent not needed for repository target

Rack

The following racks will be added as members of the Database Machine target.

Target Name
DB Machine dbm01 Rack (none)

Compute Nodes/Servers

The following compute nodes will be added as members of the Exadata Database Machine target.

Server Name	ILOM Name	Monitoring Agent	Backup Monitoring Agent
host3vm01.example.com	host3-ilom.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
host4vm01.example.com	host4-ilom.example.com	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/

Oracle Exadata Storage Server

The following Exadata Storage Servers will be added as members of the Database Machine target. SNMP will be set up automatically on Exadata Storage Servers.

Exadata Storage Server Name	Management IP	Monitoring Agent	Backup Monitoring Agent
celadm02.example.com	10.110.10.2	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
celadm03.example.com	10.110.10.3	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/
celadm04.example.com	10.110.10.4	https://host3vm01.example.com:3872/emd/main/	https://host4vm01.example.com:3872/emd/main/

Click **Submit**.

12. After the Database Machine components are discovered, a Target Creation Summary page displays to show all discovered targets and components. If there are any errors, the Status column will display a red flag. Place your cursor over the flag for a pop-up window for additional information about the error.
13. Perform post-discovery verification before you begin using Enterprise Manager. See [Post-Discovery Configuration and Verification](#).

Fresh Discovery of Exadata Database Machine on SuperCluster Systems

You can use the Exadata plug-in to discover and monitor an Oracle SuperCluster system.

You can monitor only the Database and Exadata components of the Oracle SuperCluster installed on LDOM and Zones. To monitor specific LDOM or virtualization configurations at the operating system level, please use Oracle Enterprise Systems Infrastructure Plug-in 13c.

The SuperCluster discovery wizard provided by “Systems Infrastructure Plug-in” is mainly responsible for discovery of the SuperCluster Hardware targets including ZFS server (For details please refer superCluster discovery document from emsi plug-in). It is recommended that the users should go through this flow first and then perform the Exadata Database Machine Discovery.

You can use the Exadata plug-in to discover and monitor an Oracle SuperCluster system. You can monitor *only* the Database and Exadata components of the Oracle SuperCluster installed on LDOM and Zones. To monitor specific LDOM or virtualization configurations at the operating system level, please use Oracle Enterprise Manager Ops Center.

To discover the Oracle SuperCluster as an Exadata Database Machine:

1. See *Discovering and Promoting Oracle Solaris Zones* in *Oracle® Enterprise Manager Administrator's Guide* for completing the prerequisite of pushing the agent to the control domain and zones.
2. Open a Service Request (SR) and upload the following files from your first LDOM:

/opt/oracle.SupportTools/onecommand/onecommand.params
/opt/oracle.SupportTools/onecommand/config.dat

 **Note**

Make a backup of your `/opt/oracle.SupportTools/onecommand` directory as you will overwrite it with the new configurations.

3. Oracle Support will provide you the file(s) generated based on the information provided in your configuration files. Copy all files into your `/opt/oracle.SupportTools/onecommand` directory.
4. Ensure that READ permissions are open for all files in this directory as well as in the `/opt/oracle.SupportTools/ems` directory. The agent must be able to read these files during discovery.
5. Run Self Update from Enterprise Manager to download the Solaris SPARC agent software on the OMS. Apply that download to the OMS to make it available for deployment:
 - a. From Enterprise Manager, click **Setup**, then **Extensibility**, and finally **Self Update**. Click **Agent Software**.
 - b. Select **Oracle Solaris on SPARC (64-bit)**.
 - c. If the status is Available, then click **Download** from the Actions menu.
 - d. Once downloaded, you must Apply it to the OMS from the same Self Update page.For more details on Self Update in Online or Offline mode, see *Updating Enterprise Manager in Oracle® Enterprise Manager Administrator's Guide*.
6. Install an agent on each of the Database nodes of the Oracle SuperCluster:
 - a. From Enterprise Manager, select **Setup**, then **Add Target**, and finally **Add Targets Manually**.
 - b. Select **Add Host Targets** (selected by default) and click **Add Host**.
 - c. Continue with the installation wizard and ensure that you run `root.sh` on each node once installation is completed.
7. Configure the database nodes for Database Machine discovery.

This step is needed to set descriptions of each Database node by updating the Database node IP address, Host Channel Adapter ID, and management host name in the InfiniBand configurations using `set_nodedesc.sh`. Enterprise Manager Agents look for this information when discovering the Database Machine.

Run:

```
# ibnetdiscover | grep your_db_nodes
```

If no output is returned, then run the following command to set the node descriptions from all database nodes:

```
# /bin/ib_set_node_desc_ssc.sh
```

8. Discover Exadata Database Machine from Enterprise Manager using Manual Discovery wizard. This discovery process is the same as any other Exadata Database Machine target. See [Discovering an Oracle Exadata Database Machine](#).
9. Discover the HA RAC Cluster and the cluster databases and configure each target as usual.
10. Perform post-discovery verification before you begin using Enterprise Manager to monitor the Exadata Database Machine. See [Post-Discovery Configuration and Verification](#).

Convert Database Machine Targets with 12c Target Types to 13c Using the Console

To convert an Exadata Database Machine discovered by Enterprise Manager 12c or by Enterprise Manager 13c with 12c target types to be monitored by Enterprise Manager 13c:

1. From the Database Machine menu, select **Convert 12c Member Targets**.
2. Select a **Conversion Option**.
 - (Default) Delete 12c targets and their historical data. This option halts all monitoring of legacy targets and deletes the historical data. This option allows for a clean start for all Enterprise Manager 13c target monitoring.
 - Retain 12c targets and their historical data. This option preserves the historical data for legacy targets.

Once you have selected a conversion option, click **Next**.

3. Set the **Credentials**. For the conversion, you must set the credentials for the Exadata Database Machine components to be converted. These are a subset of the components that comprise the Exadata. For further details on recommended settings, see the table [Table 3-1](#) and the related details in the corresponding step in [Discovering the Oracle Exadata Database Machine](#).

Infiniband Switch Name	*Credential
InfiniBand-switch1.example.com	Required, Not Set
InfiniBand-switch2.example.com	Required, Not Set
InfiniBand-switch3.example.com	Required, Not Set

Compute Node Server ILOM	*Credential
ilom1.example.com	Required, Not Set

Once you have set the credentials, the Credentials page will update to show the credentials set for all components to be converted. Click **Next**.

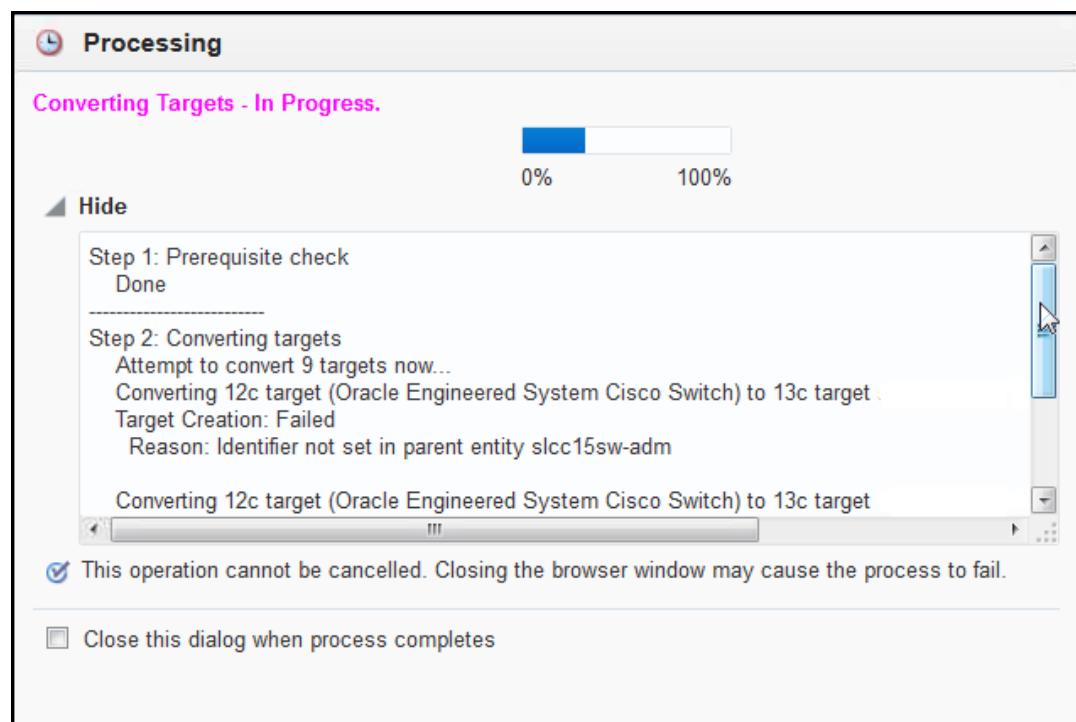
4. **Review** the components to be converted in the following three sections.

- **Summary:** This section summarizes the number of targets to be converted and if any targets will have additional monitoring information available to you after the conversion.
- **Targets to be Converted:** This table shows all Enterprise Manager 12c targets to be converted to 13c target types for monitoring by Enterprise Manager 13c and provides the new target name information.
- **Targets Not Impacted by Conversion:** All components of the Exadata Database Machine do not need to be converted to be monitored by Enterprise Manager 13c. This list identifies those components that do not need to be converted. These components are already monitored by Enterprise Manager 13c using 13c target types.

Click **Submit** to begin the conversion.

5. A **Processing** pop-up window will appear which shows you the status of the conversion and a summary of the success or failure of the components to be converted.

Do not close this window until the processing is complete!



You can select the option to close the window after the processing is complete, or click **Done**.

6. Once processing is complete, a **Conversion Results** page provides a summary of the components converted and details of each component that succeeded or failed to be converted:

DB Machine MyExadata.example.com Conversion Result

Summary

- Number of targets that have been converted successfully: 6
- Number of targets that failed to convert: 3
- 2 Targets will have additional monitoring mechanism
- IB Network `ib-network.example.com` will be monitored as IB Fabric.

Details

Conversion Status	12c Target Name	12c Target Type	New Target Name	New Target Type	Delete Target After Conversion	Comments
✗	Ethernet1	Oracle Ethernet	Ethernet1	Systems Infrastructure	No	Target Creation: Failed Reason: Identifier not set
✗	Ethernet2	Oracle Ethernet	Ethernet2	Systems Infrastructure	No	Target Creation: Failed Reason: Identifier not set
✓	ilom1.example.com	Oracle Ethernet	ilom1.example.com	Systems Infrastructure	No	Conversion succeeded.
✓	ilom2.example.com	Oracle Ethernet	ilom2.example.com	Systems Infrastructure	No	Conversion succeeded.
✗	pdu1.example.com	Oracle Ethernet	pdu1.example.com	Systems Infrastructure	No	Target Creation: Failed Reason: Identifier not set
✓	pdu2.example.com	Oracle Ethernet	pdu2.example.com	Systems Infrastructure	No	Conversion succeeded.
✓	ib-switch1.example.com	Oracle Infiniband	ib-switch1.example.com	Systems Infrastructure	No	Conversion succeeded.
✓	ib-switch2.example.com	Oracle Infiniband	ib-switch2.example.com	Systems Infrastructure	No	Conversion succeeded.

Recommendations: Launch the conversion UI again to fix the conversion error. Convert Remaining 12c Targets

• [Launch New Database Machine Home Page](#)

From this page, click **Convert Remaining 12c Targets** to repeat the conversion process for those components that failed to be converted. Click **Launch New Database Machine Home Page** to view the updated home page with all converted components monitored by Enterprise Manager 13c.

Discover the Oracle Exadata Database Machine Target Using EMCLI

EMCLI based discovery is customizable, enabling you to add or skip components in the event of component outage and select the monitoring agents. EMCLI Database Machine Discovery is idempotent, subsequent executions will make the best effort to discover any previously undiscovered targets (components) and will ignore any targets that have already been discovered and configured.

The discovery works for both physical and virtual Exadata Database Machines (IB/RoCE) using local or remote agent.

Ensure to complete the prerequisite tasks to discover the Database Machine target using EMCLI. For the list of prerequisites to complete, see [Prerequisites for Discovery Using EMCLI](#).

- [Submit the Deployment Procedure for the Discovery](#)
- [View the Discovery Status](#)
- [Convert 12c Type Database Machine Targets to 13c Using EMCLI](#)

Submit the Deployment Procedure for the Discovery

Log in to EMCLI. See login in *Oracle Enterprise Manager Command Line Interface*.

Run the following command to discover Oracle Exadata Database Machine:

```
emcli submit_procedure -name=DBMachineDiscovery -input_file="data://<input_file_absolute_path>"
```

In the above command:

- **DBMachineDiscovery**: This is the deployment procedure used for the discovery of Oracle Exadata Database Machine using EMCLI. See [Deployment Procedure for EMCLI](#).
- **input_file_absolute_path**: Full path of the input file that you created in the prerequisite steps. See [Create the Input File](#).

To receive notification regarding the discovery DP status, use the `-notification` flag as shown below:

```
emcli submit_procedure -name=DBMachineDiscovery -input_file="data://<input_file_absolute_path>" -notification="scheduled, action required, running" ;
```

Deployment Procedure for EMCLI

The Deployment Procedure (DP) **DBMachineDiscovery** is used for the discovery of Oracle Exadata Database Machine using EMCLI. This DP is composed two DPs:

- **DBMachineSystemCheck**: This DP takes the input file that you created while performing the prerequisite tasks and performs a set of prerequisite-checks. After the success of prerequisite-checks, this DP generates an output log file which is consumed by the next DP.
- **DBMachineSystemCreation**: This DP consumes the output log file generated by the DP for prerequisites-check, and performs Exadata Database Machine discovery. As a result, the steps such as creation of targets, associations, and promotion of targets are completed.

View the Discovery Status

You can monitor the status of the submitted deployment procedure using the EMCLI command or through the Enterprise Manager UI.

- [Check the Deployment Procedure Status Using EMCLI](#)
- [Check the Deployment Procedure Status Using Console](#)

In case of errors during discovery, resolve the error based on the deployment procedure output information.

See [EMCLI Discovery Issues](#).

Check the Deployment Procedure Status Using EMCLI

After you submitted the discovery deployment procedure, the execution ID is displayed in the command line as shown below:

```
emcli submit_procedure -name="DBMachineDiscovery" -input_file=data:"/home/user_jane/dpInput/emcli_discover_dp1_vir_remote"
```

```
Schedule not specified, defaults to immediate.  
6979503EF41F2447E05338C8F70A1796  
Deployment procedure submitted successfully
```

Use the execution ID to query the status and display it in a specific format. For example, if the execution ID is 69818CF4995E0C88E05338C8F70AA2C3:

```
emcli get_instance_status -instance=69818CF4995E0C88E05338C8F70AA2C3 -xml -  
details -showJobOutput -tailLength=50
```

An example out of the above command in xml format:

```
</step>  
<step isEnabled="true" name="promoteTargets" stepType="">>  
  <GUID>promoteTargets</GUID>  
  <errorModeString>stop</errorModeString>  
  <typeStep>Step</typeStep>  
  <status></status>  
  <startedOn>2018-04-10 08:56:28.0</startedOn>  
  <lastUpdatedOn>2018-04-10 08:58:16.0</lastUpdatedOn>  
  <completedOn>2018-04-10 08:58:16.0</completedOn>
```

Check the Deployment Procedure Status Using Console

The status and progress of the discovery can be viewed using the Enterprise Manager console.

From the main **Enterprise** menu, click **Provisioning and Patching** > click **Procedure Activity**. The following status information is displayed:

- Target discovery status: Success or Failure
- In case of failure, the reason for it, and additional information to correct it.

Resolve the issues displayed in deployment procedure output and re-submit the discovery through Enterprise Manager CLI or Enterprise Manager console to discover Exadata Database Machine or the remaining components of the database machine.

Convert 12c Type Database Machine Targets to 13c Using EMCLI

To convert an Exadata Database Machine discovered by Enterprise Manager 12c or by Enterprise Manager 13c with 12c target types to be monitored by Enterprise Manager 13c:

1. Create an input file by entering the information about the database machine target. Following is an example input file:

```
-----Input File credentials example-----  
# Config Map  
configMap.dbmTargetName=<DB Machine Target Name>  
configMap.retainOldTargets=false  
  
# Credentials  
credMap.ibIломCred=<IB Switch ILOM Credentials>  
credMap.ibSnmpCred=<IB Switch SNMP Credentials>  
credMap.pduHttpCred=<PDU Credentials>  
credMap.pduSnmpCred=<PDU SNMP Credentials>  
credMap.ciscoIosCred=<CISCO Switch IOS Credentials>  
credMap.ciscoSnmpCred=<CISCO Switch SNMP Credentials>  
credMap.computenodeIломCred=<Compute Node ILOM Credentials>  
credMap.computenodeAdminCred=<Compute Node Admin Credentials>
```

```
credMap.computenodeSnmpCred=<Compute Node SNMP Credentials>
credMap.computeSnmpOpt=false
credMap.agentCred=<Agent Credentials>
```

In the above example, the following parameters are used:

- configMap.dbmTargetName: The name of the Database Machine Target that you want to convert
- configMap.retainOldTargets: You can retain 12c targets and their historical data after the conversion. By default, this is set to false. This option halts all monitoring of legacy targets and deletes the historical data. Alternatively, you can set it to true to halt the monitoring of legacy targets and retain their historical data.
- credMap.ibIломCred: IB switch ILOM credentials
- credMap.ibSnmpCred: IB switch SNMP credentials
- credMap.pduHttpCred: PDU credentials
- credMap.pduSnmpCred: PDU SNMP credentials
- credMap.ciscoIosCred: CISCO switch IOS credentials
- credMap.ciscoSnmpCred: CISCO switch SNMP credentials
- credMap.computenodeIломCred: Compute node ILOM credentials
- credMap.computenodeAdminCred: Compute node Admin credentials. Provide this when credMap.computeSnmpOpt is set to true.
- credMap.computenodeSnmpCred: Compute node SNMP subscription credentials. Provide this when credMap.computeSnmpOpt is set to true.
- credMap.computeSnmpOpt: Compute node SNMP subscription flag. By default, this is set to false. If the flag is set to true, then after the conversion of ILOM, the respective ILOM SNMP is subscribed.
- credMap.agentCred: Agent credentials for compute node SNMP subscription.

For more information about the credentials, see [Named Credentials and Commands to Create Them](#).

2. Run the following EMCLI command to start the conversion of the Database Machine target from 12c type to 13c by specifying the path of the input file:

```
emcli submit_procedure -name=Convert12cTo13cDatabaseMachine -
input_file=<input_file_path> -notification="scheduled, action required,
running";
```

After running the command, note the deployment procedure (DP) number. You will need this to check the status of the conversion in the console.

3. To check the status of the deployment procedure in the console, in the **Enterprise** menu, click **Provisioning and Patching** > click **Procedure Activity**. The page lists all the deployment procedures.
4. Click the link corresponding to the deployment procedure you started in step 2. This displays the **Procedure Activity** page for your deployment procedure.
5. Under **Procedure Steps**, click **View** > click **Expand All**. You will now be able to view all the steps in the procedure like **Prerequisite**, **Conversion**, and **Summary**. The status of each step is displayed in the same section. To view more details of each step, click on it.

In case of failure in any of the steps, a detailed diagnostic is performed and information is available for you to apply the fix. After you apply the fix, re-run the EMCLI command specified in step 2 and repeat the steps after.

In the example below, you can view the summary of the conversion which indicates that one of the targets cannot be converted:

Discovering Grid Infrastructure and RAC

The process for discovering Grid Infrastructure (Cluster) and Real Application Cluster (RAC) related targets including Oracle High Availability Service and Cluster targets as well as ASM, database, listener, and related targets on Exadata is the same as on other platforms. See *Discovering and Adding Database Targets* in the *Oracle Enterprise Manager Cloud Control Administrator's Guide* for more information.

Switch from Using CellCLI to RESTful API

You can at any time lock down an Exadata Storage Server and switch to ExaCLI or RESTful API for monitoring/administration. It is recommended that you use RESTful API.

Ensure that the RESTful API user is created by following the steps in [Create an ExaCLI or RESTful API User](#).

For more information about ExaCLI, see [Using the ExaCLI Utility in Exadata Database Machine Maintenance Guide](#).

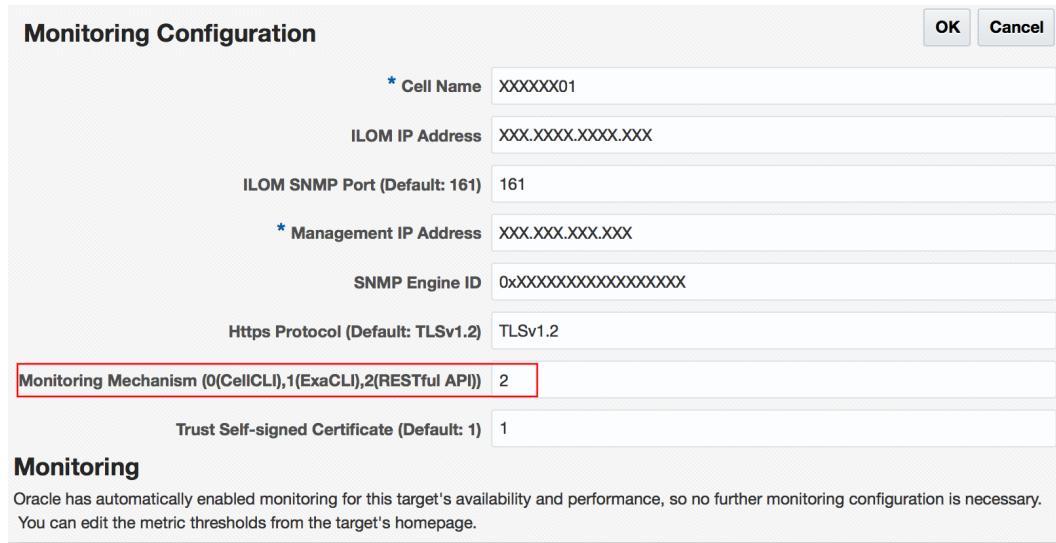
Topics:

- [Update the Monitoring Configuration Using Console](#)
- [Update the Monitoring Configuration Using EMCLI for Switching from CellCLI to RESTful API](#)

Update the Monitoring Configuration Using Console

For each Exadata Storage Server, provide the RESTful API monitoring credentials in the monitoring configuration.

1. From the Exadata Storage Server menu, select **Target Setup**, then select **Monitoring Configuration**.
The Monitoring Configuration dialog box opens.
2. Set the value of **Monitoring Mechanism** to:
 - 1 for ExaCLI
 - 2 for RESTful API
3. Set **Trust Self-signed Certificate** to 1.



Click **OK**.

4. From the main **Setup** menu, click **Security**, click **Monitoring Credentials**, and select **Exadata Storage Server** targets.

For each Exadata Storage Server, enter the RESTful API monitoring credentials.

Click **Save**.

Update the Monitoring Configuration Using EMCLI for Switching from CellCLI to RESTful API

If you have a large set of storage servers, then starting with Enterprise Manager 13.5 Release Update 11, you can use EMCLI to provide the monitoring credentials through a single input file for updating the overall monitoring configuration. This method is supported only if you want to switch to RESTful API.

Topics:

- [EMCLI Command Syntax](#)
- [Input File Parameters](#)
- [Sample Input File](#)

EMCLI Command Syntax

The following EMCLI command submits the deployment procedure *UpdateCellMonitorMechanism* and provides the monitoring configuration details through the input file *config_file_path*:

```
emcli submit_procedure -name=UpdateCellMonitorMechanism -  
input_file=<config_file_path>
```

Example command:

```
emcli submit_procedure -name=UpdateCellMonitorMechanism -  
input_file="data:/u01/oracle/exa_discovery/emcli_cell_mon_creds_conversion"
```

For more details about the *submit_procedure* command, see *submit_procedure* in *Command Line Interface*.

Input File Parameters

The following parameters must be provided in the input file for the deployment procedure:

File Parameter	Mandatory Parameter	Description
dbMachineMap.dbMachineName.N	Yes	The database machine name whose Storage Servers monitoring mechanism must be updated with a common REST API named credential. N is the index number of the database machine. For example, <i>dbMachineMap.dbMachineName.0</i> . The index number starts with 0 and increments with each new added database machine name.
dbMachineMap.agentNamedCreds.N	Yes	The agent named credentials corresponding to each database machine. As with database machine name, N is the index number of the database machine. For example, if <i>dbMachineMap.dbMachineName.0=myDatabase01</i> , then after using the same index number, the corresponding agent named credentials for the database machine may be <i>dbMachineMap.agentNamedCreds.0=sysman:DB01_CREDS</i> .
dbMachineMap.restNamedCreds	Yes	The REST API named credentials for Exadata Storage Server. The format for specifying it: <i><EM_USER>:<rest_named_creds></i>
dbMachineMap.monitoringMechanism	No	The default and only accepted value is 2 which indicates RESTful API.
dbMachineMap.selfSignedCertificate	No	The default value is 1 which indicates <i>Self-signed Certificate</i> . The other allowed value is 0.
dbMachineMap.securityProtocol	1	The default value is TLSv1.2 .

Sample Input File

In the following example, the database machines with index numbers 0 and 1 have database names `myDatabase01` and `myDatabase02`, and agent named credentials `sysman:DB01_CREDS` and `sysman:DB02_CREDS` correspondingly.

```
dbMachineMap.dbMachineName.0=myDatabase01
dbMachineMap.dbMachineName.1=myDatabase02
dbMachineMap.agentNamedCreds.0=sysman:DB01_CREDS
dbMachineMap.agentNamedCreds.1=sysman:DB02_CREDS
dbMachineMap.restNamedCreds=sysman:rest_creds
dbMachineMap.monitoringMechanism=2
dbMachineMap.selfSignedCertificate=1
dbMachineMap.securityProtocol=TLSv1.2
```

Post-Discovery Configuration and Verification

Once the Oracle Exadata Database Machine has been discovered through Enterprise Manager Cloud Control 13c, complete the following post-discovery verification:

- [Verify SNMP Subscription of the Oracle Management Agent to Exadata Storage Server](#)
- [Verify SNMP Configuration for InfiniBand Switch Targets](#)
- [Verify the Cisco Ethernet Switch SNMP Configuration](#)
- [Verify the PDU SNMP Configuration](#)
- [Convert SNMP V2 to V3 at Scale for All Hardware Components in Exadata](#)

Additionally, verify that all the components of Exadata Database Machine are discovered, up, and running. Also, ensure that the components are placed in the rack slots.

It is recommended that you review the logs to verify the successful completion of the discovery process. For the list of logs that you can review and their locations, see [Review Logs](#).

If you encounter issues during discovery, then to troubleshoot them, see [Discovery Troubleshooting](#) and [EMCLI Discovery Issues](#) in the troubleshooting section.

Verify SNMP Subscription of the Oracle Management Agent to Exadata Storage Server

Exadata Storage Server SNMP configuration is performed using the `cellcli` command and can be run in batch using `dcli` from a compute node.

Note

- During the discovery process, you can optionally provide the necessary root credentials to subscribe for SNMP traps from Exadata Storage Servers. If you have done so, then you can skip the remaining steps of this section and proceed with [Verify SNMP Configuration for InfiniBand Switch Targets](#).
- While the `public` string is used for the SNMP examples in the following sections, it is supported to use any valid SNMP community string.
- Use SNMP V3 subscriptions for IPv6-based Exadata Storage Server targets because of limitations in the SNMP V1 protocol.
- The SNMP receivelet listens on a single address and port for all monitored targets. The port is the UDP port with the same number as the TCP port used in the `EMD_URL`.

By default, the SNMP receivelet listens on all addresses; if the property `SnmpRecvletListenNIC` is set in the `emd.properties` file, the receivelet will attempt to resolve the value as either a name or IP address, and listen on only that address.

This parameter is independent of `AgentListenOnAllNICs` and `EMD_URL` because in some installations, the Agent may need to communicate with the OMS and with managed targets on different networks.

Using the `ALTER CELL` Command

While using the `ALTER CELL` command, all existing subscribers must be specified along with the new subscriber being added. Any existing subscribers not specified will be removed from the configuration. Similarly, you can also modify the `notificationPolicy` or `notificationMethod` attributes.

While using the `ALTER CELL` command, the `host=` and `community=` attribute values should be quoted, and `type=` is NOT quoted.

If you are using the DCLI utility to set up SNMP configuration, then any command containing punctuation, which will be interpreted by the local shell, must be enclosed with double quotation marks. If the command includes the following characters, then outer quotation marks and escape characters are required:

- `$` (dollar sign)
- `'` (quotation mark)
- `<` (less than)
- `>` (greater than)
- `()` (parentheses)

The backslash (`\`) is the escape character that allows the characters to be passed to the `cellcli` utility without being interpreted by the remote shell.

Exadata Storage Server supports SNMP v1/v2 and v3 version based alert notifications to the subscribers. To receive alerts in the SNMPv3 version, the monitoring agents must be added to the `snmpsubscriber` list using a `snmpuser`.

`cellcli` command to create a `snmpuser`:

```
# alter cell snmpuser = ((name=XXXXXX, authprotocol=sha, authpassword=*,  
privprotocol=AES, privpassword=*) );
```

In the above command:

- **name:** User name
- Use only * for the password values in the command.
- **authProtocol:** Authentication protocol. The supported protocol is **SHA**. The authProtocol must be specified for the `snmpuser` attribute. The system prompts for the authentication password. The authentication password must have 8 to 12 alphanumeric characters.
- **privProtocol:** Encryption protocol. The protocol options are **none**, **AES**, or **DES**. When the `privProtocol` attribute is not specified, the default value is **none**. The system prompts for an encryption password if the encryption protocol is specified. The password is exactly 8 alphanumeric characters, and they are case sensitive.

`cellcli` command to subscribe the monitoring agents for receiving SNMPv1 alert notifications:

```
# alter cell  
snmpsubscriber=((host=<Agent_Host>,port=<Agent_Port>,community=****) )
```

`cellcli` command to subscribe the monitoring agents for receiving SNMPv3 alert notifications:

```
# alter cell snmpsubscriber=((host=<Agent_Host>,port=<Agent_Port>,  
snmpUser=XXXXXX, type=v3))
```

In the above two commands:

- **Agent_Host:** Exadata Storage Server monitoring agent installed host
- **Agent_Port:** Exadata Storage Server monitoring agent port
- **snmpUser** is the one that is created using `snmpuser` attribute on Exadata Storage Server.

Check Current SNMP Configuration

Check the current SNMP configuration using the following `cellcli` commands:

1. To list the current subscribers for SNMP:

```
# cellcli -e list cell attributes snmpSubscriber
```

When correctly configured, this command should list the primary and backup agents for the Exadata Storage Server target, for example:

```
# cellcli -e "list cell attributes snmpSubscriber"  
(host=[hostname.mycompany.com],port=3872,community=public),  
(host=[hostname2.mycompany.com],port=3872,type=v3, snmpuser=v3user))
```

If the required monitoring agents are not added to the SNMP subscribers list, then add the subscription. See [Set Up SNMP Subscription for Exadata Storage Servers](#).

2. To list the currently configured notification methods:

```
# cellcli -e list cell attributes notificationMethod
```

The possible values of the `notificationMethod` attribute are **snmp**, **mail**, **none**, or a combination of **snmp** and **mail**. Specify the combination as `notificationMethod='mail,snmp'`. The default value of the attribute is **mail**. For example, when the configuration is set for a combination of **snmp** and **mail**, then the command returns the following response:

```
# cellcli -e list cell attributes notificationMethod  
snmp,mail
```

3. To list the currently configured notification policy for SNMP:

```
# cellcli -e list cell attributes notificationPolicy
```

Possible values are any or all of **Critical**, **Warning**, and **Clear**. When correctly configured, this command should return **Critical,Warning,Clear**, for example:

```
# cellcli -e list cell attributes notificationPolicy  
Critical,Warning,Clear
```

If SNMP configuration is not available, then configure it. See [SNMP Configuration Missing](#) in the troubleshooting section.

Configure SNMP Values Using `dcli` (optional)

The SNMP configuration commands can be run using `dcli` to perform the configuration in batch on all Exadata Storage Servers:

1. Set the `snmpSubscriber` value:

```
$ dcli -g cell_group -l root "cellcli -e \"alter cell snmpSubscriber=((host='[host  
name]',port=[port]),(host='[host name]',port=[port]))\""
```

Where `[host name]` and `[port]` values are Agents monitoring your Exadata Storage Server targets.

2. Set the `notificationMethod` value:

```
$ dcli -g cell_group -l root "cellcli -e \"alter cell  
notificationMethod='snmp,mail'\""
```

3. Set the `notificationPolicy` value:

```
$ dcli -g cell_group -l root "cellcli -e \"alter cell  
notificationPolicy='Critical,Warning,Clear'\""
```

Verify SSH Connectivity Using `cellcli`

To begin monitoring using `cellcli`, open a new terminal and verify whether the SSH connectivity was successfully established:

```
$ ssh -l cellmonitor <cell_ipaddress> cellcli -e 'list cell detail'
```

- If you are not prompted for any password, then you can assume that the connectivity is established.
- If you are asked to confirm whether you want to continue connecting, specify **Yes**.

Remove a Subscription

To remove the subscription, use the `ALTER CELL` command again by excluding the host name that you want to remove from the `snmpsubscriber` list.

Verify SNMP Configuration for InfiniBand Switch Targets

The SNMP configuration for Enterprise Manager monitoring of InfiniBand Switches is done automatically as part of the Enterprise Manager guided discovery process. It is good practice, however, to verify that SNMP configuration has been successful.

Note

- During the discovery process, you can optionally provide the necessary `root` credentials to set up SNMP trap for the InfiniBand Switch. If you have done so, then you can skip the remaining steps of this section.
- `ilom-admin` is the preferred user when discovering an InfiniBand Switch target for full monitoring. `root` user can also be used for full monitoring. If the switch is discovered using `root` or `ilom-admin` user, the community string provided during discovery is added to the switch, if it doesn't exist already.
- During SNMP V3 discovery, the V3 user added by the (discovery) user during discovery is added to the switch if the discovery is performed using `root` or the `ilom-admin` user.

To verify the InfiniBand Switch SNMP configuration for Enterprise Manager monitoring:

```
snmpget -v 1 -c <community_string> <hostname_of_IB_switch>
1.3.6.1.4.1.42.2.70.101.1.1.9.1.1.5
```

For example:

```
$ snmpget -v 1 -c public my_IB_switch.my_company.com
1.3.6.1.4.1.42.2.70.101.1.1.9.1.1.5
SNMPv2-SMI::enterprises.42.2.70.101.1.1.9.1.1.5 = INTEGER: 1
```

Note

If the Timeout message is displayed as output for the above command, then it means that the InfiniBand Switch is not yet configured for SNMP. To configure SNMPv3 user and SNMP community, see [Create an SNMPv3 User in IB Switches](#) and [Create SNMP Community for InfiniBand Switch Targets](#) in the prerequisites section.

Note

To remove the subscription, follow steps 1 to 5 in [Create SNMP Community for InfiniBand Switch Targets](#) (step 3 is not needed). In step 5, set the Agent compute node host name to **0.0.0.0** and the port to **0**.

Verify the Cisco Ethernet Switch SNMP Configuration

Run the `snmpwalk` command line utility or equivalent tool to verify the Cisco Switch configuration.

Run the following commands to fetch and display the data from the Cisco switch:

```
$ snmpget -v 1 -c <community_string> <hostname_of_cisco_switch> 1.3.6.1.4.1.9.2.1.56.0
$ snmpget -v 2c -c <community_string> <hostname_of_cisco_switch> 1.3.6.1.4.1.9.2.1.56.0
```

 **Note**

If a timeout message is displayed as an output for the above command, then it means that the Cisco Switch is not yet configured correctly.

If the Cisco Ethernet Switch SNMP configuration is missing, then see [Set Up SNMP Subscription for Cisco Ethernet Switch Targets](#) in the troubleshooting section.

Verify the PDU SNMP Configuration

Use the `snmpwalk` command line utility or equivalent tool to verify the PDU configuration.

Run the following command to fetch and display the data from PDU:

```
snmpget -v 1 -c <community_string> <hostname_of_pdu> 1.3.6.1.4.1.2769.1.2.3.1.1.1.0
```

 **Note**

If a timeout message is displayed as an output for the above command, then it means that the PDU is not yet configured correctly.

If the PDU SNMP configuration is missing, then see [Set Up SNMP for Power Distribution Unit \(PDU\) Targets](#) in the troubleshooting section.

Convert SNMP V2 to V3 at Scale for All Hardware Components in Exadata

Migrate the already discovered Exadata environments configured with SNMP v2 to use SNMP v3 without having to rediscover or manually perform the related operations.

1. To query the list of targets and their SNMP subscription status, create an input file by entering the information about the database machine target. Following is an example input file:

```
configMap.dbmTargetName=<DB Machine Target Name>
#
## Credentials
credMap.agentCred=<Agent Credentials>
```

```
credMap.computenodeIlomCred=<Compute Node ILOM Credentials>
credMap.computenodeSnmpV3Cred=<Compute Node SNMP V3 Credentials>

credMap.cellMonitorCred=<Storage Server Admin Credentials>
credMap.cellSnmpV3Cred=<Storage Server SNMP V3 Credentials>

credMap.pduHttpCred=<PDU Credentials>
credMap.pduSnmpV3Cred=<PDU SNMP V3 Credentials>

credMap.ciscoIosCred=<CISCO Switch IOS Credentials>
credMap.ciscoSnmpV3Cred=<CISCO Switch SNMP V3 Credentials>

credMap.ibIlomCred=<IB Switch ILOM Credentials>
credMap.ibSnmpV3Cred=<IB Switch SNMP V3 Credentials>
```

For the details about the values to provide for the parameters in the inputs file, see the table below.

2. Run the following EMCLI command to list all the targets in the Database Machine along with their SNMP subscription status:

```
emcli submit_procedure -name=DBMachineListSnmpSubscription -
input_file=<input_file_path>"
```

3. To set up SNMP V3 on all relevant targets in the Database Machine, create an input file by entering the information about the database machine target. Following is an example input file:

```
configMap.dbmTargetName=<DB Machine Target Name>

configMap.overridePassword=Y
configMap.overrideProtocols=Y

#
## Credentials
credMap.agentCred=<Agent Credentials>

credMap.computenodeIlomCred=<Compute Node ILOM Credentials>
credMap.computenodeAdminCred=<Compute Node Admin Credentials>
credMap.computenodeSnmpV3Cred=<Compute Node SNMP V3 Credentials>

credMap.cellMonitorCred=<Storage Server Admin Credentials>
credMap.cellSnmpV3Cred=<Storage Server SNMP V3 Credentials>

credMap.pduHttpCred=<PDU Credentials>
credMap.pduSnmpV3Cred=<PDU SNMP V3 Credentials>

credMap.ciscoIosCred=<CISCO Switch IOS Credentials>
credMap.ciscoSnmpV3Cred=<CISCO Switch SNMP V3 Credentials>

credMap.ibIlomCred=<IB Switch ILOM Credentials>
credMap.ibSnmpV3Cred=<IB Switch SNMP V3 Credentials>
```

For the details about the values to provide for the parameters in the inputs file, see the table below.

4. Run the following EMCLI command to set up SNMP V3 on all the targets:

```
emcli submit_procedure -name=DBMachineUpdateSnmpSubscription -
input_file=<input_file_path>
```

Details of the parameters used in the input file:

Parameter	Mandatory	Description
configMap.dbmTargetName	Yes	Database Machine name whose SNMP monitoring mechanism needs to be updated
configMap.overridePassword	No	Flag to decide if the target's authorization and privacy protocols' password should be updated with the SNMP monitoring credentials provided in the input file. Default value N. This flag is applicable only for DBMachineUpdateSnmpSubscription.
configMap.overrideProtocol	No	Flag to decide if the target's authorization and privacy protocols should be updated with the SNMP monitoring credentials provided in the input file. Default value N. This flag is applicable only for DBMachineUpdateSnmpSubscription.
credMap.agentCred	Yes	Agent credentials for compute node SNMP subscription. Format: <EM_USER>:<agent_named_creds>
credMap.cellMonitorCred	No	Storage server Admin credentials. Format: <EM_USER>:<cell_admin_name_creds> and must be of Type ExaCreds, ExaCLICred, or ExaSSHCreds.
credMap.cellnodeSnmpV3Cred	No	Storage server SNMP V3 credentials. Format: <EM_USER>:<cell_snmp_named_creds> and must be of Type SNMPV3Creds.
credMap.ciscoIosCred	No	CISCO switch IOS credentials. Format: <EM_USER>:<cisco_named_creds> and must be of Type cisco_creds, cisco_key_auth_creds, or cisco_private_key_auth_creds.
credMap.ciscoSnmpV3Cred	No	CISCO switch SNMP V3 credentials. Format: <EM_USER>:<cisco_snmp_named_creds> and must be of Type SNMPV3Creds.

Parameter	Mandatory	Description
credMap.computenodeAdminCred	No	Compute node Admin credentials. Format: <EM_USER>:<compute_admin_named_creds> and must be of Type ssh_creds.
credMap.computenodeIlomCred	No	Compute node ILOM credentials. Format: <EM_USER>:<compute_ilom_named_creds> and must be of Type ilom_creds, ilom_key_auth_creds, or ilom_private_key_auth_creds.
credMap.computenodeSnmpV3Cred	No	Compute node SNMP V3 subscription credentials. Format: <EM_USER>:<compute_snmp_name_creds> and must be of Type SNMPV3Creds
credMap.ibIlomCred	No	IB switch ILOM credentials. Format: <EM_USER>:<ib_named_creds> and must be of Type ilom_creds, ilom_key_auth_creds, or ilom_private_key_auth_creds.
credMap.ibSnmpV3Cred	No	IB switch SNMP V3 credentials. Format: <EM_USER>:<ib_snmp_named_creds> and must be of Type SNMPV3Creds.
credMap.pduHttpCred	No	PDU credentials. Format: <EM_USER>:<pdu_http_named_creds> and must be of Type http.
credMap.pduSnmpV3Cred	No	PDU SNMP V3 credentials. Format: <EM_USER>:<pdu_snmp_name_creds> and must be of Type SNMPV3Creds.

Oracle Exadata Database Machine Administration

This chapter provides the information needed to administer your Oracle Exadata Database Machine through Oracle Enterprise Manager. The following topics are discussed:

- [Creating Roles to Manage the Plug-in](#)
- [Update Monitoring User of an Already Discovered SI Target](#)
- [Exadata Database Machine Topology](#)
- [Provisioning Virtualized Exadata Clusters](#)
- [Exadata Storage Server Metrics and Alert Settings](#)
- [Exadata Storage Server Management](#)
- [Managing the InfiniBand Fabric](#)
- [Discovering and Monitoring Exascale](#)
- [Flash Cache Resource Monitoring](#)
- [Exadata Database Machine Fault Monitoring](#)
- [Monitoring Exadata Database Machine Components](#)
- [Automatic Indexing for Database Nodes](#)
- [Evaluating Compliance of Exadata Targets](#)
- [Reporting](#)

Creating Roles to Manage the Plug-in

To manage the plug-in, you need to create roles and administrators, and then assign roles to administrators. This restricts the privileges that each user has, for example in deleting the plug-in or accessing reports.

 **Note**

For security reasons, Oracle recommends that the `SYSMAN` account be used only as a template to create other accounts, and not used directly.

To manage the plug-ins, you must have the Self Update privileges. See [Assigning Self Update Privileges to Users in Cloud Administration Guide](#). For information on creating custom roles, see [Creating Custom Roles for Self Service Application Administrators and Users in Cloud Administration Guide](#).

To create roles to provide management rights to users:

1. Log in to the Enterprise Manager Cloud Control as the super administrator user.
2. Click **Setup**, then **Security**.

3. Select Roles.

On the Security page, a list of predefined roles is provided. These roles can serve as basis to define custom roles to suite specific site level requirements.

 **Note**

The predefined roles provided cannot be edited or deleted.

4. Select a role that closely matches the role you wish to create. Click **Create Like**.
5. On the Properties page, enter a name for the role you wish to create. You can optionally add a description.

Click **Next**.

6. On the Roles page, select the roles from the list of Available Roles. Click **Move** to add the role to Selected Roles.

Click **Next**.

7. On the Target Privileges page, select the privilege you want to grant to the new role.

Click **Next**.

8. On the Resource Privileges page, you can edit specific privileges to be explicitly granted. Click the **Manage Privilege Grant** edit icon to make the changes.

Click **Next**.

9. On the Administrators page, select the administrators from the list of Available Administrators that you want to grant the new role to. Click **Move** to add the administrator to Selected Administrators.

Click **Next**.

10. On the Review page, a complete summary of the new role you have created is displayed. Click **Back** to go to previous screens to make changes. Click **Finish** to complete the role creation.

When the newly created administrator logs in, unlike SYSMAN, the administrator is restricted by the privileges set.

Update Monitoring User of an Already Discovered SI Target

Follow these steps to update the monitoring user for a System Infrastructure Remote Host target that is already discovered:

1. Place the following command in the *sudoers* file on the remote host:

```
<non-root-username> ALL=(ALL) /usr/sbin/smif-ctrl, /usr/bin/ipmitool, /bin/cat
/etc/iscsi/iscsid.conf, /sbin/fdisk, /sbin/vgs, /sbin/pvs, /sbin/lvs, /sbin/
dmsetup, /usr/sbin/dbmcli, /sbin/service, /sbin/initctl, /opt/oracle.cellos/
compmon/exadata_mon_hw_asr.pl, /usr/local/bin/imageinfo, /usr/sbin/
exadata.img.hw, /usr/sbin/ibnetdiscover, /usr/sbin/smartctl
```

2. Use EMCLI to add an access point using a non-root user to the already discovered remote host. Specify the dispatch URL as `sudo-ssh://<TARGET_IP/HOST_NAME>:22` by running the following EMCLI command:

```
emcli add_target -name=<REMOTE_HOST_NAME>
-type=oracle_si_host_remote -host=<AGENT_HOST_NAME>
-
monitoring_cred=ssh_creds_set;oracle_si_host_remote_ssh;ssh_creds;username:
<USER_NAME>;userpass:<PASSWORD>;
-subseparator=properties== -properties=dispatch.url=sudo-ssh://<TARGET_IP/
HOST_NAME>:22
-access_point_name=<ACCESS_POINT_NAME> -
access_point_type=oracle_si_host_remote_ssh
```

 **Note**

Provide a `REMOTE_HOST_NAME` name which is different from the existing one. The new name is ignored and the already monitored target is automatically identified and attached to the access point. The new name will not be used.

3. Verify that the new access point is added to the existing System Infrastructure Remote Host target.
4. Once the new access point is verified to be running, remove the old access point which is using the root user for monitoring. Go to Access point overview page to remove the access point.

Exadata Database Machine Topology

Database Machine management simplifies monitoring and managing tasks by integrating all hardware and software components into one entity. You do not need to monitor each target individually, but instead you can view the whole Exadata Database Machine as a single target. You can view all critical issues in the system, monitor performance, and drill down to individual targets from the Database Machine target home page.

The following topology topics are presented in this section:

- [Drilling Down to Individual Targets of Database Machine](#)
- [Viewing Virtual Exadata Targets](#)
- [Viewing Critical Hardware Information for the Database Machine](#)
- [Viewing Exadata Resource Utilization](#)
- [Viewing Database Machine Alerts](#)
- [Viewing Metrics](#)
- [Adding Exadata Database Machine Components Manually](#)
- [Removing an Exadata Database Machine Target](#)
- [Removing a Component of a Database Machine Target](#)
- [Updating the Exadata Database Machine Schematic Diagram](#)

Use the Topology page of Database Machine to view the topology of the system by Cluster or by Database. Clusters are a complete software system starting with a RAC database, the

underlying ASM, and CRS. Clusters define one logical entity that is interconnected. The Database Machine could include several clusters, one cluster, or could just be a number of individual databases. While hardware racks define the hardware topology of the Database Machine, clusters define the logical or system topology of the Database Machine.

You can view the Topology by Cluster or Database. Click an element in the Topology and view alert data associated with the element.

You can monitor all components of the Database Machine. Database Machine monitors all subcomponent targets, whether hardware or software. This includes the database, ASM, CRS, hosts, Exadata and the InfiniBand network.

To view the topology of an existing Database Machine 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. From this page, you can add hardware components (such as Oracle Exadata Storage Servers, InfiniBand switches, Ethernet Switches, KVM switches, PDU, and compute node ILOM) in the Oracle Database Machine as managed targets.

2. From the Oracle Exadata Database Machines and Cloud Services page, select the Oracle Database Machine target whose topology you want to view.
3. From the Oracle Database Machine Home page, click the **Database Machine** menu, then select **Members** and then **Topology** from the drop-down menu.

Enterprise Manager Cloud Control displays the Configuration Topology page.

Drilling Down to Individual Targets of Database Machine

You can drill down immediately to a subcomponent target of the Database Machine (such as RAC, a database instance, or an Exadata Storage Server).

To drill down to individual targets:

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 Oracle Database Machine target whose components you want to view.

Enterprise Manager displays the Oracle Database Machine Home page showing an Overview, Schematic, and Incident section for the selected Database Machine.

3. From the Oracle Database Machine Home page, use the left navigation panel to expand the list of available targets that comprise the Database Machine.
4. Click the target to which you want to drill down.

Viewing Virtual Exadata Targets

Once discovered, the Exadata plug-in shows the virtual machines monitored by Enterprise Manager Cloud Control:



Note

The schematic diagram in the Database Machine home page is based on the content of the `databasemachine.xml` file found during discovery. The **virtual platforms** (`Dom0`) are displayed as compute nodes in the rack in the schematic diagram.

The Database Machine Software topology diagram will not display the physical Oracle Server, virtual Oracle Server targets (`DomU`), and Virtual Platform target (`Dom0`) targets. However, it will continue to show the host targets which are running in `DomU`.

The Software tab for the Exadata Database Machine target shows all clusters, ASM, and Database targets in the whole physical Database Machine grouped by clusters:

Target Name	Target Type	Status	Target Version
DB Machine	Oracle Exadata Database Machine		
+ASM	Cluster ASM		
cdb01	Cluster Database		19.4.0.0.0
Instances			
cdb01_cdb011	Database Instance		19.4.0.0.0
cdb01_cdb012	Database Instance		19.4.0.0.0
Pluggable Databases			
cluster1	Cluster		19.0.0.0.0

Viewing Critical Hardware Information for the Database Machine

You can view critical metrics for all the hardware subcomponents of the Database Machine such as DB hosts, Exadata Storage Servers, InfiniBand switches and so on. These metrics vary for different component targets. For example, database server nodes and Exadata servers include the CPU, I/O, and storage metrics.

To view critical hardware-centric information for the entire Database machine:

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

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

- From the Oracle Exadata Database Machines and Cloud Services page, select the Oracle Database Machine target whose hardware information you want to view.

3. From the Oracle Database Machine Home page, view the hardware schematic of the Database Machine.

Viewing Exadata Resource Utilization

Access the utilization graphs of compute and storage resources for cluster placement, database placement, database CPU utilization, database memory utilization, database I/O utilization, and disk group space allocation. From the Oracle Exadata Database Machine home page, select the type of resources and view the graphs accordingly.

- [Compute Resources](#)
 - [Cluster Placement](#)
 - [Database Placement](#)
 - [Database CPU Utilization](#)
 - [Database Memory Utilization](#)
- [Storage Resources](#)
 - [Exadata Storage Servers](#)
 - [Exadata Storage Server Grid](#)

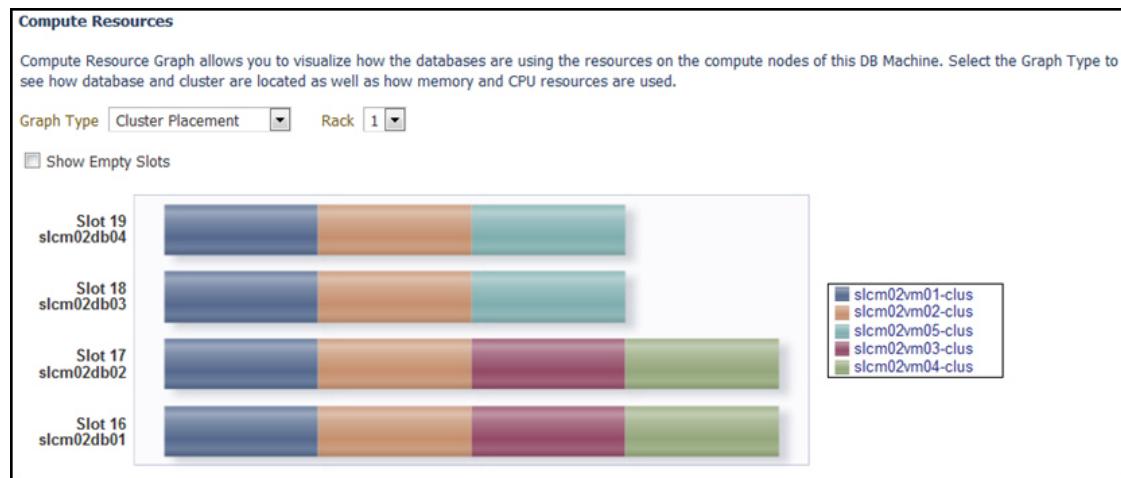
Compute Resources

From Oracle Exadata Database Machine home page, click the Database Machine drop down > select Resource Utilization > click Compute Resources.

The following compute resource allocation graphs are available in virtualized Exadata. These graphs are dependent on the virtual machine hierarchy and metric data from the VI plug-in:

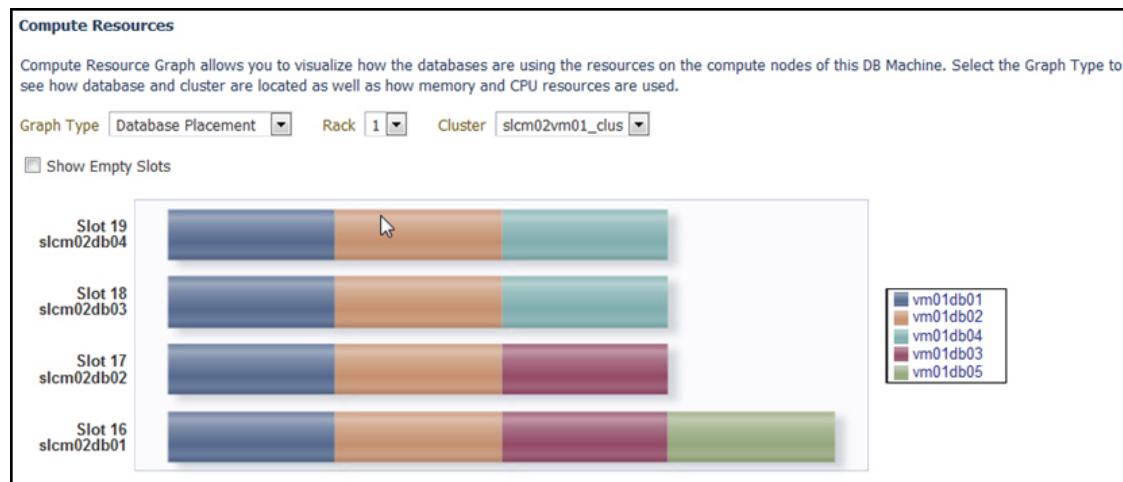
Cluster Placement

This graph shows the ClusterWare cluster placement on physical servers in a particular Exadata Database Machine rack. Since this is a placement graph, the widths of the data series reflect the number of clusters on the physical server that has the most number of clusters.



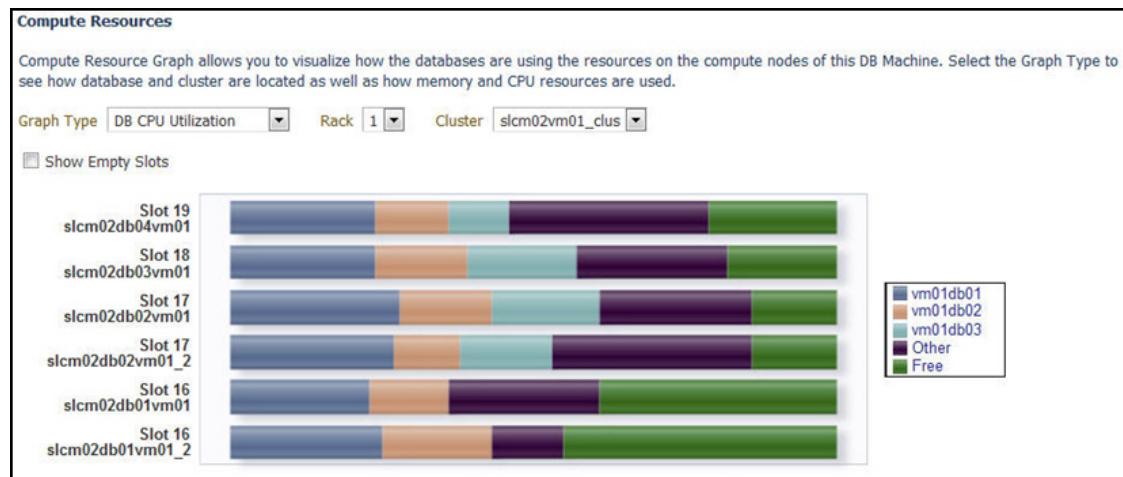
Database Placement

This graph shows the database placement on physical servers in a particular Exadata Database Machine rack for a particular DB cluster. Since this is a placement graph, the widths of the data series reflect the number of DB on the physical server that has the most number of databases for a particular DB cluster.



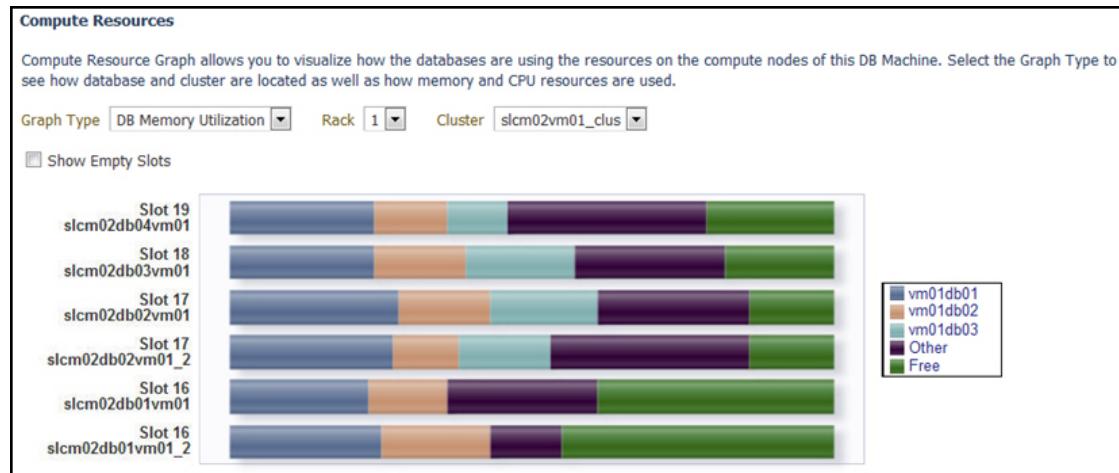
Database CPU Utilization

This graph shows the database CPU utilization per database per VM host for a particular DB cluster.



Database Memory Utilization

This graph shows the database memory utilization per database per VM host for a particular DB cluster.



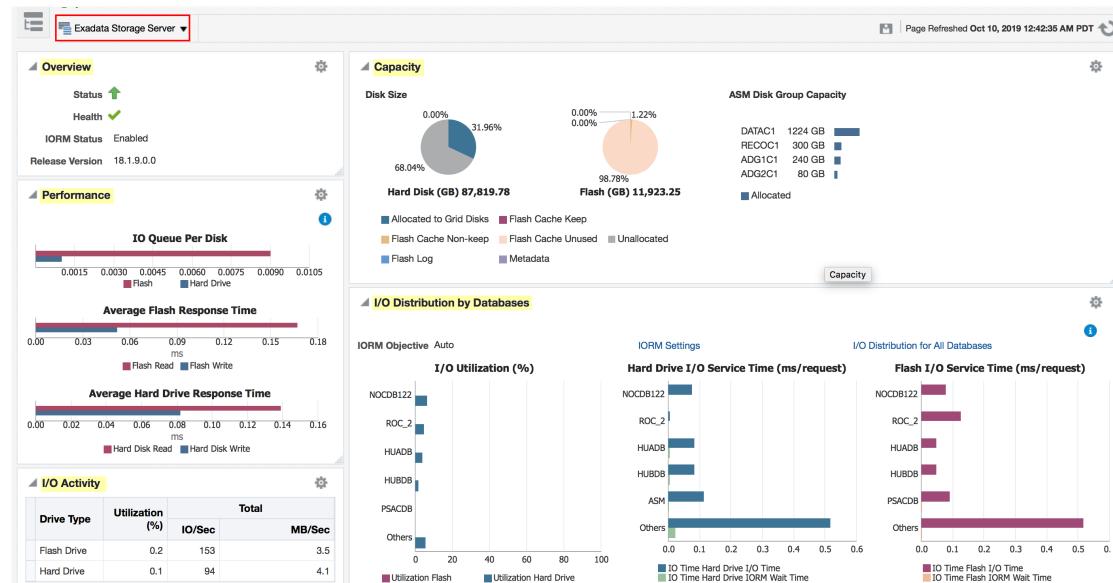
Storage Resources

From Oracle Exadata Database Machine home page, click the Database Machine drop down > select **Resource Utilization** > click **Storage Resources**.

The following storage resource allocation graphs are available in virtualized Exadata for Exadata Storage Servers and Exadata Storage Server Grid:

Exadata Storage Servers

Exadata Storage Server home page has the following sections that provide graphs and detailed information about status and performance of the target:



- **Overview:** Provides details of Exadata Storage Server availability status, health details, I/O Resource Manager status, and server version
- **Performance:** Provides hard drive and flash response time to IO requests, number of I/O's in the queue averaged per disk type.

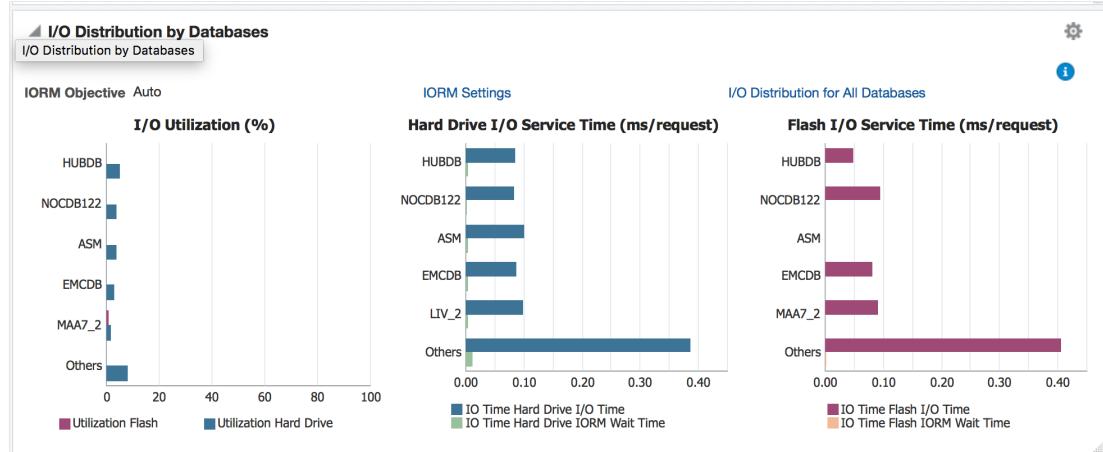
- **Capacity:** Provides capacity details of different storage types available on the storage server
 - Total available capacity of Hard drive and Flash
 - Flash Cache space information about keep objects, non-keep objects, unused space. This section also provides information about the size of the flash used for logging.
 - Capacity information for ASM disk groups
- **I/O Activity:** Provides performance information like Utilization, Latency, and Throughput for that available storage drives. This section also provides Flash Cache read and write, Smart Scan I/O throughput, InfiniBand Network throughput, Smart Log Efficiency, and IORM Boost.
- **I/O Distribution by Database:** Displays information about databases that are ranked top 5 in I/O Utilization for the available storage drives. *other* is an aggregation of remaining databases. This section also displays the Hard Drive I/O Service Time and Flash I/O Service Time for top 5 databases and *other* databases.
- **Incidents and Problems:** Lists alerts for the target

Database I/O Utilization

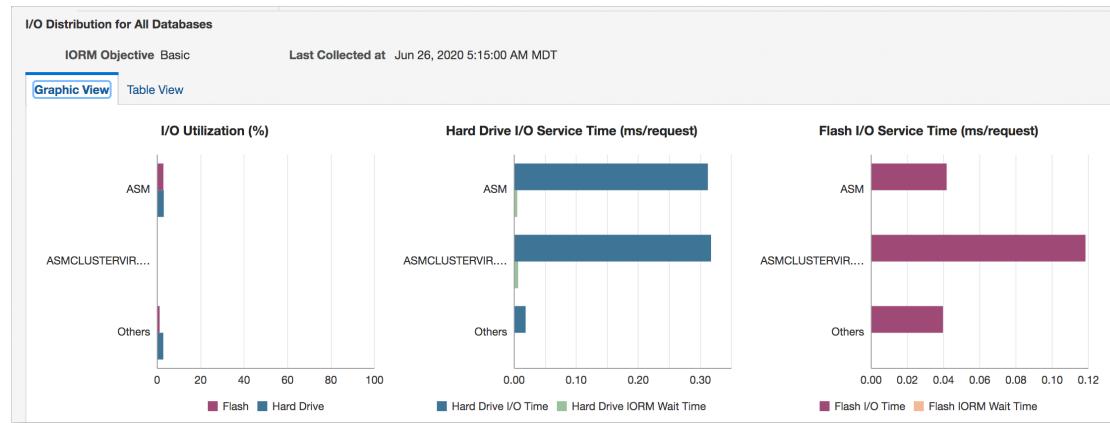
The Database I/O utilization graph shows the database I/O utilization per Exadata Storage Server of this Exadata Database Machine.

To access the Database I/O Utilization graph, from the Exadata Storage Servers home page, go to the section **I/O Distribution by Databases** > click **I/O Distribution for All Databases**.

I/O Distribution by Databases:



I/O Distribution for All Databases - Graphic View:



I/O Distribution for All Databases - Tabular View:

I/O Distribution for All Databases						
Database	I/O Utilization (%)		Hard Drive I/O Service Time (ms/request)		Flash I/O Service Time (ms/request)	
	Hard Drive	Flash	Hard Drive I/O Time	Hard Drive IORM Wait Time	Flash I/O Time	Flash IORM Wait Time
ASM	3.008	2.8427	0.312	0.005	0.0417	0
ASMCLUSTERVER.SLCM0801	0.01	0.0193	0.317	0.0063	0.1183	0
Others	2.7617	1.1073	0.0183	0	0.0397	0

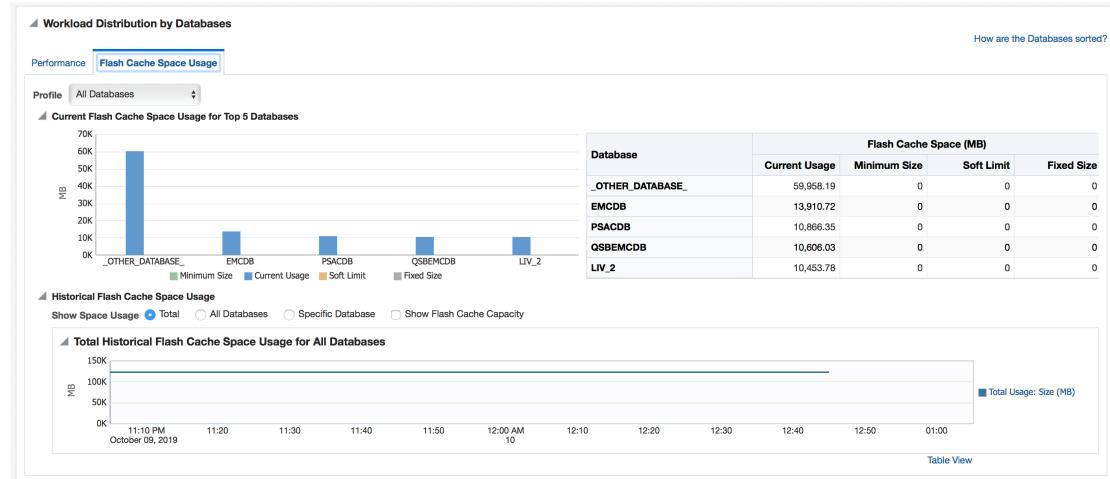
Workload Distribution by Databases

The Workload Distribution by Databases graphs show database workload performance on the hard drive and flash and the flash cache space usage.

To access the Workload Distribution by Databases graphs, from the Exadata Storage Servers home page, go to the section **I/O Distribution by Databases** > click **IORM Settings**.

To access the Flash Cache Space Usage graph, click the **Flash Cache Space Usage** tab.

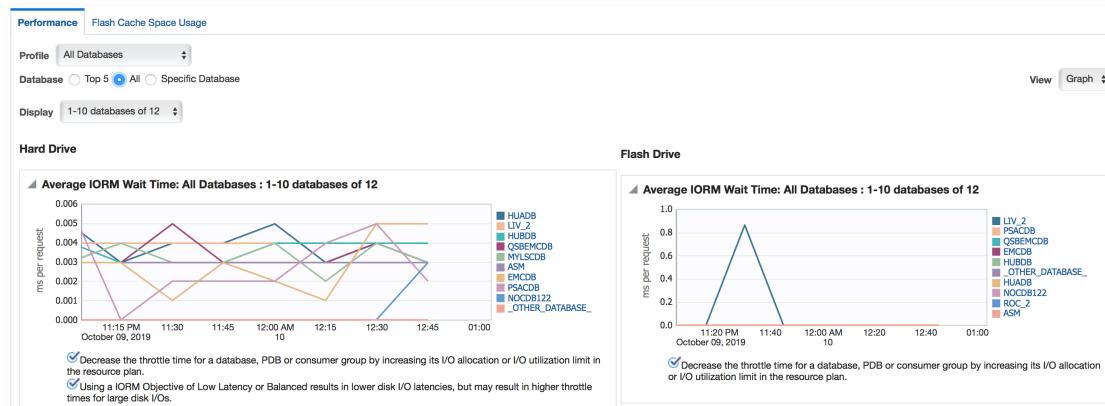
Flash Cache Space Usage:



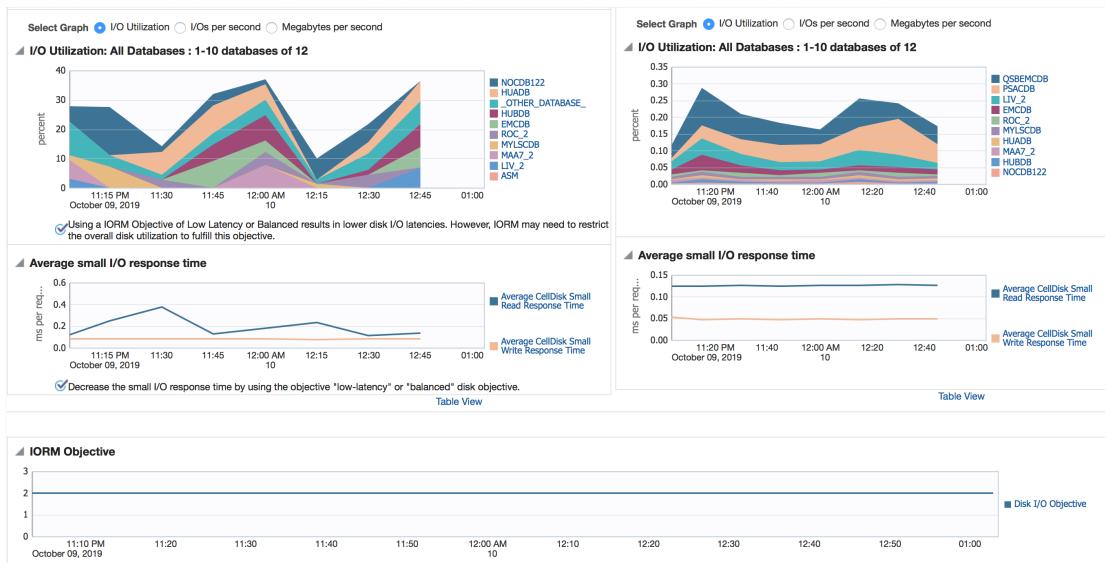
To view the graphs that correspond to the performance of hard drives and Flash, click the **Performance** tab.

This section provides database workload performance on Hard Drive and Flash. You can filter databases based on their workload performance rankings. You can view the graphs for the filters like top 5 databases, rank based pagination, or a specific database. You can drill down from database workload performance monitoring to pluggable database and consumer groups. Following are some of the example graphs:

IORM Wait Time for All Databases (Hard Drive and Flash Drive):



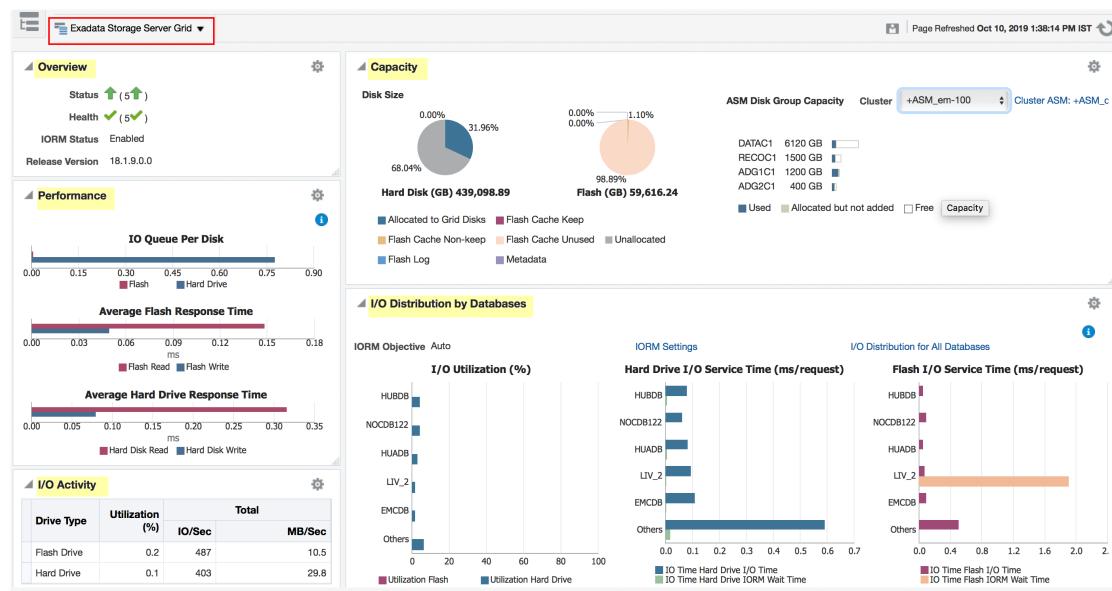
I/O Utilization and Average small I/O response time for All Databases (Hard Drive and Flash Drive):



Exadata Storage Server Grid

Exadata Storage Server Grid home page displays the aggregated status and performance data for all the Exadata Storage Servers. The sections are similar to those that are listed for individual Exadata Storage Server targets.

An example Exadata Storage Server Grid home page:



In the **Capacity** section of the Exadata Storage Server Grid home page, the ASM disk group capacity is displayed for the selected cluster.

You can view the I/O Utilization graph and Workload Distribution of Databases graph by clicking on the links **I/O Distribution for All Databases** and **IORM Settings** in the section **I/O Distribution by Databases**. The graphs that are displayed are the aggregated information from the Exadata Storage Server member targets of the grid. For details of the graphs, see [Database I/O Utilization](#) and [Workload Distribution by Databases](#).

Viewing Database Machine Alerts

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

To view Database Machine alerts:

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

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

- From the Oracle Exadata Database Machines and Cloud Services page, select the Oracle Database Machine target whose machine configuration information you want to view.

Enterprise Manager displays the Oracle Database Machine home page on which you can see all alerts associated with the current DB Machine.

Viewing Metrics

Metrics can be viewed at the component level of Exadata Database Machine. To view the Exadata Storage Server metrics:

- Navigate to the Oracle Exadata Storage Server home page by choosing the Oracle Exadata Storage Server target type from the All Targets page of Enterprise Manager.

Enterprise Manager displays the Exadata Storage Server Home page for the target you selected.

2. From the Exadata Storage Server drop-down menu, choose **Monitoring** and then **All Metrics**.

The All Metrics page allows you to view a wide variety of Exadata metrics. For a complete description of the Exadata metrics available, see the *Oracle® Enterprise Manager Oracle Database and Database-Related Metric Reference Manual*.

Adding Exadata Database Machine Components Manually

You can add Exadata Database Machine components manually using the following steps:

1. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**.
Enterprise Manager displays the Add Targets Manually page where you can choose the type of target you want to add.
2. From the Add Targets Manually section, choose **Add Using Guided Process**.
3. From the Add Using Guided Process box, choose **Oracle Exadata Database Machine** and click **Add**.
4. From the Oracle Exadata Database Machine Discovery page, click on the option **Discover newly added hardware components in an existing Database machine as targets**.
From the Select a Database Machine drop down list, select the Database Machine that has the new components and then click **Discover Targets**.
5. Choose the monitoring agent using the search option.
6. Provide the proper Infiniband Switch Credentials and click **Next**.
7. A prerequisite check is run. Click **Next**.
8. A list of available components for this Database Machine will be displayed. Select the new component(s) that need to be added and click **Next**.
9. Select the monitoring agent and backup monitoring agent for the new component(s).
10. On the Credentials page, set the credentials for the new components and click **Next**.
11. On the Review page, review all of the information and then click **Submit**.

Removing an Exadata Database Machine Target

To remove an Exadata Database Machine target, you can

- Remove all members of the Oracle Exadata Database Machine. If the member target is shared with another Database Machine target, then the member target will not be deleted and will continue to be monitored. In other words, the member targets will be deleted if they are associated with only this Database Machine target.
- Remove only system members of the Oracle Exadata Database Machine. The other member targets will not be deleted and will continue to be monitored. They can be associated to another Oracle Exadata Database Machine, if required.

Note

You must remove the SNMP configuration on the InfiniBand switch, ILOM, Cisco switches (RoCE and Admin), and PDU **manually** if you remove these targets from Enterprise Manager.

SNMP unsubscription is performed automatically for the targets like Oracle Exadata Storage Servers and Oracle Exadata Compute Nodes based on the options selected.

To remove an Exadata Database Machine target:

1. From the Database Machine target menu, select **Target Setup**, then select **Remove Target**.
2. On the Remove page, select whether to remove all members or to remove only the system members.

If you choose to remove all members, then you also have an additional option to unsubscribe SNMP subscriptions for member Oracle Exadata Storage Servers and Oracle Exadata Compute Nodes. Specify the appropriate user credentials to unsubscribe the SNMP subscriptions.

3. Click **Yes** to initiate the remove job.

Note

Host targets for the compute nodes and any targets that are also member targets of another Oracle Exadata Database Machine target will not be removed. System and non-system targets include:

System Targets:

- Oracle Exadata Database Machine
- Oracle Infiniband Fabric (*Enterprise Manager 13c target*)
- Oracle Exadata Storage Server Grid

Non-System Targets:

- Oracle Exadata Storage Server
- Oracle Exadata KVM
- Systems Infrastructure Switch
- Systems Infrastructure PDU
- Systems Infrastructure Rack
- Oracle Infiniband Switch (*Enterprise Manager 13c target*)
- Oracle Engineered System Cisco Switch (*Enterprise Manager 13c target*)
- Oracle Engineered System PDU (*Enterprise Manager 13c target*)
- Oracle Engineered System ILOM Server (*Enterprise Manager 13c target*)

Removing a Component of a Database Machine Target

If you need to remove a component of an Exadata Database Machine target, you can perform this task within Enterprise Manager Cloud Control 13c:

1. From the Database Machine home page's Target Navigation pane, right-click on the target item you wish to remove.
2. Select **Target Setup**, then **Remove Target**.
3. A warning page will display to confirm the target deletion. Click **Yes** to continue.

Updating the Exadata Database Machine Schematic Diagram

The schematic diagram illustrates the placement and status of the components of the Exadata Database Machine.

In some cases, the Exadata Database Machine schematic diagram is not displaying the components correctly. For example:

- You may have successfully discovered the Exadata Database Machine, but some components are not displaying correctly in the Exadata schematic diagram. Instead, an empty slot is shown in place of the component.
- The Exadata Database Machine schematic diagram shows the status of the component as "red/down" where as individual components would really show that they are up and running fine.
- You want to re-locate or rearrange the order of the components in the slots of Exadata Database Machine schematic diagram.

To accomplish these tasks, you will need to drop a component from the schematic diagram and add the correct one:

- [Drop a Component from the Exadata Database Machine Schematic Diagram](#)
- [Add Components to the Exadata Database Machine Schematic Diagram](#)

Drop a Component from the Exadata Database Machine Schematic Diagram

To drop a component from the Exadata Database Machine schematic diagram:

1. From the Targets menu, select **Exadata**.
2. From the Oracle Exadata Database Machines and Cloud Services page, select the Oracle Database Machine target that you want to modify.

Enterprise Manager displays the Oracle Exadata Database Machine home page of the selected Database Machine.

3. On the Exadata Database Machine schematic diagram right click on the component you wish to remove, click **Remove Target**.
4. Ensure you have selected the correct component in pop-up dialog box and click **OK**.
5. The Exadata Database Machine schematic diagram will refresh to show the empty slot.

Add Components to the Exadata Database Machine Schematic Diagram

To add a component from the Exadata Database Machine schematic diagram:

1. From the Targets menu, select **Exadata**.
2. From the Oracle Exadata Database Machines and Cloud Services page, select the Oracle Database Machine target you want to modify.
Enterprise Manager displays the Oracle Database Machine home page of the selected Database Machine.
3. On the Exadata Database Machine schematic diagram, select the **Photo-Realistic** view option.
4. Right-click on an empty slot of the photo-realistic diagram to display option the **Place Target**. A pop-up window **Place Target into Rack** is displayed. In the pop-up window, select the target to place it in the rack.
5. Once you see the component added in the slot you specify, click **Done** on the Exadata Database Machine schematic diagram.

Provisioning Virtualized Exadata Clusters

Provisioning involves repeatable, reliable, automated, unattended, and scheduled mass deployment of a RAC Cluster including virtual machines (VMs), Oracle Database (DB), Grid Infrastructure, and ASM on Virtualized Exadata.

 **Note**

Exadata Virtualized Provisioning is only available for virtual Exadata configured with InfiniBand networking.

The Exadata plug-in's virtualization provisioning functionality provides the following capabilities:

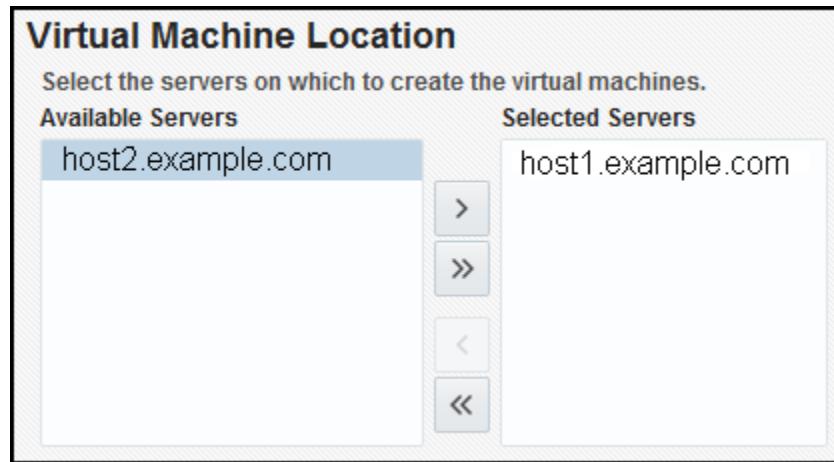
- [Creating a Database Cluster](#)
- [Scaling Up a Database Cluster](#)
- [Scaling Down a Database Cluster](#)
- [Deleting a Database Cluster](#)

Creating a Database Cluster

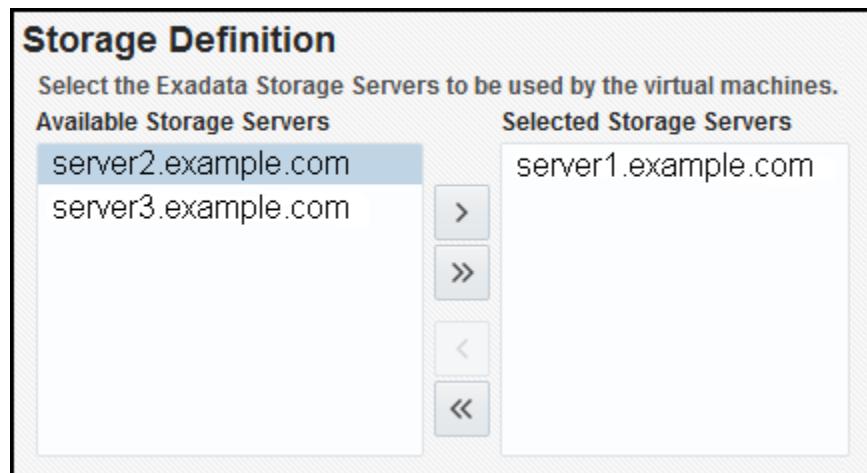
To create a database cluster, select the servers on which you want to create the virtual machines (VM), and select the Exadata Storage servers that can be used by the virtual machines. Provide additional information like the credentials to access them, VM definition, software locations, operating system users and groups, and networking information to complete the setup.

1. From the Database Machine target menu, select **Provisioning**, then select **Create Cluster**.
The Exadata Provisioning Wizard will be displayed.
2. On the Exadata Provisioning: Cluster page, provide the information for:
 - **Cluster Definition:** Create a cluster name. Click **Show existing clusters** to display a list of database clusters already provisioned.

- **Virtual Machine Location:** Select the servers on which you want to create the virtual machines. Select one or more from the **Available Servers** pane, then click the move button to move the selected server to the **Selected Servers** pane:



- **Storage Definition:** Select the Exadata Storage Servers to be used by the virtual machines. Select one or more from the **Available Storage Servers** pane, then click the move button to move the selected server to the **Selected Storage Servers** pane:



Click **Next**.

3. On the Credentials page, set the credentials for:

Exadata Provisioning: Credentials

Compute Server (Dom0)
Specify root credential for each Server.

Set Credential		Test Credentials
Server Name	Credential	Test Credential Status
Compute_server1	Not Defined	Credential not set
Compute_server2	Not Defined	Credential not set

Exadata Storage Server
Specify cell administrator credential for each Exadata Storage Server.

Set Credential		Test Credentials
Cell Name	Credential	Test Credential Status
Storage_server1	Not Defined	Credential not set
Storage_server2	Not Defined	Credential not set
Storage_server3	Not Defined	Credential not set

Specify agent host credential for each monitoring agent

Set Credential		Test Credentials
Agent URL	Credential	Test Credential Status
https://admin01vm01.example.com:20001/emd/main/	Not Defined	Credential not set

- **Compute Server:** Set the credentials for the `root` user. From the Set Credentials drop-down, select **All** to apply the settings to all servers, or select one or more servers from the list and select **Selected** to apply the settings to only the selected servers.

In the Set Server Credential pop-up window, select an existing named credential or select **New** to create a new one.

Click **Test Connection** to verify that the credentials are properly set. If successful, the Test Credential Status will update to show `Succeeded`.

- **Exadata Storage Server:** Set the credentials for the Exadata Storage Server administrator.

In the Set Server Credential pop-up window, select an existing named credential or select **New** to create a new one.

Click **Test Connection** to verify that the credentials are properly set. If successful, the Test Credential Status will update to show `Succeeded`.

- **Agent:** If Exadata Storage Server is monitored using RESTful API, then you must provide the agent host users and passwords for all the agents. From the Set Credentials drop-down, select an existing named credential or select **New** to create a new one.

Click **Test Connection** to verify that the credentials are properly set. If successful, the Test Credential Status will update to show `Succeeded`.

Click **Next** to define the Virtual Machines.

4. On the Virtual Machines (VMs) page, provide the following information:

① Note

When available, the Exadata Provisioning Wizard will pre-populate the fields with the appropriate defaults. You can change this information to suit your environment.

- **Virtual Machine Definition:** In this region, provide the details to define the VM:
 - Virtual Machine Size: Select Small, Medium, or Large. Click Customize to adjust the number of CPU cores and available memory.
 - Root password: Create a password in the two password fields.
 - Verify the Prefix, DNS, and NTP fields.
 - Time Zone: select the appropriate time zone from the drop-down list.

Virtual Machine Definition

Virtual Machine Size	Small	CPU cores: 2	Memory: 16GB	Customize...												
Available resources		CPU cores: 12	Memory: 41GB	Show details...												
<p>* Root password</p> <p>* Confirm root password</p> <p>Prefix s1cm02</p> <table border="1"><tr><td>DNS</td><td>11.22.33.444</td><td>11.22.33.555</td><td>11.22.33.666</td></tr><tr><td>NTP</td><td>10.200.3.4</td><td>10.200.3.5</td><td>10.200.3.6</td></tr><tr><td>* Time Zone</td><td colspan="3">(UTC-07:00) Denver - Mountain Time (MT)</td></tr></table>					DNS	11.22.33.444	11.22.33.555	11.22.33.666	NTP	10.200.3.4	10.200.3.5	10.200.3.6	* Time Zone	(UTC-07:00) Denver - Mountain Time (MT)		
DNS	11.22.33.444	11.22.33.555	11.22.33.666													
NTP	10.200.3.4	10.200.3.5	10.200.3.6													
* Time Zone	(UTC-07:00) Denver - Mountain Time (MT)															

- **Software Locations:** When available, the Exadata Provisioning Wizard will pre-populate the fields; otherwise, provide the following information:
 - Inventory Location
 - Grid Infrastructure Home
 - Database Home Location
 - Agent Installation Base Directory
 - Agent Port (optional)
 - Software Language (optional)

Software Locations

* Inventory Location	/u01/app/oralInventory	
* Grid Infrastructure Home	11.2.0.4 BP13	/u01/app/11.2.0.4/grid
* Database Home Location	11.2.0.4 BP13	/u01/app/oracle/product/11.2.0.4/dbhome_1
* Agent Installation Base Directory	/u01/app/oracle/agent	
Agent Port	3872	
Software Language	en	

- **Operating System Users and Groups:** Like the Software Location region, the Exadata Provisioning Wizard will pre-populate the fields except for the password fields. Otherwise, provide the following information:
 - User name: including the ID, password, and home directory.
 - DBA group name and ID.
 - OINSTALL group name and ID.

Operating System Users and Groups

* User name	oracle	* ID	1001	* Password		* Confirm Password		* Home directory	/home/oracle
* DBA Group Name	dba	* ID	1002						
* OINSTALL Group Name	oinstall	* ID	1003						

Click **Next** to provide the Network details.

5. On the Network page, specify the IP address, name and domain used for the Admin, Client and Private network. Provide the details for the following information:
 - **Gateway and Domain Details:** The domain and subnet mask details should already be supplied by the Exadata Provisioning Wizard. Enter a valid IP address for the Admin and Client gateway.
If a private network is used, then select the **Custom** button and specify the gateway. Typically it is *eth0*, *eth1*, or *bondeth0*.
If the Admin or Client or both networks are configured to use VLAN tagging, then the VLAN IDs are detected and displayed.
 - **Virtual Machine 1:** For the first virtual machine in the cluster, the Exadata Provisioning Wizard will auto-complete the Prefix and Start ID field. You can enter an optional Suffix. Enter a valid IP address.
 - **Generated Virtual Machine Name and IP:** After you have entered the information for the first VM, click **Generate** to create the information for the other VMs in the cluster. You can always enter the details manually.

Generated Virtual Machine Name and IP

Virtual Machine 1 From slcm02db06.us.oracle.com

Admin Network Name	slcm0202.example.com	IP	255.0.0.1
Client Network Name	slcm0204.example.com	IP	255.0.0.3
VIP Name	slcm0206.example.com	IP	255.0.0.5
Private Network Name	slcm0208.example.com	IP	255.0.0.7

Once you have entered all the necessary information, click **Validate IP** at the top of the page to verify the IP addresses. Click **Next** to continue to the enter the details for Grid Infrastructure and to create the Initial Database.

6. On the Grid Infrastructure and Initial Database page, enter the details for:
 - Grid Infrastructure:** Enter the details for the Cluster (SCAN name and port) and verify the Disk Group information. The Exadata Provisioning Wizard will pre-populate the details, but you can adjust them as needed.
 - Initial Database:** Click the check box to **Create Initial Database**. Additional information for Database Identification and Administrator Credentials will be required. Click **Advanced** to expand the region for additional details for Memory Parameters, Processes, and Character Sets:

Advanced

Memory Parameters	Processes
Memory Management: Automatic Shared Memory Management	Specify processes parameters.
<input checked="" type="checkbox"/> Specify Memory Settings as Percentage of Available Memory	Processes: 150
Percentage of Available Memory: <input type="text" value="40"/>	
Total Physical Memory(MB): 10384	Character Sets
Total SGA(MB): 4914	Database Character Set: Use Default - The default character setting based on the language setting of the operating system
Total PGA(MB): 1638	National Character Set: AL16UTF16 - Unicode UTF-16 Universal Character Set

Click **Next** to set the schedule.

7. On the Schedule page, the Exadata Provisioning Wizard will create the Deployment Instance value. Select a schedule start and notification options:
 - Schedule:** Select to initiate the creation immediately or later. If you select later, then you will be prompted to select a date and time.
 - Notification:** Select the notification statuses for which you will be notified.

Click **Review** to review the settings and initiate the job.

8. On the Review page, review the selection in the summary displayed. To change any section, return to the previous page and edit the selection.

Click **Submit** to begin the creation job.

Scaling Up a Database Cluster

You can scale up the compute ability of your database cluster by creating and adding more virtual machines to the cluster.

1. From the Database Machine target menu, select **Provisioning**, then select **Scale Up Cluster**.

The Exadata Provisioning Wizard will display.

2. On the Cluster page, enter the cluster name or click the Search icon to select a cluster from the list.

Select one or more from the **Available Servers** pane, then click the move button to move the selected server to the **Selected Servers** pane.

Click **Next** to set the credentials.

3. On the Credentials page, set the credentials for:

- **Compute Server:** Set the credentials for the `root` user. From the Set Credentials drop-down, select **All** to apply the settings to all servers, or select one or more servers from the list and select **Selected** to apply the settings to only the selected servers.

In the Set Server Credential pop-up window, select an existing named credential or select **New** to create a new one.

Click **Test Connection** to verify that the credentials are properly set. If successful, the Test Credential Status will update to show **Succeeded**.

- **Virtual Machines:** Set the credentials for the DomU Host, DomU Root, Cluster ASM, and Cluster Database.

You have the option to use preferred credentials or to override the preferred credentials.

Click **Test Connection** to verify that the credentials are properly set.

Click **Next**.

4. On the Virtual Machines page, click **Next**.

5. On the Network page, specify the IP address, name and domain used for the Admin, Client and Private network. Provide the details for the following information:

- **Gateway and Domain Details:** The domain and subnet mask details should already be supplied by the Exadata Provisioning Wizard. Enter a valid IP address for the Admin and Client gateway.

If a private network is used, then select the **Custom** button and specify the gateway. Typically it is `eth0`, `eth1`, or `bondeth0`.

If the Admin or Client or both networks are configured to use VLAN tagging, then the VLAN IDs are detected and displayed.

- **Virtual Machine 1:** For the first virtual machine in the cluster, the Exadata Provisioning Wizard will auto-complete the Prefix and Start ID field. You can enter an optional Suffix. Enter a valid IP address.
- **Generated Virtual Machine Name and IP:** After you have entered the information for the first VM, click **Generate** to create the information for the other VMs in the cluster. You can always enter the details manually.

Generated Virtual Machine Name and IP		
Virtual Machine 1 From slcm02db06.us.oracle.com		
Admin Network Name	slcm0202.example.com	IP 255.0.0.1
Client Network Name	slcm0204.example.com	IP 255.0.0.3
VIP Name	slcm0206.example.com	IP 255.0.0.5
Private Network Name	slcm0208.example.com	IP 255.0.0.7

Once you have entered all the necessary information, click **Validate IP** at the top of the page to verify the IP addresses.

Click **Next** to continue to schedule the scale up.

① Note

The step for Grid Infrastructure and Initial Database are skipped because they do not need to be set up again.

6. On the Schedule page, the Exadata Provisioning Wizard will create the Deployment Instance value. Select a schedule start and notification options:
 - **Schedule:** Select to initiate the scale up immediately or later. If you select later, then you will be prompted to select a date and time.
 - **Notification:** Select the notification statuses for which you will be notified.Click **Review** to review the settings and initiate the scale up.
7. On the Review page, review the selection in the summary displayed. To change any section, return to the previous page and edit the selection.
Click **Submit** to begin the scale up.

Scaling Down a Database Cluster

To scale down a database cluster, the Virtual Machine is removed from the cluster:

1. From the Database Machine target menu, select **Provisioning**, then select **Scale Down Cluster**.

The Exadata Provisioning Wizard will display.

2. Enter the cluster name you want to scale down or click the Search icon to select an available cluster.

Once you have selected a cluster, you will be prompted to:

- Select nodes to delete.
- Verify or enter new named credentials for DomU (host, root, and Exadata Server) and for the Exadata Storage Server. Click **Test Credentials** to verify the credentials have been selected properly.

Click **Next** to schedule the scale-down job.

3. On the Schedule page, the Exadata Provisioning Wizard will create the Deployment Instance value. Select a schedule start and notification options:
 - **Schedule:** Select to initiate the scale down immediately or later. If you select later, then you will be prompted to select a date and time.
 - **Notification:** Select the notification statuses for which you will be notified.
- Click **Review** to review the settings and initiate the scale down.
4. On the Review page, review the selection in the summary displayed. To change any section, return to the previous page and edit the selection.
- Click **Submit** to begin the scale down.

Deleting a Database Cluster

You can delete an existing database cluster and free the compute and storage resources used by the cluster. In the provisioning wizard, select the cluster to delete and provide the credentials to access the Exadata Storage Server and the virtual machine hosts.

1. From the Database Machine target menu, select **Provisioning**, then select **Delete Cluster**.

The Exadata Provisioning Wizard will display.
2. On the Cluster page, enter the cluster name you want to delete or click the Search icon to select an available cluster.

The page will update to show the nodes to be deleted

Verify or enter new named credentials for DomU (host, root, and Exadata Server) and for the Exadata Storage Server. Click **Test Credentials** to verify the credentials have been selected properly.

Click **Next** to schedule the delete job.
3. On the Schedule page, the Exadata Provisioning Wizard will create the Deployment Instance value. Select a schedule start and notification options:
 - **Schedule:** Select to initiate the scale down immediately or later. If you select later, then you will be prompted to select a date and time.
 - **Notification:** Select the notification statuses for which you will be notified.
- Click **Review** to review the settings and initiate the delete job.
4. On the Review page, review the selection in the summary displayed. To change any section, return to the previous page and edit the selection.
- Click **Submit** to begin the delete job.

Exadata Storage Server Metrics and Alert Settings

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

1. From the Enterprise Manager home page, select **All Targets**. Select an Exadata Storage Server.
2. 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.

3. Click **OK** to save any changes.

Exadata Storage Server Management

This section provides introductory instructions for managing Exadata Storage Servers. The following topics are presented:

- [About Exadata Storage Server](#)
- [Viewing an Exadata Storage Server Topology](#)
- [Performing Administration Tasks on Exadata Storage Servers](#)
- [Managing the I/O Resource](#)
- [Diagnosing Exadata Storage Server Alerts](#)
- [Accessing Oracle Support Workbench for Exadata Storage Server](#)
- [Removing a Component of a Database Machine Target](#)

About Exadata Storage Server

An Exadata Storage Server is a highly optimized storage server that runs Oracle Exadata System Software to store and access Oracle Database data. Use the Exadata Storage Server Home page to manage and monitor the Oracle Exadata Storage Server as an Enterprise Manager Cloud Control target. You can discover and consolidate management, monitoring and administration of a single or a group of Oracle Exadata Storage Servers in a datacenter using Enterprise Manager.

Note

As the Exadata Database Machine has evolved, there have been different names used to reference the Exadata Storage Server, including storage cell and cell server. Any reference with these names in this or related documentation refers to Exadata Storage Server.

Exadata Storage Servers can be discovered automatically or manually. The individual Exadata Storage Server is monitored and managed as an Enterprise Manager target and provides the exception, configuration and performance information.

Grouping of Exadata Storage Servers is used for easy management and monitoring a set of Storage Servers. You can group them both manually and automatically. The grouping function provides an aggregation of exceptions, configuration and performance information of the group of Exadata Storage Servers.

You can view performance analysis by linking Exadata performance both at the Exadata Storage Server level and group level to ASM and database performance. You can drill down to Exadata configuration and performance issues from both the database and ASM targets.

Storage Grid (for example, multiple database/ASM instances sharing the same Exadata Storage Server) is supported to the same extent as dedicated storage.

Viewing an Exadata Storage Server Topology

You can view the topology of an Oracle Exadata Storage Server target by following the steps below:

1. Navigate to the Exadata Storage Server home page by choosing the Exadata target for which you want to view the configuration.

Enterprise Manager displays the Exadata Storage Server Home page for the target you selected.

2. From the Target menu, choose **Configuration** and then **Topology**.

Enterprise Manager displays the Configuration Topology page for the selected Exadata Storage Server. The topology page provides a visual layout of the target's relationships with other targets. From this page you can:

- Do a target search filtered by target status/events/target type
- Select from a set of relationships to represent in the graph
- Select annotations to display in the graph, such as alerts and link labels
- Select from a set of options: view navigator, expand or collapse all, toggle graph layout, reload topology
- Print
- Zoom via the slide control
- Pan via the navigator control
- Toggle the presentation from graph to table

When you hover over a node or group member, a popup displays detailed information about the entity. A link can appear in the popup to more detailed information such as customer documentation.

Performing Administration Tasks on Exadata Storage Servers

Note

It is best practice to manage and configure Exadata Storage Servers consistently. Enterprise Manager provides the ability to perform Administration Tasks on all Exadata Storage Servers in an Exadata Database Machine from the corresponding Exadata Storage Server Grid target home page, or on individual Exadata Storage Server targets from the Exadata Storage Server target home page. To the maximum extent possible, consider managing from the Exadata Storage Server Grid target to manage the servers consistently, selecting a subset of servers if necessary. The instructions in this section can be used for either the Exadata Storage Server Grid or Exadata Storage Server targets.

To perform an administration operation on an Exadata Storage Server, such as executing a Exadata Storage Server command, follow these steps:

1. Navigate to the Exadata Storage Server home page by choosing the Exadata target for which you want to perform an administrative task from the All Targets page.

Enterprise Manager displays the Exadata Storage Server Home page for the target you selected.

2. Click **Exadata Storage Server, and then select **Administration**.**

From this menu you can choose either **Execute Cell Command**, or **Manage I/O Resource**.

3. Click **Execute Cell Command.**

The Command page of the Exadata Storage Servers Administration wizard appears. Enter a CELLCLI command as the administrative command to be executed on the Exadata Storage Server. You must read the Command Instructions before you enter a command. Only a single cellCLI command is allowed to execute. You must enter the command without the cellcli -e prefix, which is automatically appended when you submit the command. Finally, you cannot use the following characters: ; / ' < > / |.

4. Click **Next to continue.**

Enterprise Manager displays the Admin Credentials page. Select or enter the Administration credentials to execute the command. The credentials you enter are used when submitting the operation. You can choose between Preferred Credentials, Named Credentials, and New Credentials. You can also click **More Details** to view information about Credential Type, Last modified, Credential Name, Credential Owner, Last Modified Date, Last Modified By, and Preferred Credentials Set At.

5. Click **Next.**

Enterprise Manager displays the Schedule page. Use the Schedule page to schedule the administration task. Enter the Job Name and the Job Description, then provide the job information in the Schedule the Administration Job section. You can choose to begin the job immediately or enter the time you want the job to begin.

6. Click **Next to continue.**

The Summary page displays. Use the Summary page to ensure you have entered the correct values and then submit the command. The Summary page lists the Job Name, Description, Command to Execute, when the job is Scheduled, and the Selected Cell.

7. Click **Submit Command to submit the job.**

The Job Status page displays. Use the Job Status page to link to the Job Detail page of the administration task.

Managing the I/O Resource

Oracle Exadata Storage Servers are added as targets during the Exadata Database Machine discovery workflow (see [Exadata Database Machine Discovery](#)) and are grouped automatically under the group **Exadata Storage Server Grid**. Monitoring and Managing I/O Resource can be performed at Exadata Storage Server Grid and Exadata Storage Server target level. The information provided below is for Grid and Server target.

To access the IORM Performance page:

- 1. Select an Exadata Storage Server:**
 - a. From the Targets menu, select **Exadata**.**
 - b. Select a DB Machine from the list of Target Names.**
 - c. In the Target Navigation pane, expand the Exadata Grid item and click one of the Exadata Storage Servers.**

2. Once you have selected an Exadata Storage Server, click the **Exadata Storage Server** menu, select **Administration**, then **Manage I/O Resource**.

The IORM page has the following sections:

- **I/O Resource Manager Settings:** This section helps to view and update the IORM configuration of Exadata Storage Server. Use **Get Latest** button to view the latest Exadata Storage Server IORM configuration. Use the same section to edit and update the IORM configuration on Exadata Storage Server.
- **Database Resource Management(DBRM) Settings:** This section display the list of databases and their basic configuration, that are using Exadata Storage Server resources. From this page, you can launch the database resource management page.
- **Workload Distribution by Databases:** The IORM Monitoring section of the page provides a view of the performance statistics of Disk I/O (Wait, IOPS, MBPS, Utilization, Latency, and Objective charts). These statistics help to identify which databases and consumer groups are using the available resources. They also help to adjust the IORM configuration (using IORM Settings section on the same page) as needed.

IORM Configuration includes IORM Objective and Inter-Database Plan.

For further details on managing I/O resources, refer to the *Managing I/O Resources* chapter in the *Oracle® Exadata Storage Server Software User's Guide*.

Add/Update IORM Configuration

To update the I/O Resource Manager (IORM) settings (for Exadata Storage Server software release 12.1.2.1.0 and later):

1. Navigate to the IORM page as described above. The image shows the I/O Resource Manager (IORM) Settings pane.

Note

You can also update a single Exadata Storage Server. Expand the Exadata Grid group to view all Exadata Storage Servers associated with the group. Click the Exadata Storage Server you want to update.

The steps to update the IORM settings is the same for a single Exadata Storage Server or group of Exadata Storage Servers.

2. Update the inter-database plan by editing existing values or add/remove directives to Inter-database plan using the Add/Remove buttons available on Inter-database table.
 - a. From the Database Name column, select a database from the drop-down menu or enter a database name.
 - b. Enter a value for the Hard Disk I/O Utilization Limit column.
 - c. Enter a value for the Database I/O Share column.
 - d. Enter minimum and maximum values (in MB) for the Flash Cache column.
 - e. Fill the remaining attributes like Type and Role.
3. In the Disk I/O Objective drop-down menu, select an objective from the list (**Auto** is the default):

- **Low Latency** - Use this setting for critical OLTP workloads that require extremely good disk latency. This setting provides the lowest possible latency by significantly limiting disk utilization.
- **Balanced** - Use this setting for critical OLTP and DSS workloads. This setting balances low disk latency and high throughput. This setting limits disk utilization of large I/Os to a lesser extent than Low Latency to achieve a balance between good latency and good throughput.
- **High Throughput** - Use this setting for critical DSS workloads that require high throughput.
- **Auto** - Use this setting to have IORM determine the optimization objective. IORM continuously and dynamically determines the optimization objective, based on the workloads observed, and resource plans enabled.
- **Basic** - Use this setting to disable I/O prioritization and limit the maximum small I/O latency.

 **Note**

If the *Auto* and *Balanced* objectives are not applicable (i.e., they have no impact on I/Os in Extreme Flash Exadata Storage Servers), they will not appear in the IORM page.

Click **Update**. The Exadata Storage Server Administration Wizard will appear prompting you for the information necessary to complete the Disk I/O Objective configuration:

- a. On the wizard's **Command** page, the Cell Control Command-Line Interface (CellCLI) value should be as shown when the Disk I/O Objective selected is *auto*:

```
# alter iormplan objective = 'auto'
```

Click **Next**.

- b. On the wizard's **Admin Credentials** page, enter the user name and password for the selected Exadata Storage Servers.

Click **Next**.

- c. On the wizard's **Schedule** page, enter a job name (required) and job description (optional). Select an option to start **Immediately** or **Later**. If you select the *Later* option, enter the time that you want the job to run.

Click **Next**.

- d. On the wizard's **Review** page, verify that the settings are correct. To submit the command, click **Submit Command**.

After the job is successfully submitted, the Job Status page is displayed.

Click **Return** to return to the I/O Resource Manager (IORM) Settings pane.

4. Click **Get Latest** to refresh the IORM configuration from Exadata Storage Server, which will include the updated Disk I/O Objective.
5. To confirm the Disk I/O Objective setting, for example if it was selected as *auto*, run the following command from the command line:

```
# dcli -g cell_group cellcli -e "list iormplan attributes objective"
```

Output should show the value `auto`:

```
cell01: auto
cell02: auto
cell03: auto
.
.
.
cell14: auto
```

The Inter-Database Plan

An inter-database plan specifies how resources are allocated by *percentage* or *share* among multiple databases for each Exadata Storage Server. The directives (rows) in an inter-database plan specify allocations to databases, rather than consumer groups. The inter-database plan is configured and enabled with the CellCLI utility at each Exadata Storage Server. Only one inter-database plan can be active on a Exadata Storage Server at any given time.

Percentage based Inter-database plan: Use **allocations** to specify percentage for I/O distribution for each database. With Percentage based plan, IORM can manage up to 32 databases. You can allocate resources across 8 different levels. A percentage based inter-database plan is similar to a database resource plan, in that each directive consists of an allocation amount and a level from 1 to 8. For a given percentage based plan, the total allocations at any level must be less than or equal to 100 percent. An inter-database plan differs from a database resource plan in that it cannot contain sub-plans and only contains I/O resource directives.

Share based inter-database plan: This is the recommended method for Inter-database plans. Use share to specify the relative priority for a database. A higher share value implies higher priority and more guaranteed I/O resources. Each database is given a share value which is an integer between 1 and 32. Use the limit method to restrict the I/O utilization for a database to a specified utilization limit. The sum of the shares can be greater than 100. Share-based plans support up to 1024 directives within the Inter-database plan.

I/O Resource Management can provide predictable performance by guaranteeing space in Exadata Smart Flash Cache. For this use Flash Cache Space columns to specify space allocation for each database in Inter-database plan.

Specify the attributes values for each directive in Inter-database plan as per requirement.

Note

If the Exadata Storage Server version is 11.2.3.1.0 or earlier, the Share, Percentage based inter-database plan radio buttons are not available. You can view only Percentage-based options (that is, the drop-down only displays the Basic, Advance options).

When considering an inter-database plan:

- If Oracle Exadata Storage Server is hosting only one database, then an inter-database plan is not needed.
- If an inter-database plan is not specified, then all the databases receive an equal allocation.

Diagnosing Exadata Storage Server Alerts

Enterprise Manager listens for Exadata Storage Server alerts sent from the Exadata Storage Server Management Server; so, any hardware failure or Exadata Storage Server error will be reported in Enterprise Manager. For information about the error messages you might encounter on the Exadata Database Machine, refer to the [Error and Alert Messages](#) in *Oracle Exadata Database Machine Error Message Reference Guide*.

Managing the InfiniBand Fabric

All InfiniBand Switches are discovered automatically during the database machine discovery workflow (see [Exadata Database Machine Discovery](#)) and are grouped automatically under the group **IBFabric@<switch-name>**.

 **Note**

InfiniBand Fabric target is not available for RoCE Exadata.

1. From the Enterprise Manager home page, select **Targets**, then **Oracle Exadata Database Machines and Cloud Services**.
2. In the Target Navigation pane, select **InfiniBand Fabric** from the list.
3. In the IB Fabric pane, you can view an overview and activity summary for all InfiniBand Switches.
4. Click **Refresh** for an On Demand refresh of the InfiniBand schematic. Updates reflect the real-time data.

The following topics address managing your InfiniBand network:

- [InfiniBand/RoCE Switch Metrics](#)
- [Performing Administration Tasks on InfiniBand Networks](#)
- [Setting Up Alerts](#)

InfiniBand/RoCE Switch Metrics

The Enterprise Manager Agent runs a remote SSH and remote SNMP GET call to collect metric data for the InfiniBand switch. InfiniBand metrics provides operational details such as:

- Status / Availability
- Port status
- Vital signs: CPU, Memory, Power, Temperature
- Network interface various data
- Incoming traffic errors, traffic Kb/s and %
- Outgoing traffic errors, traffic Kb/s and %
- Administration and Operational bandwidth Mb/s

The following metrics are available for your InfiniBand Fabric:

Switch Aggregated Status

The Aggregate Sensor takes input from multiple sensors and aggregates the data to identify problems with the switch that require attention. Whenever the sensor trips into an "Asserted" state (indicating a problem) or "Deasserted" (indicating that the problem is cleared) for a component on the switch, associated Enterprise Manager events will be generated.

Response

This is the main metric indicating availability of the InfiniBand/RoCE switch. It is collected every 60 seconds by default through the management interface of the switch.

Switch Configuration

This metric captures the switch configuration. The information collected is valuable only to Oracle Support, which will use it to assist in debugging situations.

Switch Basic Status

This metric gives basic status of the switch like Booted on, Locator light status, Power status and overall status of the switch.

Sensor Status

This metric gives the status of various sensors available in the switch like power supply, fan, motherboard, and cooling.

Switch Port Statistics

This metric provides information on number the of incoming and outgoing errors, incoming and outgoing octets.

Component State

This metric gives the state of various components in the switch like Fan, Motherboard, Power Supply and various InfiniBand and Ethernet ports.

Network Port InfiniBand performance

This metric gives performance data of each InfiniBand port.

Performing Administration Tasks on InfiniBand Networks

Note

Administrative tasks are not allowed to be performed on RoCE switch.

To perform an administration operation on an InfiniBand Network, follow these steps:

1. Navigate to the Database Machine home page of the InfiniBand Network by choosing the Database Machine for which you want to perform an administrative task from the **All Targets** page.

Enterprise Manager displays the Database Machine Home page for the target you selected.

2. Navigate to **System Infrastructure Switch** home page for which you want to perform an administrative task.
3. Go to **Administration**, and select **Switch Operations**.
4. Select the administrative operation you want to execute (Enable/Disable port, Clear performance/Error counters, Switch LED on/off, Set up SNMP subscription). The available operations from which you can select are dependent on the target type and target you selected. Once you choose the operation, you may need to select a value that will appear after choosing the operation.
5. Click **Next** to continue.

Enterprise Manager displays the Credentials & Schedule page. Select or enter the credentials to execute the command. The credentials you enter are used when submitting the operation. You can choose between Preferred Credentials, Named Credentials, and New Credentials. Schedule the administration task. Provide the job information in the **Administration Job Schedule** section. You can choose to begin the job immediately or enter the time you want the job to begin.

6. Click **Next** to continue.

The Review page appears. Use the Review page to ensure you have entered the correct values and then submit the command. The Review page lists the Job Name, Description, Command to Execute, when the job is Scheduled, the Target Type, and the Selected Target.

7. Click **Submit Command** to submit the job.

When you click Submit Command, a popup is shown if the job is successful. You can go to the Job Detail Page or back to the page from where this wizard was launched.

Setting Up Alerts

After configuring the InfiniBand Switch targets to send SNMP alerts, set up alerts in Enterprise Manager Cloud Control.

1. Log in to Enterprise Manager Cloud Control.
2. Click **Targets**, then **All Targets**. All discovered targets will display.
3. In the *All Targets* page, click **Systems Infrastructure Switch**.
4. Click the target you are interested in. The target home page appears.
5. In the drop-down menu for the Systems Infrastructure Switch, select **Monitoring** and then **Metric and Collections Settings**.
6. In the *Metric and Collection Settings* page, you can modify metric threshold values, edit monitoring settings for specific metrics, change metric collection schedules, and disable collection of a metric.

You can modify the thresholds directly in the table or click the edit icon (pencil icon) to access the Edit Advanced Settings page. For more information on the fields displayed in this page and how the thresholds can be modified, click **Help** from the top-right corner of this page.

Discovering and Monitoring Exascale

Exascale is a collection of Exadata software services that manage Exadata storage servers and provide storage to multiple database server clusters. Enterprise Manager provides Exascale discovery and monitoring via a dedicated *Oracle Exadata Exascale* target type.

For more information on Oracle Exadata Exascale, see [What is Oracle Exadata Exascale?](#) in *Exascale User's Guide*.

Topics:

- [Prerequisites](#)
- [Discovering Exascale](#)
 - [Exascale Discovery as Part of Oracle Exadata Database Machine Discovery](#)
 - [Exascale Discovery After Oracle Exadata Database Machine Discovery](#)
- [Monitoring Exascale](#)
 - [Capacity](#)
 - [Serviced Clients](#)
 - [Performance](#)
 - [Incidents and Compliance](#)

The Exascale target home page provides monitoring capabilities such as availability and health monitoring, as well as storage and IOPS utilization monitoring, both aggregated and distributed by serviced clients.

Advanced performance monitoring consists of an interactive illustration of IOPS allocations across vaults, storage pools and databases. Advanced filtering options provide the ability to visually identify IOPS bottlenecks and view them in context with historical Exascale IOPS utilization in order to triage performance issues. The advanced monitoring capabilities are available if the Exadata Management Pack is enabled. See Exadata Exascale Advanced Monitoring in *Advanced Management and Monitoring for Engineered Systems*.

Prerequisites

- Exascale is supported from Exadata System Software version 24.1.
- Exascale is supported from Oracle Enterprise Manager version 13.5GC Release Update 23 (both Agent and Management Server)
- Exascale is supported for Oracle Database 23ai (23.5 or higher).
- The Exascale administrator should create a user with `cl_monitor` privilege using Exascale command line `escli`. The required private or public key is created using `escli` and saved in a PEM file. This key can also be stored in the digital wallet. Oracle Enterprise Manager requires this user name and corresponding private key to monitor Exascale. This information can be copied from the PEM file directly or extracted from the wallet. See [User Privileges](#), [Create a User](#), and [Start and Use ESCLI](#) in *Exascale User's Guide*.

Discovering Exascale

Enterprise Manager uses a dedicated target type *Oracle Exadata Exascale* to monitor Exascale. Exascale targets can be discovered either as part of Exadata Database Machine discovery or separately after Database Machine discovery.

Topics:

- [Exascale Discovery as Part of Oracle Exadata Database Machine Discovery](#)
- [Exascale Discovery After Oracle Exadata Database Machine Discovery](#)

Exascale Discovery as Part of Oracle Exadata Database Machine Discovery

1. Follow the steps *1 through 10* in [Discover the Oracle Exadata Database Machine Target Using the Console](#) for discovering a Database Machine and its components.
2. After providing credentials for monitoring Agents and storage servers, the discovery wizard checks for the presence of Exascale on the Database Machine. If present, the Storage Management page is displayed to enable Exascale discovery. The Exascale target name is prepopulated with the Exascale cluster name.

The screenshot shows the Oracle Enterprise Manager Cloud Control 13c interface. The title bar says 'ORACLE Enterprise Manager Cloud Control 13c' and 'SYSMAN'. The main title is 'Oracle Exadata Database Machine Discovery'. The progress bar shows 'Storage Management' is the current step. The sub-section is 'Database Machine Discovery: Storage Management'. The 'Oracle Exadata Exascale' section contains fields for 'Target Name' (set to 'IO5esclu1'), 'Target Type' (set to 'Oracle Exadata Exascale'), 'Control Server Endpoint', and 'Control Server Credential' (with a 'Set Credential' button). The 'Monitoring Details' section shows monitoring agent URLs: 'Monitoring Agent' (https://IO5adm01vm02:5501) and 'Backup Monitoring Agent' (https://IO5adm02vm02:5501). There are 'Back', 'Step 7 of 8', 'Next', and 'Cancel' buttons at the bottom.

Provide the **Monitoring Agent** and **Backup Monitoring** Agent to monitor the Exascale target.

3. Provide the **Control Server Endpoint**, as provided by the Exadata administrator. This will be in the format *<fully qualified hostname>:port*.
4. Specify the **Control Server Credential** using the name and private key used in [Create an Exascale Monitoring User](#).

Copy the private key you previously generated and paste into the **Private Key** and **Confirm Private Key** fields. Click **OK** to save. Click **Next** to proceed.

5. Continue with steps *11 through 12* in [Discover the Oracle Exadata Database Machine Target Using the Console](#).

You can access the Exascale target from the Target Navigation tree. The Exascale target can also be accessed from the All Targets page in the **Servers, Storage and Network** category.

Exascale Discovery After Oracle Exadata Database Machine Discovery

If you want to add Exascale to a Database Machine after the latter is discovered by Oracle Enterprise Manager, then Exascale can be discovered as a newly added component and associated with the existing Database Machine target. In the Database Machine discovery workflow, select the **Discover newly added components in an existing Database Machine as targets** option. This is the same workflow used to discover other new components, such as storage servers.

Monitoring Exascale

The Exascale home page has four tabs for monitoring:

Topics:

- [Capacity](#)
- [Serviced Clients](#)
- [Performance](#)
- [Incidents and Compliance](#)

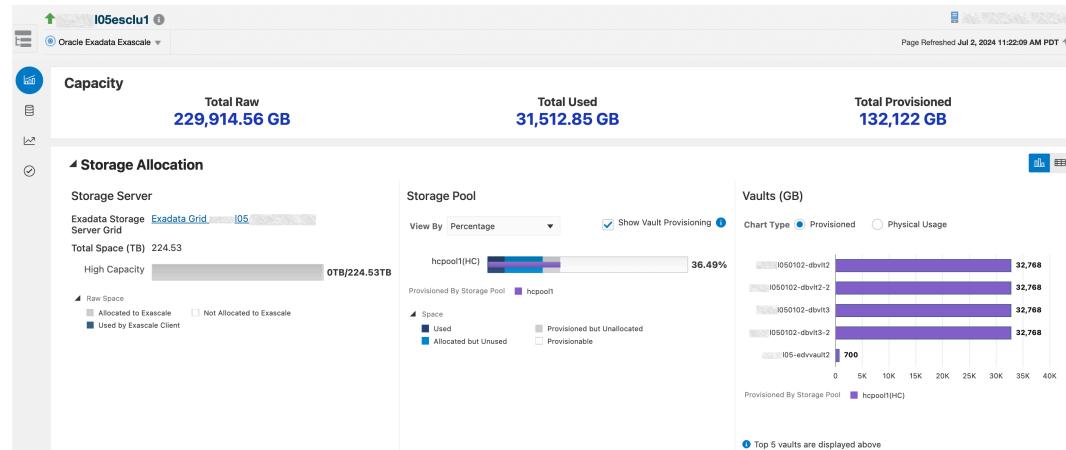
The advanced monitoring capabilities are available if the Exadata Management Pack is enabled. See Exadata Exascale Advanced Monitoring in *Advanced Management and Monitoring for Engineered Systems*.

Capacity

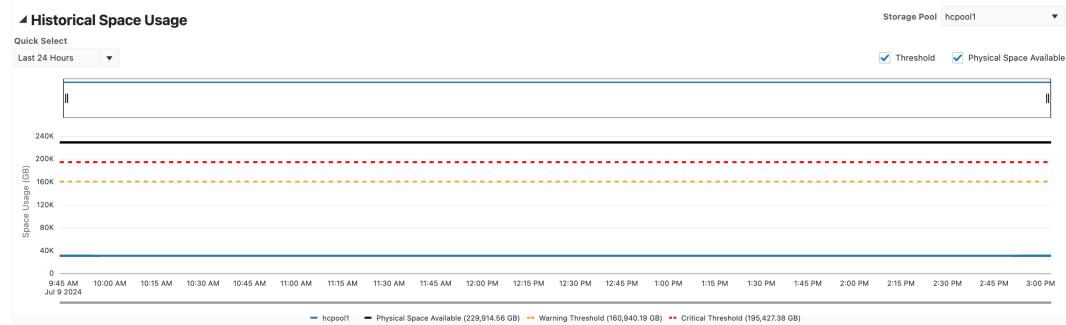
The **Capacity** tab has three sections:

- **Capacity:** This section displays the following summary statistics:
 - **Total Raw:** Total raw space available, in GB
 - **Total Used:** Total space used by clients, in GB
 - **Total Provisioned:** Total space provisioned to vaults, in GB
- **Storage Allocation:** This section shows the total available storage, the storage that has been provisioned by the storage pools to the vaults, and the storage used by Exascale clients. The section has three charts:
 - **Storage Server:** Displays information related to the Exadata Storage Server Grid that provisions storage to the Exascale, including total pool disk space provisioned and the pool disk space used by Exascale clients
 - **Storage Pool:** Displays storage pool information. Media type is color-coded in the charts. A storage pool consists of all pool disks of the same media type, either High Capacity (*HC*) or Extreme Flash (*EF*). The charts provide an option to view the space that is provisioned to the vaults and the space that is used by Exascale clients. Space can be viewed as a percentage or an absolute value, in GB. Detailed chart information is displayed upon mouse-over.

- **Vaults:** Displays how the vaults have been provisioned and used. Storage for the vaults can be provisioned from more than one storage pool. Media type is color-coded in the charts using the same colors as the storage pool charts.



- **Historical Space Usage:** This section displays storage pool space usage over a period of time. The following reference lines can be displayed on the chart:
 - **Threshold:** Displays the storage pool usage warning and critical thresholds, if they are set in Oracle Enterprise Manager.
 - **Physical Space Available:** Displays storage pool physical space



Serviced Clients

The **Serviced Clients** tab has three sections:

- **Serviced Clients:** This section contains an overview of the Exascale clients that includes the following information:
 - **Databases:** The number of databases provisioned on this Exascale and their status. This data includes all client databases regardless of whether they are monitored by Oracle Enterprise Manager.
 - **Volumes:** The number of volumes and the count by availability
- **Serviced Databases:** This section contains a table listing all databases provisioned on this Exascale along with the following information:
 - **Status:** Database availability status in Oracle Enterprise Manager

- **Allocated Space Used:** Space used by the database client as a percentage of the allocated space
- **Space Used (GB):** Space used by database client in GB
- **Vaults:** The name of the vaults where the database datafiles are stored
- **Number of Vaults:** The number of vaults used by the database. Oracle recommends that a database uses a single vault. A warning icon is displayed next to databases that use more than one vault.

The table can be filtered by database target name and/or vault name.

- **Serviced Volumes:** This section contains a table listing volumes that have been created from the vaults along with following information:
 - **Media Type:** The media type of the volume - *HC*, *EF*, *NULLDEVICE*, or *RAMDEVICE*
 - **Redundancy:** The redundancy of the volume - *high* (3-way mirroring), *normal* (2-way mirroring), or *none*
 - **State:** The state of the volume - *UNAVAILABLE*, *AVAILABLE*, *DELETING*, *DELETED*, *ATTACHING*, or *DETACHING*
 - **Size:** The size of the volume in GB
 - **Attachment Type:** The attachment type of the volume - *EDV* or *iSCSI*
 - **Device Name:** The name of the storage device associated with the attachment, shown only for EDV attachment type
 - **Device Path:** The absolute path of the storage device on the EDV node, shown only for EDV attachment type
 - **GI Cluster/Host:** The name of the cluster that contains the EDV attachment or name of the host associated with the attachment

The table can be filtered by volume name and/or vault name.

The screenshot shows the Oracle Exadata Exascale interface with the following sections:

- Serviced Clients:**
 - Databases:** Total: 1 (↑1)
 - Volumes:** Total: 2 (↑2)
- Serviced Databases:**

Name	Status	Allocated Space Used (%)	Space Used (GB)	Vaults	Number of Vaults
esdb23_1	↑	49.70	138.73	l050102-dbvl1t2	1

Allocated Space Used (%) Used: Free

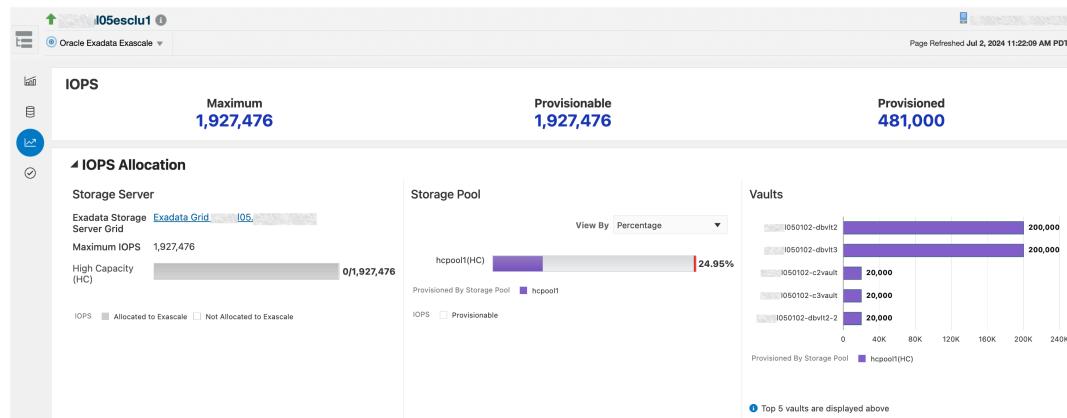
The space shown here is logical space and does not include recovery data. The database needs to be monitored. ⚠ Indicates that the database uses multiple vaults and will be unable to perform CDB snapshots.
- Serviced Volumes:**

Name	Vault	Media Type	Redundancy	State	Size (GB)	Attachment Type	Device Name	Device Path	GI Cluster/Host
edv2	l05-edvvault2	HC	high	AVAILABLE	50	EDV	esvol1	/dev/exc/esvol1	
edv3	l05-edvvault3	HC	high	AVAILABLE	50	EDV	esvol1	/dev/exc/esvol1	

Performance

The **Performance** tab has two monitoring sections:

- **IOPS:** This section displays an overview of IOPS provisioning, including the following information:
 - **Maximum IOPS:** The maximum IOPS available for the Exascale
 - **Provisionable IOPS:** The IOPS that can be provisioned while creating new vaults
 - **Provisioned IOPS:** The IOPS that are already provisioned
- **IOPS Allocation:** This section displays performance information in terms of the maximum IOPS, the IOPS that are provisioned by storage pools to vaults, and the IOPS used by clients:
 - **Storage Server:** The maximum IOPS per media type available with the storage servers that can be used by Exascale
 - **Storage Pool:** Shows the IOPS per media type provisioned to the vaults by percentage and absolute values. The media type is color-coded.
 - **Vaults:** Shows how IOPS are provisioned for the vaults. IOPS for vaults can be provisioned from more than one storage pool. The media type is color-coded using the same colors as the storage pool charts.



Incidents and Compliance

The **Incidents and Compliance** tab has three sections:

- **Summary:** This section displays a status summary for the Exascale target, including the following information:
 - **Status:** The status of the Exascale cluster
 - **Version:** The release version of the Exadata system software
 - **Storage Server Status:** The status of the storage servers that provision pool disks to the Exascale storage pool
- **Storage Pool Health:** This section displays a tree table with storage pools as the root, which can be expanded to show all the storage servers that provision pool disks for the storage pool. It can be expanded further to display the individual pool disks. The tree table displays the following information:
 - **Cell Status:** The status of the Exadata storage server target
 - **Storage Server:** The storage server that provisions the pool disk
 - **Pool Disk:** The name of the pool disk provisioned to the storage pool

- **Pool Disk Status:** The status of the pool disk
- **Used (%):** The pool disk storage used percentage
- **Size (GB):** The size of the pool disk
- **Healthy:** Health of the pool disk

Storage Pool	Cell Status	Storage Server	Pool Disk	Pool Disk Status	Used (%)	Size (GB)	Healthy
hcpool1	green	i05celadm01	hcpool1_CD_00_i05celadm01	ONLINE	11.94	6,386.52	true
	green		hcpool1_CD_01_i05celadm01	ONLINE	11.93	6,386.52	true
	green		hcpool1_CD_02_i05celadm01	ONLINE	11.91	6,386.52	true
	green		hcpool1_CD_03_i05celadm01	ONLINE	11.95	6,386.52	true
	green		hcpool1_CD_04_i05celadm01	ONLINE	12.01	6,386.52	true
	green		hcpool1_CD_05_i05celadm01	ONLINE	11.95	6,386.52	true
	green		hcpool1_CD_06_i05celadm01	ONLINE	11.97	6,386.52	true
	green		hcpool1_CD_07_i05celadm01	ONLINE	12.00	6,386.52	true
	green		hcpool1_CD_08_i05celadm01	ONLINE	12.02	6,386.52	true
	green		hcpool1_CD_09_i05celadm01	ONLINE	12.01	6,386.52	true
	green		hcpool1_CD_10_i05celadm01	ONLINE	11.97	6,386.52	true

- **Incidents:** This section contains a table showing all the incidents related to the Exascale target.

Severity	Summary	Target Name	Target Type	Status	Escalation Level	Age	Time Since Last Update
No data to display.							

Flash Cache Resource Monitoring

The Oracle Exadata plug-in release 13.1.0.1.0 and later provides flash cache resource monitoring for Oracle Exadata Storage Servers. The following information is displayed in different sections and pages:

I/O Distribution by Databases: This section on the home page gives the latest performance statistics for the Top 5 databases based on their I/O Utilization.

I/O Distribution for All Databases: This page gives the latest performance statistics for all the databases in the order of I/O Utilization. This page provides both Graph View and Table View. To launch this page, from home page, go to the section **I/O Distribution by Databases** > Click **I/O Distribution for All Databases**.

Workload distribution by Databases: This tab on IORM page provides historical performance statistics by databases for both Hard Drive and Flash. You can specify Top 5 databases or specific databases to view their graphs.

Flash Cache Space Usage: This tab on IORM page provides historical statistics specific to Flash Cache Space Usage by databases.

To view the graphs, go to [Storage Resources](#).

Exadata Database Machine Fault Monitoring

Oracle Enterprise Manager Cloud Control provides hardware fault monitoring for Oracle Exadata Database Machine. For a detailed list of the available hardware alert messages, see [Hardware Alert Messages for Oracle Exadata Database Machine and Oracle Exadata Rack](#) in *Exadata Database Machine Error Message Reference Guide*.

Monitoring Exadata Database Machine Components

Enterprise Manager collects details for the following components:

- [Exadata Storage Servers](#)
- [Cisco/RoCE Switch \(Systems Infrastructure Switch\)](#)
- [ILOM Targets \(Systems Infrastructure Server\)](#)
- [Power Distribution Units \(PDUs\)](#)
- [InfiniBand Switch \(Systems Infrastructure Switch\)](#)

Exadata Storage Servers

An Enterprise Manager Agent could collect Exadata Storage Server metrics through `cellcli` calls via ssh, ExaCLI or RESTful API. SNMP traps are sent to the Enterprise Manager Agent for subscribed alert conditions.

Monitoring with `cellcli` requires SSH equivalence to be set up between the agent user and the Exadata Storage Server monitor user on each Exadata Storage Server.

On the home page, rich storage data is collected, including:

- Aggregate storage metrics.
- Exadata Storage Server alerts via SNMP (PUSH).
- Capacities.
- Database-level metrics.

Cisco/RoCE Switch (Systems Infrastructure Switch)

An Enterprise Manager Agent runs a remote SNMP get call to collect metric data for the Cisco/RoCE switch, including details on:

- Status / Availability
- Port status
- Vital signs: CPU, Memory, Power, Temperature
- Network interface various data

- Incoming traffic errors, traffic kb/s and %
- Outgoing traffic errors, traffic kb/s and %
- Administration and Operational bandwidth Mb/s

ILOM Targets (Systems Infrastructure Server)

An Enterprise Manager Agent runs remote SSH calls to each SI server target.

The following details are collected:

- Response – availability
- Sensor alerts
 - Temperature
 - Voltage
 - Fan speeds
- Configuration Data: Firmware version, serial number, and so forth.

Monitoring Systems Infrastructure ILOM Server Target Using Non-root User

Follow these steps to create the non-root user for monitoring the ILOM server target:

1. Login to the ILOM server from the browser console.
2. Go to **ILOM Administration**. In the left side panel, click **User Management**.
3. In the **Users** section, click **Add User**. The Add User dialog box is displayed.
4. In the Add User dialog box, create the new ILOM non-root user. Provide **user name**, **password**, and enable the role permissions **Console (c)** and **Read Only (o)**.
5. Save this new user and verify the user details.

Note

When the non-root user is used, create the SNMPv3 user / SNMP v2 community string and SNMP subscription on the server manually to receive the SNMP traps.

Power Distribution Units (PDUs)

An Enterprise Manager Agent runs remote SNMP get calls and receives SNMP traps (PUSH) from each PDU. Collected details include:

- Response and ping status.
- Phase values.

InfiniBand Switch (Systems Infrastructure Switch)

An Enterprise Manager Agent runs remote SSH calls to the InfiniBand switch to collect metrics. The InfiniBand Switch sends SNMP traps (PUSH) for all alerts.

Monitoring requires ssh equivalence for the root or ilom-admin user for metric collections such as:

- Response
- Various sensor status
 - Fan
 - Voltage
 - Temperature
- Port performance data
- Port administration

Automatic Indexing for Database Nodes

The automatic indexing feature automates the index management tasks in database nodes with Oracle Database 19.3 or later. Automatic indexing automatically creates, rebuilds, and drops indexes in a database based on the changes in application workload, thus improving database performance.

This feature includes the following options:

- **Settings:** On the **Automatic Indexing Configuration Settings** page, you can enable and disable automatic indexing, specify the retention period for unused indexes and automatic indexing logs, and specify the schemas to be included or excluded from using automatic indexing.
- **Activity Report:** On the **Automatic Indexing Activity Report** page, you can enter the following details and click **Generate Report** to view the details of the auto indexes generated in the database:
 - **Report Format:** Select the format in which you want the report to be generated.
 - **Sections:** Select the sections that you want displayed in the report. The **Summary**, **Index Details**, **Verification Details**, and **Errors** options are selected by default in the **Sections** field, and you can opt to remove any of these sections.
 - **Time Period:** Select the monitoring time period for which you want the report to be generated.

If you select the default options, namely the **HTML** report format and all the options in the **Sections** field for a specific time period, then the following sections are displayed in the **Report Summary**:

- **Overview of Executions:** This section displays the overall performance improvement factor as a result of the auto indexes, the number of times the auto index operations were completed, the number of times the auto index operations were interrupted, and fatal errors, if any. Note that the **Overview of Executions** section is displayed irrespective of which other section is selected in the **Sections** field.
- For the **Summary** section option, the following bar graphs are displayed:
 - * **Summary of Auto Indexes Actions**
 - * **Summary of Auto Indexes SQL Actions**
 - * **Summary of Manual Indexes**
- For the **Errors** section option, the **Error Summary** pie chart is displayed.
- For the **Index Details** section option, the following sections are displayed:

- * **Index: Created**
- * **Index: Dropped**

You can click a row in the **Index: Created** and **Index: Dropped** sections to view more details such as the ID of the index that was created or dropped, the **Key**, and **Type**. You can also use the **Download All Index Created Details** and **Download All Index Dropped Details** options given in these sections to download the index details in a .csv format.

- For the **Verification Details** section option, the **Verification Details** section is displayed, which includes the details of the SQLs for which auto indexes were generated and used. You can click a row in this section to view more details such as **SQL ID**, **SQL Text**, and **Improvement Factor**.

If you select the **TEXT** report format in the **Report Format** field, then the same information is displayed in a plain text format and can be downloaded by clicking the **Text Download** option.

For information on managing auto indexes, see [Managing Auto Indexes](#) in *Oracle Database Administrator's Guide*.

Evaluating Compliance of Exadata Targets

Oracle Enterprise Manager 13c Release 5 Update 3 (13.5.0.3) integrates Oracle Autonomous Health Framework (AHF) EXAchk for Exadata Engineered Systems X8 and X8M. Oracle EXAchk is a lightweight and non-intrusive health check framework for stack of software and hardware components in Exadata.

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.

In order to use AHF EXAchk Compliance Standards, there is no need for an Enterprise Manager plug-in. These EXAchk standards are available out-of-box. When AHF is upgraded, the associated EXAchk compliance standards are automatically updated to the corresponding version.

To monitor compliance of the Exadata targets, use the **Engineered Systems** tab in the **Compliance Dashboard**. See Evaluate Compliance in *Cloud Control Database Lifecycle Management Administrator's Guide*.

For more information on AHF EXAchk Compliance Standards, see AHF EXACHK Compliance Standards in *Oracle Enterprise Manager Cloud Control Oracle Compliance Standards Reference*.

Reporting

To monitor the performance of the database machine or the storage server by viewing the metrics and diagnostic data, access the support workbench or create a dashboard for access the database machine reports and capacity planning reports.

Topics:

- [Accessing Oracle Support Workbench for Exadata Storage Server](#)
- [Accessing Oracle Analytics Server Reports](#)
- [Oracle Exadata Database Machine Dashboard Creation](#)

Accessing Oracle Support Workbench for Exadata Storage Server

You can access the Oracle Support Workbench for the current Exadata Storage Server to access diagnostic data for problems and incidents related to the Exadata Storage Server.

To access the Support Workbench for a single Exadata Storage Server, follow these steps:

1. From the **Target** menu of the Exadata Storage Server for which you want to access the Oracle Support Workbench, select **Diagnostics**, then select **Support Workbench**.
2. If you are logging in to the Support Workbench for the first time, you must first initialize the Support Workbench properties. For first-time log in, the Initialize Support Workbench Properties window will be displayed.

Enter the Automatic Diagnostic Repository (ADR) Base Path:

The ADR Home Path is a target-specific, writable path on the agent where the generated package is to be stored. During initializing, it is not editable as it is generated automatically once the ADR base path is specified.

Click **Create**.

3. Once the Support Workbench properties have been initialized, the Support Workbench page displays showing a list of any problems for the Database Machine:

Select	Details	ID	Problem Key	Description	Number Of Incidents	Last Incident	Last Comment	Active	Packaged	SR#
<input type="checkbox"/>	> Show	13	Other 13	Cell Issues	1	January 21, 2016 1:32:54 PM UTC		Yes	No	
<input type="checkbox"/>	> Show	12	Other 12	Database Crash	1	January 21, 2016 10:48:10 AM UTC		Yes	No	
<input type="checkbox"/>	> Show	11	Other 11	Database Crash	1	January 21, 2016 10:42:13 AM UTC	Set SR :	Yes	Yes	3-123456
<input type="checkbox"/>	> Show	10	Other 10	Database Crash	1	January 21, 2016 5:55:38 AM UTC		Yes	Yes	3-123789

Accessing Oracle Analytics Server Reports

Access the Oracle Analytics Server to run the Exadata reports *Exadata Capacity Planning Details* and *Exadata vCPU Consumption 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 Cloud Control Administrator's Guide*.
2. Expand the **Exadata Reports** tree.

The links to the following reports are displayed:

- **Capacity Planning Details:** The capacity planning report for Oracle Exadata Database Machine.
- **Exadata Cloud Service Capacity Planning Report:** The capacity planning report for Oracle Exadata Cloud. See Exadata Cloud Service Capacity Planning Report in *Oracle Enterprise Manager Cloud Control for Oracle Exadata Cloud*.

Note that the report for Oracle Exadata Cloud has a subset of applicable sections of the Capacity Planning Details report used for Oracle Exadata Database Machine, specifically: **Configuration**, **Global Summary**, **Storage Capacity Planning Details**, **CPU and Memory Capacity Planning Details**, and **IO Capacity Planning Details**.

- **Exadata vCPU Consumption Report:** Total daily vCPU consumption and daily vCPU consumption by compute node and virtual host for Exadata.
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 Capacity Planning Details

Capacity Planning Details report enables you to analyze and visualize the utilization and health of Exadata resources like storage, CPU, memory, IO and network. 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

Exadata Machine Configuration

This section reports the name of the Exadata Database Machine 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 Database Machine:

	Last 24 hours ¹				Last 7 Days ²				Last 31 Days ³				Last 365 Days ³			
	Maximum		Average		Maximum		Average		Maximum		Average		Maximum		Average	
	CPU	16.26%	8.92%	26.85%	16.89%	26.85%	11.79%	26.85%	11.79%	79.73%	79.15%	77.21%	79.73%	77.21%	77.21%	77.21%
Network Interface - Gb/s																
bondeth0	Receive	Transmit	Receive	Transmit	Receive	Transmit	Receive	Transmit	Receive	Transmit	Receive	Transmit	Receive	Transmit	Receive	Transmit
bondib0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
docker0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
eth0	0.00	0.00	0.00	0.00	0.02	0.00	0.00	0.00	0.02	0.00	0.00	0.00	0.02	0.00	0.00	0.00
eth2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
eth3	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
eth4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
eth5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ib0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
ib1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Storage Capacity Planning Details

As an administrator, you can use the *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

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 y-axis against time which is on x-axis:

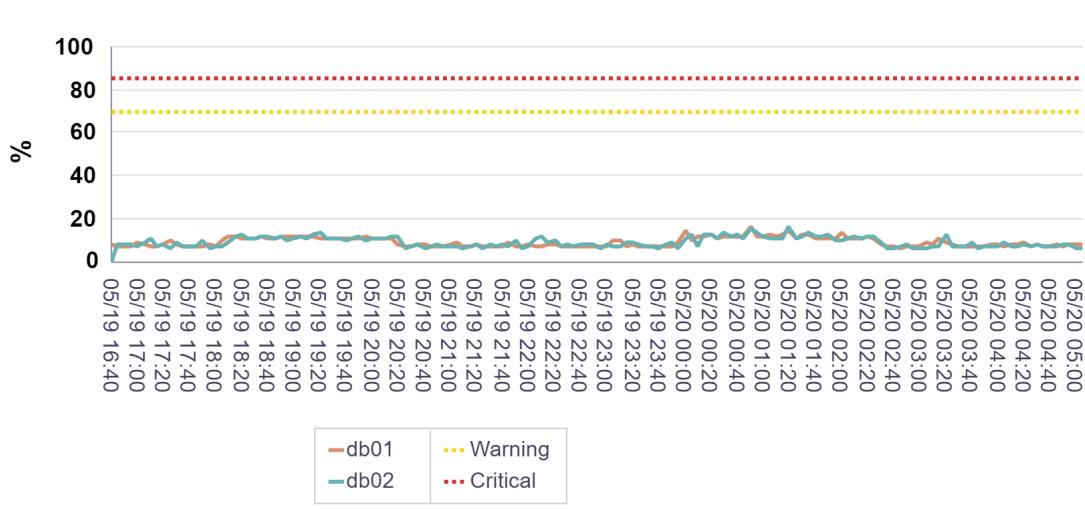
- Historical Space Usage By Disk Group (GB) - Usable
- Historical Space Usage By Disk Group (%) - Usable
- Historical Total Space Usage (GB) - Raw
- Historical Total Space Usage (GB) - Usable
- Historical Space Usage by the Top 5 Serviced Databases - As seen daily

CPU and Memory Capacity Planning Details

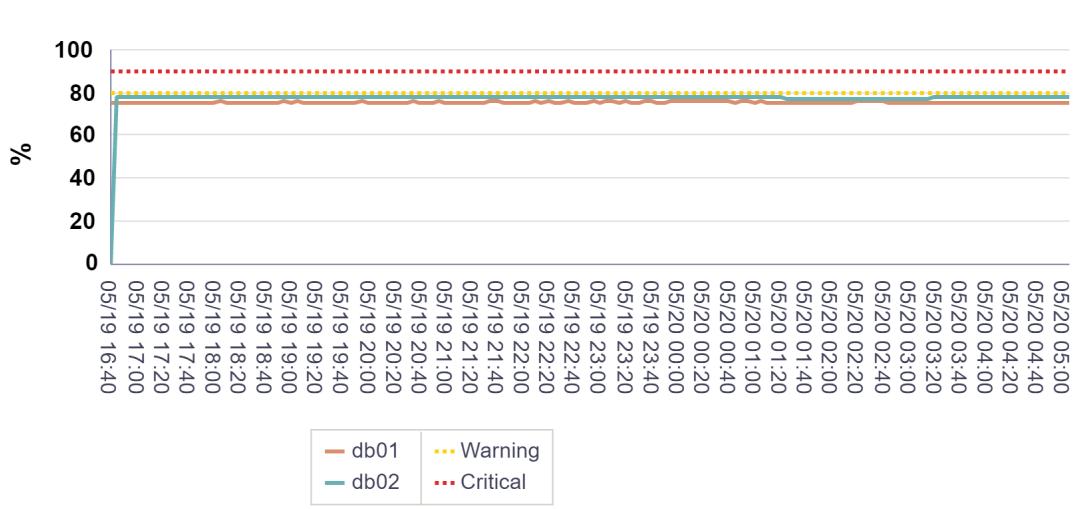
Your CPU and Memory utilizations should not be at critical limits during normal load hours. You should determine CPU and Memory utilizations based on your application needs, including CPU and Memory cycles for peak usage. If your CPU and Memory utilizations are optimized at 100% during normal load hours, 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 24 hours time range:



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



IO Capacity Planning Details

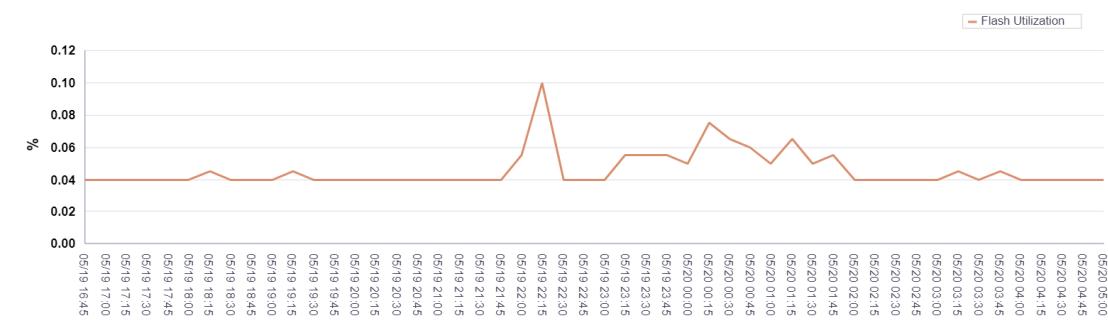
The **IO Capacity Planning Details** section displays the current and historical IO, throughput and disk utilization aggregated across all the storage servers for the selected Exadata Database Machine. It summarizes the amount of data being stored in the disks, the read/write throughput rate, and load handling in response to normal and peak demands. Use this section 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*:

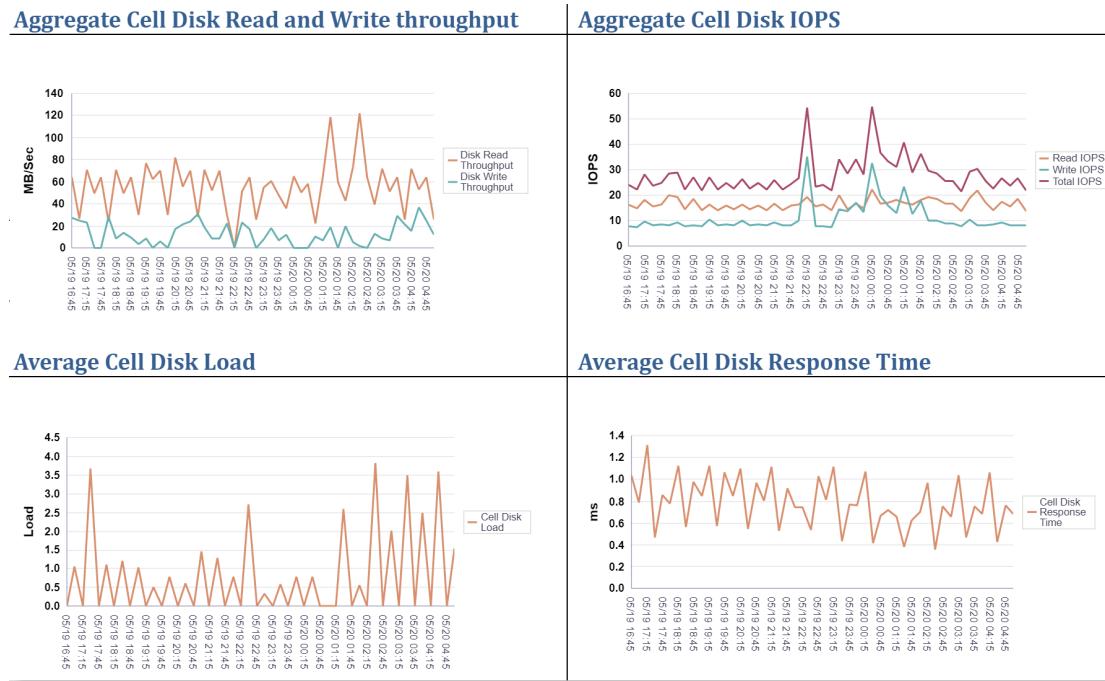
- Average Cell Disk Utilization
- Aggregate Cell Disk Read and Write throughput
 - Disk Read Throughput
 - Disk Write Throughput
- Aggregate Cell Disk IOPS
 - Read IOPS
 - Write IOPS
 - Total IOPS
- Average Cell Disk Load
- Average Cell Disk Response Time
- Average Flash Utilization
- Aggregate Flash Read and Write throughput
 - Flash Read Throughput
 - Flash Write 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 24-hour time range:

Average Flash Utilization



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



Network Capacity Planning

The *Network Capacity Planning* section displays the **Bonded** (logical) interfaces and **Slave** (Physical) interfaces throughput. Network bonding corresponds to a single IP address and is the single logical bonded interface formed by aggregating multiple physical network interfaces. Use this section to determine the amount of network traffic handled at various loads.

The section displays the following *Transmit* and *Receive* throughputs in Gb/sec for bonded and slave interfaces measured and plotted in line charts for the time ranges *Last 24 Hours*, *Last 7 Days*, *Last 31 Days*, and *Last 365 Days*:

- Transmit - Bonded Interfaces
- Receive - Bonded Interfaces
- Transmit - Slave Interfaces
- Receive - Slave Interfaces

This section does not show the RDS protocol traffic metrics.

Exadata vCPU Consumption Report

Use the *Exadata vCPU Consumption Report* to track the vCPU usage of your virtualized Exadata Database Machine. Compare the vCPU consumption for different payloads and plan the vCPU allocations for future requirements.

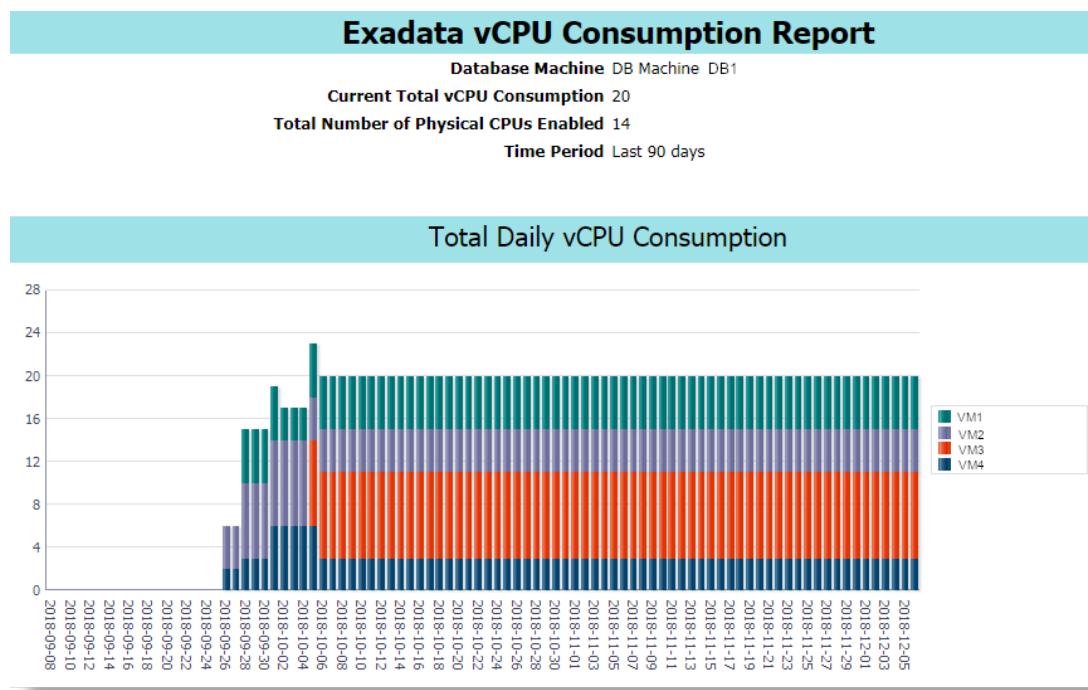
The two parameters used in the report are *Exadata target name* and *time period*. The report is generated based on an Exadata target, and defined for the time ranges like *7 Days*, *30 Days*, *90 Days*, *180 Days*, and *1 Year*, *2 years*, *3 years*, and *4 years*. The default time range is *30 Days*. If there is no data on a specific day, then the default value is 0.

Report Summary: The summary includes the selected Exadata target name, selected time period, current total vCPU consumption, and total number of physical CPUs enabled. The last two values add all the current values together from each virtual machine.

The following metrics are generated in the report for the specific Database Machine:

Total Daily vCPU Consumption

Total vCPU consumption is reported based on the time period that you select. For example, if you select 30 days, the data for 30 days including the current date is reported. The details like Exadata target, compute nodes, and hosts are obtained from the Exadata target name and target association relationship. The CPU count is obtained from the historical host configuration metrics. Within a day, only the high value of the number of vCPUs is used for the daily CPU count. On the whole, the CPU count is aggregated by the virtual host for each day. The chart that displays the total vCPU consumption for the time period of 90 days is shown below:



Date	Total vCPU Count
Dec 14, 2018	18
Dec 13, 2018	12
Dec 12, 2018	12
Dec 11, 2018	0
Dec 10, 2018	0
Dec 9, 2018	0
Dec 8, 2018	0

Daily vCPU Consumption by Compute Node and Virtual Host

This table is similar to the table for total vCPU consumption, but without the aggregation of the count. The section presents the data grouped by compute nodes and host.

Following is the daily vCPU consumption of the listed compute nodes and their hosts:

Compute Node	Host	Date	vCPU Count
ComputeNode1	Host1	Dec 14, 2018	8
		Dec 13, 2018	5
		Dec 12, 2018	5
		Dec 11, 2018	0
		Dec 10, 2018	0
		Dec 9, 2018	0
		Dec 8, 2018	0
ComputeNode2	Host2	Dec 14, 2018	10
		Dec 13, 2018	7
		Dec 12, 2018	7
		Dec 11, 2018	0
		Dec 10, 2018	0
		Dec 9, 2018	0
		Dec 8, 2018	0

Oracle Exadata Database Machine Dashboard Creation

You can create a Database Machine dashboard to monitor the performance and usage metrics of the Database Machine system, its sub-components, as well as all the database system components residing on the Database Machine.

- [Exadata Database Machine Dashboard Creation](#)
- [Find All Database Machine Reports](#)
- [How to Make the Report "Public"](#)

Exadata Database Machine Dashboard Creation

1. Log in to Enterprise Manager.
2. From the Enterprise Manager home page, click the **Enterprise** menu. Select the **Job** submenu, then **Library**.
3. Select **Database Machine Services Dashboard** from the drop-down menu next to the Create Library Job option.
4. Click **Go**.

5. Enter a name of the job (for example, CREATE_DBM_DASHBOARD_1).
6. Click **Add** and select the **DBMachine** target. After adding, make sure the target is selected by clicking on the check box next to it. Note that you can select more than one DBMachine and that a dashboard report will be created for each one.
7. Click the **Parameters** tab.

Three options are provided through the drop-down:

- Select **Create** if it is a new report.
- Select **Update** for updating an existing report with new components.
- Select **Cleanup** to remove services created by the Create job executed earlier.

8. Finally, click **Save to Library** to save the job.

A message indicating the job was created successfully should display along with a link to the running job.

You can monitor the job by clicking on the link corresponding to the job.

9. To run the job, from the Job Library page, select the newly created job and click **Submit**. This will open the job to allow you to select the Targets to run against. Select the correct targets and click **Submit**.

Find All Database Machine Reports

To find all Database Machine reports:

1. Log in to Enterprise Manager.
2. From the Enterprise Manager home page, click the **Enterprise** menu. Select **Reports**, then **Information Publisher Reports**.
3. Search for the report name. Dashboard report names, one for each Database Machine, are displayed in the following format:
[DBMachine Name]_DASHBOARD_REPORT
4. Click on the report to view the dashboard report.

How to Make the Report "Public"

The generated report is accessible only by the Enterprise Manager user who creates it. To make the report public:

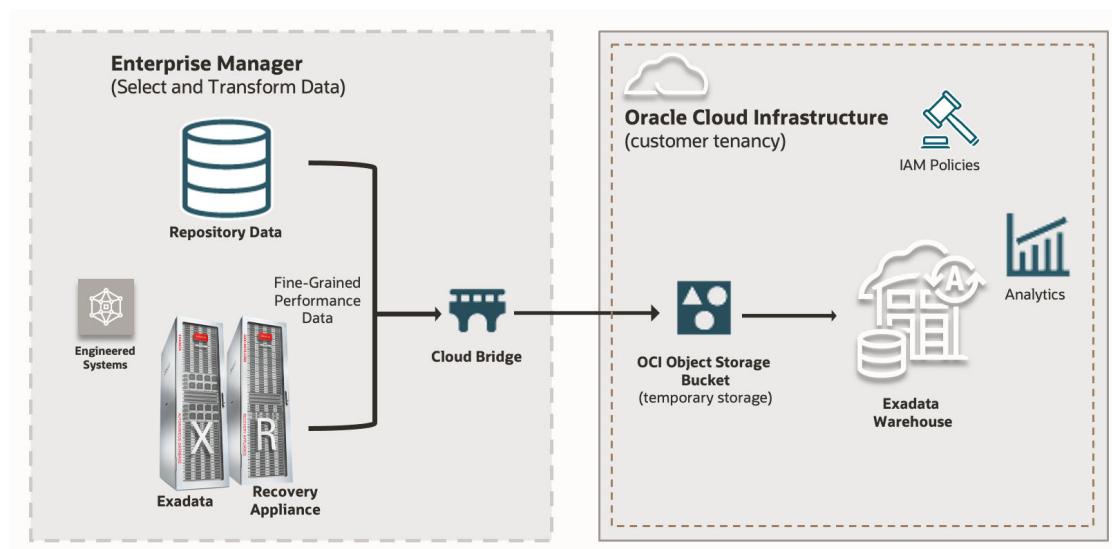
1. Select the dashboard report from the list of Information Publisher reports shown after following the steps mentioned above.
2. Click **Edit**.
3. Select the **Run report using target privileges of the report owner** option under Privileges section in General tab.
4. Click the **Access** tab.
5. Select the **Allow viewing without logging in to Enterprise Manager** option.
6. Click **OK**.

Exadata Warehouse

Set up Exadata Warehouse through Enterprise Manager Cloud Control such that it serves as a repository for fine-grained performance metric data collected from on-premises and cloud-based Oracle Engineered Systems like Exadata Database Machine and Zero Data Loss Recovery Appliance.

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

The following image illustrates how data flows 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](#)
- [Configuring Purge Settings for Old Metric Data](#)

Setting Up Exadata Warehouse and Analyzing Performance Data

For steps to set up Exadata Warehouse through Enterprise Manager Cloud Control, 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 Enterprise Manager Cloud Control (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 customer-owned Autonomous Databases in customer's tenancy
- Store fine-grained metric data with per-minute granularity. Offers up to 1,440 data point 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 Compute Nodes 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 Compute Nodes	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

Target	Oracle Storage Server	Disk Type	Total IOPS
DB Machine sample4.example.com	sample4adm11.example.com	FD	14605.37
DB Machine sample4.example.com	sample4adm10.example.com	FD	15722.89
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 Enterprise Manager Cloud Control, 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:

The screenshot shows the 'Parameters' tab of the 'Create 'Exadata Warehouse Auto Purge Settings' Job' dialog. It includes fields for retention periods and a table of target properties.

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

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

Exadata Metrics

This chapter describes how to access the various metrics for the Oracle Exadata Database Machine and provides examples of key metrics including:

- [Accessing Exadata Metrics](#)
- [Example Metrics](#)
 - [Aggregated Exadata FlashDisk and HardDisk Metric](#)
 - [Exadata Storage Server Metric](#)
 - [Exadata Key Performance Indicators Metrics](#)

Accessing Exadata Metrics

For a complete list of metrics available for the Oracle Exadata Database Machine and related targets, see the following reference manuals for the plug-ins that together provide Exadata support:

- *Oracle Exadata in Enterprise Manager Oracle Engineered Systems Metric Reference Manual*
- *Systems Infrastructure Plug-in Metric Reference Manual*
- *Virtual Infrastructure Plug-in Metric Reference Manual*

Accessing Exadata Metrics

To access the available Exadata Metrics:

1. From the Targets menu, select **Exadata**.
2. On the Oracle Exadata Database Machines page, select an Exadata Database Machine from the list.
3. From the Target Navigation tree, select a target of interest.
4. From the target's home page, select the target type name, then **Monitoring**, and then **All Metrics**.

Example Metrics

The following metric sections provide examples of the available metrics and how they are displayed.

Aggregated Exadata FlashDisk and HardDisk Metric

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

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 Exadata Storage Server disk, or **Total CellDisk I/O Load**, which gives an indication of the total input/output load to the celldisk.

Average CellDisk Read Throughput:

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	-

Total CellDisk I/O Load:

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 Storage Server Metric

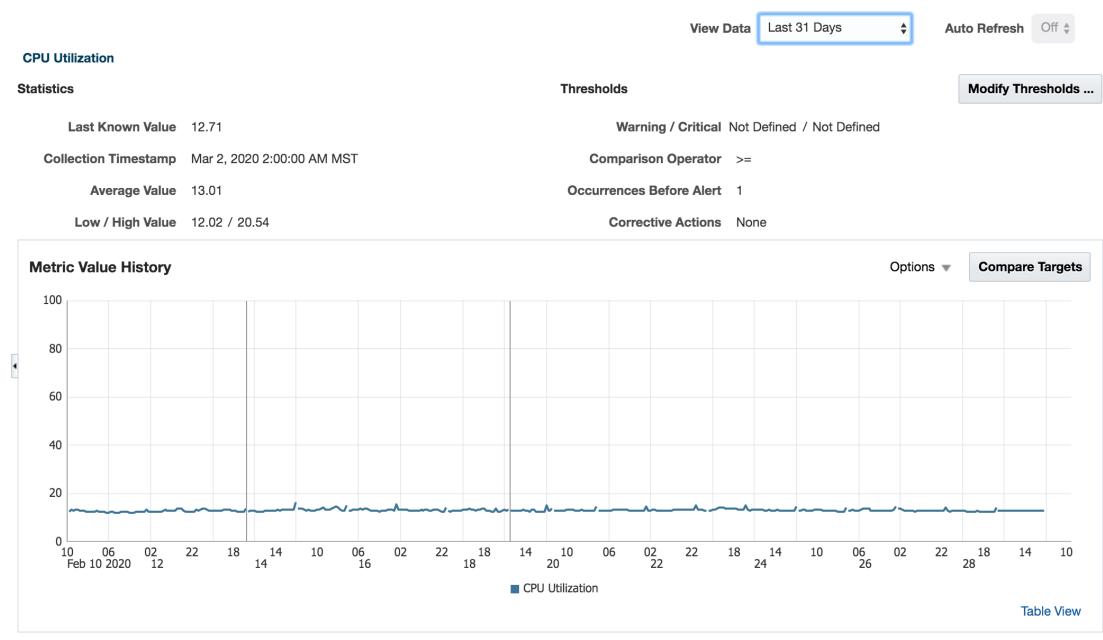
This metric category contains the performance metrics collected at the Exadata Storage Server level for each Exadata Storage Server, such as CPU utilization and memory utilization. Selecting this metric from the All Metrics page generates a high-level summary, as shown in the image:

Exadata Cell Metric**Collection Schedule** Every 15 Minutes **Upload Interval** Every Collection**Last Upload** Mar 2, 2020 2:00:00 AM MST

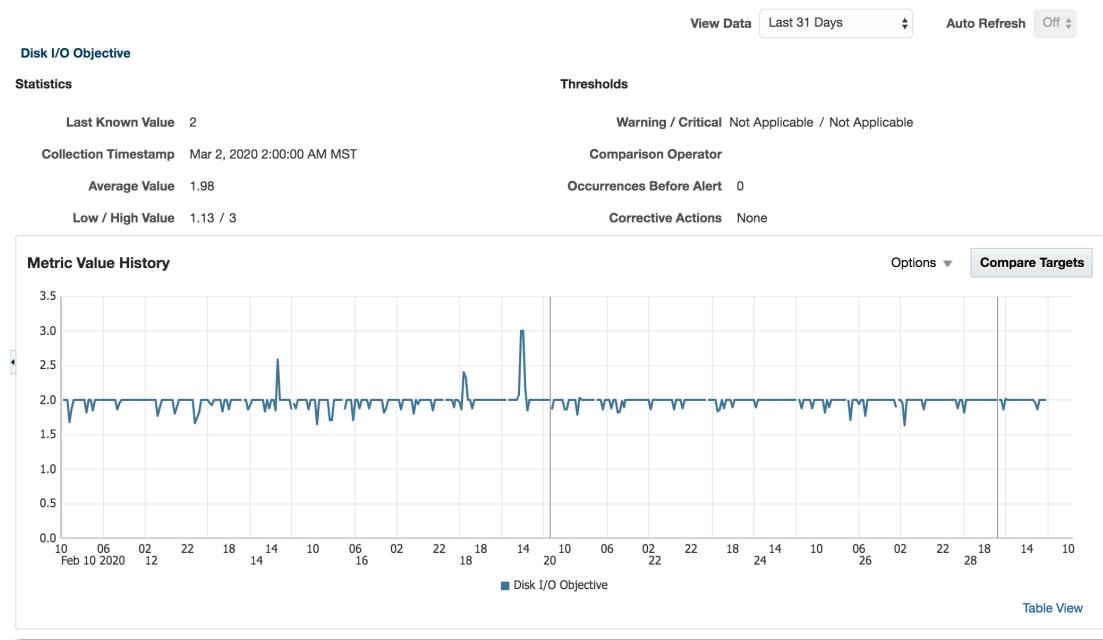
Metric	Thresholds	Real Time Value
Cell Name	Not Applicable	sicc15cel01
CPU Utilization	Not Set	12.7
Disk I/O Objective	Not Applicable	2
Exadata Run Queue Length	Not Set	3.37
Exadata Temperature Lower Threshold	Not Set	21
Exadata Temperature Reading	Not Applicable	21
Exadata Temperature Upper Threshold	Not Set	21
IORM Boost	Not Applicable	0
LED Status	Not Applicable	off
Memory Utilization	Not Set	96.07
Network Received	Not Applicable	0.08
Network Sent	Not Applicable	4.83
Offload Efficiency	Not Set	

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).

CPU Utilization:



Disk I/O Objective:



Exadata Key Performance Indicators Metrics

With Enterprise Manager 13c Release 4 Update 4 (13.4.0.4) or higher version, you can access the Key Performance Indicators Metrics for Exadata Storage Servers and Exadata Storage Server Grid to view the metrics of the key 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

 Response

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:

Average Hard Disk IO Load
Description : Average IO load across all hard disks on the Exadata Storage Server.
Average Value 1.33

Occurrences Before Alert 1

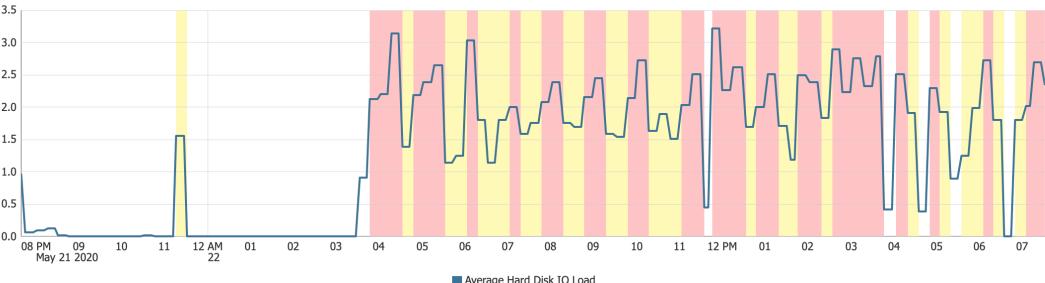
Low / High Value 0 / 3.22

Corrective Actions None

Thresholds last modified by USER_ADMIN at May 20, 2020 4:04:57 PM MDT.

Metric Value History

Options  



Average Hard Disk IO Load

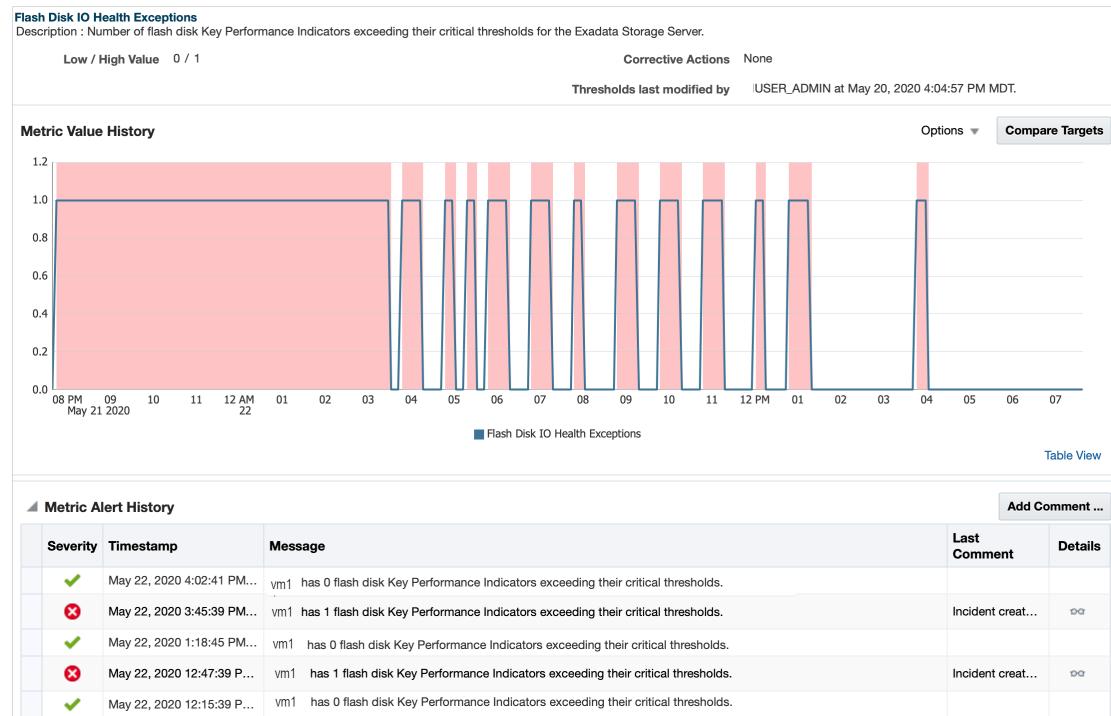


Metric Alert History

Add Comment ...

Severity	Timestamp	Message	Last Comment
	May 22, 2020 7:04:39 PM...	Average hard disk IO load for adm01 is 2.023, crossed warning (1) or critical (2) threshold.	Incident created by rule (Nam)
	May 22, 2020 6:49:39 PM...	Average hard disk IO load for adm01 is 1.809, crossed warning (1) or critical (2) threshold.	
	May 22, 2020 6:34:39 PM...	Average hard disk IO load for adm01 is 0.001, fallen below warning (1) or critical (2) threshold.	
	May 22, 2020 6:19:39 PM...	Average hard disk IO load for adm01 is 1.808, crossed warning (1) or critical (2) threshold.	
	May 22, 2020 6:04:39 PM...	Average hard disk IO load for adm01 is 2.733, crossed warning (1) or critical (2) threshold.	Incident created by rule (Nam)

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



Troubleshooting

This chapter provides troubleshooting tips and techniques on installing, discovering, and configuring the Exadata plug-in. The following sections are provided:

- [Discovery Troubleshooting](#)
- [EMCLI Discovery Issues](#)
- [Troubleshooting the Exadata Database Machine Schematic File](#)
- [Exadata Database Machine Management Troubleshooting](#)
- [Oracle Database Storage Server System Target Status Pending](#)
- [Target Status Issues](#)
- [Metric Collection Issues](#)
- [Status: Pending Issues](#)
- [Monitoring Agent Not Deployed for IPv6 Environments](#)
- [Monitoring: Error Fetching Data for Grid and CellSys Targets IORM Page Charts](#)
- [Not able to receive SNMP traps on Exadata Storage Servers using IPv6](#)
- [Exascale Discovery and Monitoring Issues](#)

Discovery Troubleshooting

Very often, the error message itself will include the cause for the error. Look for error messages in the OMS and agent logs (case insensitive search for `dbmdiscovery`) or in the Discovery window itself.

It is recommended that you review the logs for troubleshooting issues. For the list of logs that you can review and their locations, see [Review Logs](#).

The following sections are provided:

- [Hardware Availability](#)
- [Discovery Failure Diagnosis](#)
- [Exadata Storage Server is not Discovered](#)
- [Compute Node or InfiniBand Switch is not Discovered](#)
- [Extra or Missing Components in the Newly Discovered Exadata Database Machine](#)
- [ILOM, PDU, or Cisco Switch is not Discovered](#)
- [Test Connection Fails in PDU](#)
- [Target Does not Appear in the Components Step of Guided Discovery](#)
- [Target Status is Down or Metric Collection Error After Discovery](#)
- [ILOM Credential Validation Fails During Discovery](#)
- [Discovery Process Hangs](#)

- [SNMP Configuration Missing](#)
 - [Set Up SNMP Subscription for Exadata Storage Servers](#)
 - [Set Up SNMP Subscription for Cisco Ethernet Switch Targets](#)
 - [Set Up SNMP for Power Distribution Unit \(PDU\) Targets](#)
- [Time Out Error During SNMP Subscription](#)

Hardware Availability

All the hardware components must be "known" and reachable; otherwise, communication failures will occur. Use the `ping` command for each hardware component of the Exadata rack to make sure all names are resolved.

The MAP targets in Enterprise Manager Cloud Control 13c may fail while collecting the correlation identifier. This failure can happen if the credentials are incorrect OR if a target (for example, the ILOM) is too slow in responding.

ILOM can be slow when the number of open sessions on ILOM has exceeded the limit. You can resolve this issue by temporarily closing the sessions on the ILOM.

The rack placement of targets can fail:

- If `examan` did not return valid rack position for the target.
- If there is an existing target in the same location.

Discovery Failure Diagnosis

Should discovery of your Oracle Exadata Database Machine fail, collect the following information for diagnosis:

- Any `examan-*.xml`, `examan-*.html`, `targets-*.xml`, and `examan*.log` files from the `AGENT_ROOT/agent_inst/sysman/emd/state/exadata` directory.
- Agent logs: `emagent_perl.trc` and `gcagent.log`
- OMS logs: `emoms.trc` and `emoms.log`
- Any snapshot (screen capture) of errors shown on the target summary page.
- If the **Prerequisite Check** page shows critical errors, then discovery can be retried by clicking on the **Retry** menu and selecting the **Retry, Static Only** option.

Exadata Storage Server is not Discovered

If the Exadata Storage Server itself is not discovered, possible causes could be:

- The `/etc/oracle/cell/network-config/cellip.ora` file on the compute node is missing or unreadable by the agent user or Exadata Storage Server not listed in that file.
- Management Server (MS) or `cellsrv` is down.
- Exadata Storage Server management IP is changed improperly. Bouncing both `cellsrv` and MS may help.
- To check that the Exadata Storage Server is discovered with a valid management IP, run the following command from the grid infrastructure or database home on the compute node used for discovery:

```
$ORACLE_HOME/bin/kfod op=cellconfig
```

- 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.

Compute Node or InfiniBand Switch is not Discovered

If there are problems with discovery of the compute node or the InfiniBand switch, possible causes could be:

- The InfiniBand switch host name or ilom-admin password is incorrect.
- The connection from the compute node to the InfiniBand switch through SSH is blocked by a firewall.
- The InfiniBand switch is down or takes too long to respond to SSH.

To resolve problems with the compute node or InfiniBand switch discovery, try:

- If the InfiniBand switch node is not discovered, the InfiniBand switch model or switch firmware may not be supported by EM Exadata. Run the ibnetdiscover command. Output should look like:

```
Switch 36 "S-002128469f47a0a0" # "Sun DCS 36 QDR switch switch1.example.com"  
enhanced port 0 lid 1 lmc 0
```

- If the compute node is not discovered, then run the ibnetdiscover command on the compute node. The output should look similar to that shown below. If the output shows missing or invalid values, refer the issue to your network administrator.

```
Ca 2 "H-00212800013e8f4a" # " xdb1db02 S 192.168.229.85 HCA-1"
```

Extra or Missing Components in the Newly Discovered Exadata Database Machine

If the list of components on the **Components** page of the discovery wizard contains extra components or missing components, then check the following troubleshooting steps:

- For extra components, examine them for Exadata Database Machine membership. Deselect any unnecessary pre-selected components on the components page.
- Verify which schematic file that was used for discovery. Ensure that Enterprise Manager can read the latest xml file (for example, databasemachine.xml) on the compute node.
- For missing components, check the schematic file content.
- To generate a new schematic file, see [Oracle Support Doc ID 1684431.1](#).

ILOM, PDU, or Cisco Switch is not Discovered

If the ILOM, PDU, or Cisco switch is not discovered, the most likely cause is that the Exadata Database Machine Schematic file cannot be read or has incorrect data. See [Troubleshooting the Exadata Database Machine Schematic File](#).

Test Connection Fails in PDU

If Test connection gives an error message Wrong credentials Please provide correct PDU user name and password..., ignore the message and proceed with the discovery. The PDU target will be in the down status initially. To correct the failure, you can try one of the following methods:

- Add the SNMP subscriptions to the PDU manually by following the steps in [Enable SNMPv3 on PDUs](#). The target will be backup when SNMP starts collecting the availability metrics.
- Execute the below command from the agent that is monitoring the target:

```
emctl control agent runCollection <pdu-target-name>:oracle_si_pdu
oracle_si_pdu_snmp_config emctl upload agent
```

Target Does not Appear in the Components Step of Guided Discovery

Even though no error may appear on the **Prerequisite Check** page during the Exadata Database Machine guided discovery process, the target does not appear on the Components page. Possible causes and solutions include:

- Check the All Targets page to make sure that the target has not been added as an Enterprise Manager target already:
 - Log in to Enterprise Manager.
 - Select **Targets**, then **All Targets**.
 - On the All Targets page, check to see if the target that was not shown on Components selection page, appears in the list.
- A target that is added manually may not be connected to the Exadata Database Machine system target through association. To correct this problem:
 - Delete these targets before initiating the Exadata Database Machine guided discovery.
 - Alternatively, use the `emcli` command to add these targets to the appropriate system target as members.

Target Status is Down or Metric Collection Error After Discovery

After the Exadata Database Machine guided discovery, an error that the target is down or that there is a problem with the metric collection may display. Possible causes and recommended solutions include:

- For the Exadata Storage Server or InfiniBand switch, the setup of SSH may not be configured properly. To troubleshoot and resolve this problem:
 - The agent's SSH public key in the `<AGENT_USER_HOME>/ .ssh/id_rsa.pub` file is not in the `authorized_keys` file of `$HOME/.ssh` for `ilom-admin`.

For the Exadata Storage Server monitored using CellCLI: The agent's SSH public key in the `<AGENT_USER_HOME>/ .ssh/id_rsa.pub` file is not in the `$HOME/.ssh/authorized_keys` for the `cellmonitor` user on Exadata Storage Server.

For the Exadata Storage Server monitored using ExaCLI/RESTAPI: Exadata Storage Server credentials are not set as monitoring credentials for the Exadata Storage Server target.

For InfiniBand Switch: The agent's SSH public key in the `<AGENT_USER_HOME>/ .ssh/id_rsa.pub` file is not in the `$home/.ssh/authorized_keys` file for `ilom-admin` user on Infiniband Switch.

- Verify permissions. The permission settings for `.ssh` and `authorized_keys` should be:
`drwx----- 2 cellmonitor cellmonitor 4096 Oct 13 07:06 .ssh`
`-rw-r--r-- 1 cellmonitor cellmonitor 441842 Nov 10 20:03 authorized_keys`
- Resolve a `PerformOperationException` error. See [Troubleshooting the Exadata Database Machine Schematic File](#) for more information.
- If the SSH setup is confirmed to be properly configured, but the target status is still down, then check to make sure there are valid monitoring and backup agents assigned to monitor the target. To confirm, click the **Database Machine** menu and select **Monitoring Agents**.
- After confirming that valid agents are assigned to the Exadata Storage Server, check **Target Down Reason** in the Response metric to understand the root cause. To view the Response metric, go to Exadata Storage Server target menu, click **Monitoring**, click **All Metrics**, and select **Response**.
- For the ILOM, PDU, or Cisco switch, possible causes include:
 - The Exadata Database Machine Schematic Diagram file has the wrong IP address.
 - Monitoring Credentials is not set or incorrect.
- To verify or to set Monitoring Credentials:
 - Log in to Enterprise Manager.
 - Click **Setup**, then **Security**, and finally **Monitoring Credentials**.
 - On the Monitoring Credentials page, select your target type, and click **Manage Monitoring Credentials**. Verify that the monitoring credentials are set for the target. If not, then set the monitoring credentials.

ILOM Credential Validation Fails During Discovery

ILOM Credential Validation Failure Errors

ILOM credential validation may fail while performing a 13c discovery. The following error may occur:

Authentication failed

The possible cause of the error is that the credentials provided are invalid.

To resolve the issue, use valid credentials.

Discovery Process Hangs

If the discovery process for the Exadata Database Machine hangs, then check the following troubleshooting steps:

- Examine your network to verify:
 - That the host name can be resolved.
 - That the Agent(s) can access the OMS.
 - That a simple job can be executed from the console.
- If the OMS reported any errors, then check the following log file:

```
$MW_HOME/gc_inst/sysman/log/emoms.log
```

- For Repository issues, check the Repository database's `alert.log` file.
- For Agent issues, check the following log file on the monitoring agents:

```
$AGENT/agent_inst/sysman/log/gcagent.log
```

SNMP Configuration Missing

Set Up SNMP Subscription for Exadata Storage Servers

Exadata Storage Server supports SNMP v1/v2 and v3 version based alert notifications to the subscribers. To receive alerts in the SNMP V3 version, the monitoring agents must be added to the `snmpsubscriber` list using a `snmpuser`. To set the values of `snmpSubscriber`, `notificationMethod` and `notificationPolicy`:

1. Set the `snmpSubscriber` value:

- **SNMP V1**

```
# cellcli -e "alter cell snmpSubscriber=((host='[host name]',port=[port]),(host='[host name]',port=[port]))"
```

Where `[host name]` and `[port]` values are Agents monitoring your Exadata Storage Server targets.

- **SNMP V3**

The process is same as above, but along with the host and port, an SNMP V3 user and type must be provided:

```
ALTER CELL snmpSubscriber=((host='host1.example.com',port=[port],snmpUser='[v3user]', type=V3))
```

Before you run the command, make sure that a `v3user` is a valid SNMP user on the Exadata Storage Server. You can check by running following command:

```
cellcli -e "list cell attributes snmpUser"
```

If there is no valid SNMP user on the Exadata Storage Server, then add a SNMP V3 user. See [Create SNMPv3 Users on Compute Nodes and Storage Servers](#).

Note

Take special care not to overwrite existing settings for `snmpSubscriber`. If there are existing subscribers, then append the agent subscriptions. Get the list of existing SNMP subscribers using the command `cellcli -e list cell attributes snmpSubscriber` as below:

```
# cellcli -e list cell attributes snmpSubscriber  
((host=ilm-asr.example.com,port=162,community=public,type=asr))
```

Then you must append the new agent subscriptions to the existing ones and run the command `alter cell snmpsubscriber` as shown below:

```
#cellcli -e "alter cell snmpSubscriber=((host='ilm-asr.example.com',port=162,community=public,type=asr),(host='[host name]',port=[port]),(host='[host name]',port=[port]))"
```

2. For SNMP V3, update the following in Enterprise Manager:
 - Ensure that the SNMP V3 monitoring credentials set on the Exadata Storage Server target in Enterprise Manager match the SNMP V3 user on the Exadata Storage Server:
 - Log in to Enterprise Manager.
 - From the **Setup** menu, select **Security**, then select **Monitoring Credentials**.
 - Select the **Exadata Storage Server** type and click **Manage Monitoring Credentials**.
 - Select the Exadata Storage Server and set the SNMP V3 credentials.
 - Ensure that the Monitoring Configuration property *SNMP Engine ID* is set up for the Exadata Storage Server target in Enterprise Manager:
 - Log in to Enterprise Manager.
 - Go to the Exadata Storage Server home page.
 - From the Exadata Storage Server menu, select **Monitoring Configuration**.
 - If the Engine ID property is not set, then set it manually in Enterprise Manager. Obtain it from the Exadata Storage Server with the following command

```
CellCLI> list cell attributes snmpengineidhex ;
```

All of the above steps are automatically taken care of, if the SNMP V3 subscription is selected in the Exadata Storage Server Credentials region that appears as part of the Exadata Database Machine discovery wizard.

3. Set the **notificationMethod** value:

```
# cellcli -e "alter cell notificationMethod='snmp,mail'"
```

4. Set the **notificationPolicy** value:

```
# cellcli -e "alter cell notificationPolicy='Critical,Warning,Clear'"
```

Set Up SNMP Subscription for Cisco Ethernet Switch Targets

The Cisco Ethernet Switch must be configured to allow the Agents that monitor it to be able to both poll the switch and to receive SNMP alerts from the switch. To allow this, perform the following steps (swapping the example switch name `dm01sw-ip` with the name of the Cisco Ethernet Switch target being configured):

Note

This procedure is valid for SI targets if the monitoring of the switch is performed with a non-administrator user. If Enterprise Manager monitors the switch with an administrator user, the following procedure is automatically performed as part of the discovery process.

1. Log in to the Cisco switch and enter Configure mode:

```
# ssh cisco-switch-host -l <admin user>
User Access Verification Password:
dm01sw-ip> enable
Password:
dm01sw-ip# configure terminal
```

Enter configuration commands, one per line. End with CNTL/Z.
dm01sw-ip(config)#

2. Enable access to allow the Agents monitoring Cisco Switch target to poll the switch.

In the command, [EMagentIPAddr] is the IP address of the server where the Enterprise Manager Agent is running. See [Create SNMPv3 Users on Cisco Ethernet Switches](#) and [Create SNMP Community on Cisco Ethernet Switches](#).

3. Set the monitoring Agent as the location where SNMP alerts are delivered. The SNMP community specified for a monitoring agent on the Cisco Switch must match the value provided during Enterprise Manager Cisco Switch Management setup:

```
snmp-server host <EM_Agent_IP_Address> version { 1 | 2c | 3 { auth |  
noauth | priv }} community-string [ udp-port <EMagentRecvltListenPort> ]
```

In the above command:

- **EM_Agent_IP_Address**: IP Address of the monitoring agent
- **version { 1 | 2c | 3 { auth | noauth | priv }}**: The SNMP version can be SNMPv1, SNMPv2c, or SNMPv3. The optional keywords {auth | noauth | priv} are used to specify the authentication level.
- **udp-port <EMagentRecvltListenPort>**: [EMagentRecvltListenPort] is the EMD_URL port of the emagent or SnmpRecvletListenNIC property value if it is enabled. Specifying udp-port is optional.
- **community-string**: For SNMPv1 and SNMPv2c, enter the password-like community string. In case of SNMPv3, enter the SNMPv3 user name. For example, in case of SNMPv3, the above command is changed to:

```
snmp-server host <EM_Agent_IP_Address> version {3 { auth | noauth |  
priv }} SNMPv3-user-name [ udp-port <EMagentRecvltListenPort> ]
```

4. Verify settings and save the configuration:

```
dm01sw-ip(config)# end  
dm01sw-ip# show running-config  
dm01sw-ip# copy running-config startup-config
```

Set Up SNMP for Power Distribution Unit (PDU) Targets

To enable Enterprise Manager to collect metric data and raise events for the PDU target, you must configure the PDU to accept SNMP queries from the Agents that monitor the PDU target. Also, appropriate threshold values for different phase values must be set on the PDU.

For steps to enable SNMPv3 on the PDU, see [Enable SNMPv3 on PDUs](#).

Time Out Error During SNMP Subscription

During the discovery of Exadata Database Machine, when the Exadata Storage Servers have ILOM 5.0 or *later*, and the SNMP subscription step throws a *Time Out* error, the following two fixes are recommended:

1. Upgrade Enterprise Manager Exadata plugin to 13.4 RU12, 13.5 RU01, or *later*.
2. Upgrade Exadata Storage Server software to 21.2.0.0.0 or *later* version.

Review Logs

You can review the logs to verify the proper discovery, functioning, and troubleshooting. Following are some of the useful logs and their locations:

- **OMS logs**

Location: `$INSTANCE_HOME/sysman/log`

- `emoms.log`: The main log file for the OMS Console application
- `emoms.trc`: The main trace file for the OMS Console application
- `emoms_pbs.log`: The main log file for the OMS platform application

- **EM Managed Server logs**

Location: Relative path from EM instance base `user_projects/domains/GCDomain/servers/EMGC_OMS1`

- `EMGC_OMS1.log`: The EMGC_OMS1 instance writes all messages from its subsystems and applications to this log file.
- `EMGC_OMS1-diagnostic.log`: This log file contains application-related security errors.

- **Monitoring agent logs**

Location: Relative path from AGENT INSTANCE HOME `sysman/log`

- `gcagent.log`: This log file contains trace, debug, information, error, or warning messages from the Agent. It can be used for debugging Agent framework issues.
- `gcagent_errors.log`: It is similar to `gcagent.log`, but it contains only the log messages of ERROR and FATAL levels. The log has no size limit.
- `emagent_perl.trc`: Trace file for the PERL scripts. EM uses Perl scripts to gather some metrics and target discovery. The log level can be changed by modifying `EMAGENT_PERL_TRACE_LEVEL` variable in `sysman/config/emd.properties` file.

EMCLI Discovery Issues

The deployment procedure output displays the error messages for *targets not promoted, tasks not performed* (like *SNMP subscription, access point creation, etc.*) and information to resolve those errors. After you resolve the underlying cause for the error, submit the deployment procedure again.

To triage an issue, enable the debug mode on OMS and the agent and resubmit the discovery to get more information:

- On OMS:

```
emctl set property -name log4j.rootCategory -value "DEBUG, emlogAppender, emtrcAppender" -sysman_pwd sysman
```

- On the agent:

```
emcli set_agent_property -agent_name=<agent_target_name> -value=DEBUG -name="Logger.sdklog.level"
```

The following are some of the errors that may be encountered and resolved:

- [System Infrastructure Targets in Pending Status](#)
- [Discovery Status Summary States "Completed with errors"](#)
- [Target Discovered but Association Not Created](#)
- [Component not Available or Accessible Before Discovery](#)
- [Trouble Using ExaCLI during Discovery](#)
- [Input File Errors](#)
- [Unclear Message Shown on the Deployment Procedure Output](#)

System Infrastructure Targets in Pending Status

System Infrastructure targets are in Pending status for a while until the first configuration metrics are collected.

If the status *Pending* is displayed for the discovered targets even after 15 minutes, then go to the target home page and investigate the cause and attempt to fix them manually. For example, metric collection errors, etc.

In most cases, the issue gets resolved without rediscovery if you wait long enough for the target metrics to be collected.

Discovery Status Summary States "Completed with errors"

If the EMCLI Deployment Procedure step **System Creation** completes with a status of *Completed with Errors*, then the summary information about the errors encountered will be displayed in the procedure output. For example:

```
Number of targets that have been created successfully:8
Number of targets that have errors during discovery:1
```

```
List of targets with errors during discovery:
Exatarget_1.example.com
ExpressionEvaluationException: Below error message was returned while
executing this Computational step.
-----
Some targets are not successfully discovered. Please refer the above target
creation details for the detailed errors.
-----
```

Target Discovered but Association Not Created

If the discovery status reports that a target is created but its associations are not, then create associations manually by following these steps:

1. Add the target manually to the Exadata Database Machine Schematic from the database machine home page. This will associate the target with database machine rack.
2. Associate the target with Exadata Database Machine using the EMCLI command as in the following example:

```
emcli create_assoc -assoc_type='app_composite_contains' -
source='DB_Machine1.example.com:oracle_dbmachine'
-dest='DB_Machine1.example.com:host'
```

Component not Available or Accessible Before Discovery

If the Exadata Database Machine member target is not available or accessible, then try resolving the issue through one of the following options:

- Make component available before performing Database Machine discovery.
- If the component's availability is unresolved for an indefinite time, then exclude it from the Exadata Database Machine discovery. Use the input file parameter *component.skipComponentList* and proceed with the discovery. After the component is available, remove it from the *component.skipComponentList* parameter and resubmit the discovery to add the target to the existing Exadata Database Machine.

Trouble Using ExaCLI during Discovery

If you selected ExaCLI as the monitoring mechanism for Exadata Storage Servers and if the monitoring agent is installed on the compute nodes, then ExaCLI should already be installed. If not, ensure that ExaCLI is installed on the agent that you have selected.

To verify if ExaCLI is installed properly on the agent, run the following command:

```
[oracle@myagent ~]$ which exacli
/usr/local/bin/exacli
[oracle@myagent ~]$ exacli -c mycell.example.com -l exacli_user --xml -e
'list cell attributes name'
Password: *****
<?xml version="1.0" encoding="utf-8" ?>
<cli-output>
<context cell="mycell"/>
<cell> <name>mycell</name>
</cell>
</cli-output>
```

Input File Errors

- **Mandatory property name or value not specified in the input file**

The error is displayed indicating the properties that are missing as part of the input file. For example, Property `outputFileLoc` is not defined.

- **The specified value is not applicable for a property**

The error message lists out all allowed values for the property. For example, Check `cell metric` source input. It needs be `CellCLI` or `ExaCLI`.

- **The named credential specified in the input file does not exist**

The error is displayed indicating the named credential which is not available in Enterprise Manager. For example, Credential not found owner `SYSMAN` name `BLR_BOX_CREDS_1`.

- **The specified named credential is not valid**

The error is displayed indicating the credential failed during verification. For example, Credential test of cell target `Exa_Storage_Server_1` failed..

Resolve the issue based on the additional information available with the error, and rediscover the targets.

Resolve the above errors by updating the information in the input file by using the resolution steps suggested in the Deployment Procedure output. Resubmit the Deployment Procedure to discover targets.

Unclear Message Shown on the Deployment Procedure Output

Typically, detailed information about the error and resolution is displayed for identified use cases. In case an unclear message is displayed, to learn the cause of the error:

- Check the Oracle Management Service logs (instance logs) in the location `$INSTANCE_HOME/sysman/log` where the information would be logged.
- Check the Enterprise Manager server logs stored in the relative path from the Enterprise Manager instance base `user_projects/domains/GCDomain/servers/EMGC_OMS1`, for more information.

Troubleshooting the Exadata Database Machine Schematic File

As part of any discovery troubleshooting, possible causes and recommended resolution with the schematic file can include:

- The schematic file on the compute node is missing or is not readable by the agent user.

For Exadata Release 11.2.3.2 and later, the schematic file is:

```
/opt/oracle.SupportTools/onecommand/databasemachine.xml
```

Ensure that the schematic file exists in the specified location and that the file and directory can be accessed by the agent installation user account.

- If a `PerformOperationException` error appears, the agent NMO is not configured for `setuid-root`:

— From the OMS log:

```
2019-11-08 12:28:12,910 [[ACTIVE] ExecuteThread: '6' for queue: 'weblogic.kernel.Default (self-tuning)']  
ERROR model.DiscoveredTarget logp.251 -  
ERROR: NMO not setuid-root (Unix only)  
oracle.sysman.emSDK.agent.client.exception.PerformOperationException:
```

— As root, run:

```
# <AGENT_INST>/root.sh
```

- In the `/etc/pam.d` file, `pam_ldap.so` is used instead of `pam_unix.so` on compute nodes.

— Even though the agent user and password are correct, this error appears in the agent log:

```
oracle.sysman.emSDK.agent.client.exception.PerformOperationException:  
ERROR: Invalid username and/or password
```

To resolve the issue, in `/etc/pam.d` file use `pam_unix.so` on compute nodes.

- If the schematic file is blank, then:
 - Check your browser support and Enterprise Manager Cloud Control 13c.
 - Run through discovery again and watch for messages.
 - Check the `emoms.log` file for exceptions at the same time.
- If components are missing, then:

- Add manually to the schematic page (click **Edit**).
- Check for component presence in Enterprise Manager. Check to see if it is monitored.

Exadata Database Machine Management Troubleshooting

If data is missing in Resource Utilization graphs, then run a "view object" SQL query to find out what data is missing. Common problems include:

- Schematic file is not loaded correctly.
- Cluster, Database, and ASM are not added as Enterprise Manager targets.
- Database or Exadata Storage Server target is down or is returning metric collection errors.
- Metric is collected in the Enterprise Manager repository, but has an `IS_CURRENT != Y` setting.

Oracle Database Storage Server System Target Status Pending

The **Pending** state of the Oracle Database Storage Server System indicates that the storage server associations with the system could not be derived. The derived association rule is based on the ECM data of Exadata Storage Servers, DB, and ASM. Generally, after adding a target, it may take up to 30 minutes for ECM data to appear.

The **Last Collected** information on the **Configuration** page indicates the ECM data availability. To view the information, from the target menu, select **Configuration**, and click **Latest**.

If the configuration data is not present for any of the targets like Database Machine, DB, and ASM, then perform a **Refresh** on the configuration page.

Target Status Issues

If the Target status shows DOWN inaccurately, then:

- If you are using `cellCLI` to monitor the Exadata Storage Server, you can check SSH equivalence (`cellmonitor user`) with the following command:

```
ssh -i /home/oracle/.ssh/id_dsa -l cellmonitor <cell name> -e cellcli list cell
```

Output should be: `<cell name>`
- For the **PDU**: Check to make sure you can access the PDU through your browser to verify that it is connected to your LAN:

```
http://<pdu name>
```
- For the **Cisco Switch**: Check for proper SNMP subscriptions. See [Set Up SNMP Subscription for Cisco Ethernet Switch Targets](#) for details.

Metric Collection Issues

If the Target status shows a Metric Collection Error, then:

- Hover over the icon or navigate to Incident Manager.
- Read the full text of the error.

- Visit the Monitoring Configuration page and examine the settings. From the **Setup** menu, select **Monitoring Configuration**.
- Trigger a new collection: From the **Target** menu, select **Configuration**, then select **Last Collected**, then **Actions**, and finally select **Refresh**.
- Access the monitoring Agent Metric through your browser:
`https://<Agent_Host_Name>/emd/browser/main`
Click **Target >** and then click **Response** to evaluate the results. You may need to log a service request (SR) with Oracle Support.

Status: Pending Issues

For those issues where components are in a Pending status, see the following troubleshooting steps:

- [Cells sys Targets](#)
- [Database Machine Target or Any Associated Components](#)

Cells sys Targets

If the Cells sys target seems to be in a Pending status for too long, then:

- Verify that there is an association for the Cluster ASM, Database, and Exadata Storage Servers.
- Check and fix the status of the associated target database.
- Check and fix the status of the associated target ASM cluster.
- Ensure an **UP** status for all Exadata Storage Server targets.
- Delete any unassociated cells sys targets.

Database Machine Target or Any Associated Components

If the Exadata Database Machine target or any associated components are in a Pending status for too long, then:

- Check for duplicate or pending delete targets. From the **Setup** menu, select **Manage Cloud Control**, then select **Health Overview**.
- Check the target configuration. From the target's home page menu, select **Target Setup**, then select **Monitoring Configurations**.
- Search for the target name in the agent or OMS logs:

```
$ grep <target name> gcagent.log or emoms.log
```

Monitoring Agent Not Deployed for IPv6 Environments

Problem: For IPv6 environments, the monitoring agent is not deployed.

Cause: If the IPv6 address is not included in the `/etc/hosts` file, then the agent will not be deployed.

Resolution: Edit the `/etc/hosts` of the compute node (or the VM in case of virtual Exadata) to map the OMS host name to an IPv6 address.

Monitoring: Error Fetching Data for Grid and CellSys Targets IORM Page Charts

When you encounter **Error Fetching Data** in the IORM page *Workload distribution by Databases* of Grid and CellSys targets charts in the first 24 hours of Database Machine discovery, to resolve, you can wait 24 hours after the Database Machine is discovered or perform the following steps:

1. Ensure that the Grid and CellSys member targets configuration collection is collected. If the configuration is not available, perform the below steps:
 - a. Go to Exadata Grid target menu, **Configuration, Latest**, click **Refresh**.
 - b. Select the member targets configuration collection and click **OK**.
 - c. Wait until the configuration collection is collected successfully.
2. Update the collection frequency of the **Exadata Storage Type** metric to 1 minute for every member of the Grid and CellSys target. Follow the below steps:
 - a. Go to Exadata Storage Server target menu, **Monitoring**, and **Metric and Collection Settings**.
 - b. Click the **Other Collected Items** tab. Update the **Exadata Storage Type** metric collection frequency to 1 minute and save it.
3. Wait for 10 minutes, launch the Grid and CellSys target home page, and ensure that the **Overview** section does not display the message *Configuration collection is not available*. If the message continues to appear, then repeat *Steps 1 and 2*.

After the issue is resolved for Grid and CellSys IORM page charts, set the **Exadata Storage Type** metric collection frequency to 1440 minutes.

Not able to receive SNMP traps on Exadata Storage Servers using IPv6

For IPv6 environments, the Enterprise Manager agent needs have SNMP v3 subscription to the Exadata Storage Servers, for complete monitoring.

Prerequisites:

- SNMP V3 user created on Exadata Storage Server. See **Step 1** for instructions.
- Exadata Storage Server version must be 12.1.2.2 or later to support SNMP V3 subscription.

If the SNMP subscription was missed during Exadata discovery, you can follow these steps:

To configure SNMP v3 subscription on an Exadata Storage Server:

1. Create an SNMP V3 user on Exadata Storage Server Target.
 - Suppose you need a user called "v3user". You can check whether this user is already present on the Exadata Storage Server using this command:

```
cellcli -e "list cell attributes snmpUser"
```

- If this user is not present on the Exadata Storage Server, create a new user. See [Create SNMPv3 Users on Compute Nodes and Storage Servers](#).

 **Note**

The Enterprise Manager Agent supports DES and AES privacy protocols in addition to SHA authentication protocol.

2. If you need to edit the SNMP V3 monitoring credentials of the Exadata Storage Server after discovery, you can log in to Enterprise Manager and:
 - Select **Setup**, then **Security**, then **Monitoring Credentials**.
 - Select **Oracle Exadata Storage Server** type, and click on **Manage Monitoring Credentials** button.
 - Select the corresponding row for the **Exadata Storage Server** target and **SNMP V3 Monitoring Credentials** credential set and click the **Set Credentials** button.
 - Populate the **Enter Monitoring Credentials** form with the appropriate information and click **Save**.

 **Note**

Ensure that the SNMP V3 monitoring credentials are setup on the Exadata Storage Server target in Enterprise Manager which match the SNMP V3 user on the Exadata Storage Server.

Exascale Discovery and Monitoring Issues

- **Exascale target status is down**
Go to the **Oracle Exadata Exascale**, **Target Setup**, select **Monitoring Configuration**. Check the REST URL which should be of the form `https://<fqdn_of_controlserver>:<port>/exascaleapi/v1`.
- **Exascale target status is *Availability Evaluation Error (Failed to get private key)***
Go to **Setup**, **Security**, select **Monitoring Credentials**. Select **Oracle Exadata Exascale** row and click **Manage Monitoring Credentials**. Select the Exascale target that has the issue and click **Set Credentials**. Fix the credential by providing the correct value.
- **Exascale target status is *Availability Evaluation Error (the oracle_exascale target <target name> does not exist)***
Ensure that both the discovery and monitoring components of the Exadata plug-in are patched with Oracle Enterprise Manager 13.5GC RU23.

Index

A

alerts
 InfiniBand, [33](#)
ALTER CELL command, [2](#)

C

Cell management, [25](#)
Cisco Ethernet Switch
 SNMP configuration, [7](#)
creating roles
 user roles
 creating, [1](#)

D

dashboard
 public report, [54](#)
dashboard creation, [53](#)
Database Machines
 alerts, [12](#)
delete Exadata Database Machine component, [15](#)
deploy plug-in, [13](#)
discovery
 SPARC SuperCluster, [9](#)
 switch troubleshooting, [3](#)
Disk I/O Objective, [28](#)

E

Exadata Cell management, [25](#)
Exadata component
 delete, [15](#)
Exadata Database Machine
 configuration support, [4](#)
 dashboard, [53](#)
 remove target, [13](#)
 troubleshooting, [13](#)
 view topology, [3](#)
Exadata Storage Server administration, [27](#)
Exadata Storage Server SNMP configuration, [1](#)

F

Fault Monitoring, [41](#)
features, [1](#)
 hardware support, [2](#)
 monitoring and notification, [1](#)
 target discovery, [3](#)
firewall configuration, [11](#)

H

hardware
 supported hardware, [4](#)
 view critical information, [5](#)
hardware not supported, [7](#)
hardware support features, [2](#)

I

InfiniBand network management, [31](#)
InfiniBand Switch
 configure for SNMP, [5](#)

M

manually, [13](#)
metrics and alert settings, [24](#)
monitoring and notification features, [1](#)
multi-rack support, [5](#)

N

names resolution, [10](#)

O

Oracle SuperCluster support, [5](#)

P

prerequisites
 Exadata Storage Server software, [9](#)
 firewall configuration, [11](#)
 InfiniBand Switch, [9](#)
 PDU firmware, [10](#)

R

remove Exadata Database Machine target, [13](#)
remove SNMP subscription, [4](#)

S

schematic file
 troubleshooting, [12](#)
SNMP
 configure using cellcli, [6](#)
 configure using dcli, [4](#)
SNMP configuration
 Cisco Ethernet Switch, [7](#)
 InfiniBand Switch, [5](#)
SNMP subscription
 remove, [4](#)
SPARC SuperCluster
 discovery, [9](#)
SSH connectivity, [4](#)

supported hardware, [4, 7](#)

 partitioned Exadata Database Machine, [5](#)
supported operating systems, [7](#)

T

target discovery features, [3](#)

troubleshooting
 discovery, [1](#)
 Exadata Database Machine, [13](#)
 schematic file, [12](#)

U

user roles, [13](#)

V

view topology, [3](#)