

Oracle® Exadata Database Machine

Installation and Configuration Guide for Exadata Database Machine



20.1.0
F29249-06
February 2021

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Oracle Exadata Database Machine Installation and Configuration Guide for Exadata Database Machine, 20.1.0

F29249-06

Copyright © 2008, 2021, Oracle and/or its affiliates.

Primary Authors: Peter Fusek, James Spiller

Contributing Authors: Craig Brown, Glenn Maxey

Contributors: Doug Archambault, Leo Agranonik, Andrew Babb, Arturo Ceron, Nilesh Choudhury, Henry Chow, Ravindra Dani, Boris Erlikhman, Jaime Figueroa, Gurmeet Goindi, Roger Hansen, Leslie Keller, Frank Kobylanski, René Kundersma, Holger Leister, Yang Liu, Juan Loaiza, Barb Lundhild, Catherine Luu, Philip Newlan, Adrian Ng, Dan Norris, Michael Nowak, Gavin Parish, Hector Pujol, Dmitry Potapov, Darryl Presley, Ashish Ray, Richard Scales, Oliver Sharwood, Jia Shi, Kesavan Srinivasan, Krishnadev Telikicherla, Cliff Thomas, Alex Tsukerman, Kothanda Umamageswaran, Doug Utzig, Zheren Zhang

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 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" or "commercial computer software documentation" 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 and Java 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	xiii
Documentation Accessibility	xiii
Related Documentation	xiii
Conventions	xv

1 Site Requirements for Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack

1.1	General Environmental Requirements	1-2
1.1.1	General Environmental Requirements for Oracle Exadata Database Machine X6 and Later	1-2
1.1.2	General Environmental Requirements for Oracle Exadata Database Machine X5-2	1-3
1.1.3	General Environmental Requirements for Oracle Exadata Database Machine X5-8	1-5
1.1.4	General Environmental Requirements for Oracle Exadata Database Machine X4-2	1-6
1.1.5	General Environmental Requirements for Oracle Exadata Database Machine X4-8 with Exadata Storage Server X5-2 Servers	1-9
1.1.6	General Environmental Requirements for Oracle Exadata Database Machine X4-8 Full Rack	1-10
1.1.7	General Environmental Requirements for Oracle Exadata Database Machine X3-2	1-11
1.1.8	General Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers	1-14
1.1.9	General Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X3-2 Servers	1-15
1.1.10	General Environmental Requirements for Oracle Exadata Database Machine X2-2	1-17
1.1.11	General Environmental Requirements for Oracle Exadata Database Machine X2-8 Full Rack	1-18
1.1.12	General Environmental Requirements for Oracle Exadata Storage Expansion Rack X6-2 and Later	1-19
1.1.13	General Environmental Requirements for Oracle Exadata Storage Expansion Rack X5-2	1-19

1.1.14	General Environmental Requirements for Oracle Exadata Storage Expansion Rack X4-2	1-21
1.1.15	General Environmental Requirements for Oracle Exadata Storage Expansion Rack X3-2	1-23
1.1.16	General Environmental Requirements for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers	1-25
1.1.17	General Environmental Requirements for Single Servers	1-25
1.2	Space Requirements	1-29
1.2.1	Space Requirements for Oracle Exadata Database Machine X7-2 and Later Systems	1-30
1.2.2	Space Requirements for Racks up to Oracle Exadata Database Machine X6	1-30
1.3	Receiving, Unpacking, and Access Route Requirements	1-30
1.3.1	Rack Weights for Oracle Exadata Database Machine	1-31
1.3.1.1	Rack Weights for Oracle Exadata Database Machine 2-Socket Servers	1-32
1.3.1.2	Rack Weights for Oracle Exadata Database Machine 8-Socket Servers	1-33
1.3.2	Rack Weights for Oracle Exadata Storage Expansion Rack	1-35
1.4	Maintenance Access Requirements	1-37
1.5	Flooring Requirements	1-38
1.6	Electrical Power Requirements	1-38
1.6.1	PDU Power Requirements	1-39
1.6.1.1	Low-voltage 15 kVA Single Phase PDUs for North America, South America, Japan and Taiwan	1-39
1.6.1.2	Low-voltage 15 kVA Three Phase PDUs for North America, South America, Japan and Taiwan	1-40
1.6.1.3	High-voltage 15 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	1-41
1.6.1.4	High-voltage 15 kVA Three Phase for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	1-42
1.6.1.5	Low-voltage 22 kVA Single Phase PDUs for North America, South America, Japan and Taiwan	1-43
1.6.1.6	High-voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	1-44
1.6.1.7	Low-voltage 24 kVA Three Phase PDUs for North America, South America, Japan and Taiwan	1-45
1.6.1.8	High-voltage 24 kVA Three Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	1-45
1.6.2	Facility Power Requirements	1-46
1.6.3	Circuit Breaker Requirements	1-46

1.6.4	Electrical Grounding Guidelines	1-47
1.7	Temperature and Humidity Requirements	1-47
1.8	Ventilation and Cooling Requirements	1-48
1.9	Network Connection and IP Address Requirements for Oracle Exadata Database Machine	1-50
1.9.1	Network Connection Requirements for Oracle Exadata Database Machine	1-50
1.9.2	DNS Configuration for Oracle Exadata Database Machine	1-53
1.9.3	IP Address Requirements for Oracle Exadata Database Machine	1-54
1.9.3.1	IP Address Requirements for Oracle Exadata Database Machine Systems with RoCE Network Fabric (X8M)	1-55
1.9.3.2	IP Address Requirements for Oracle Exadata Database Machine X5 to X8 Systems with InfiniBand Network Fabric	1-59
1.9.3.3	IP Address Requirements for Oracle Exadata Database Machine X4 Systems	1-64
1.9.3.4	IP Address Requirements for Oracle Exadata Database Machine X3 Systems	1-68
1.9.3.5	IP Address Requirements for Oracle Exadata Database Machine X2 Systems	1-73
1.9.4	Calculating the IP Address Requirements when Using Oracle VMs	1-76
1.10	Network Connection and IP Address Requirements for Oracle Exadata Storage Expansion Rack	1-79
1.11	Ensuring That the Site is Ready	1-82

2 Understanding the Network Requirements for Oracle Exadata Database Machine

2.1	Overview of Network Requirements	2-1
2.2	Network Channel Bonding Support	2-8
2.2.1	Bonded Network Configurations	2-8
2.2.1.1	Bonded Network Configuration for Oracle Exadata Database Machine Two-Socket Systems	2-8
2.2.1.2	Bonded Network Configuration for Oracle Exadata Database Machine X7-8, X8-8, and X8M-8	2-10
2.2.1.3	Bonded Network Configuration for Oracle Exadata Database Machine X4-8, X5-8, and X6-8	2-11
2.2.1.4	Bonded Network Configuration for Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack	2-11
2.2.2	Non-bonded Network Configuration	2-12
2.2.2.1	Non-bonded Network Configuration for Oracle Exadata Database Machine Two-Socket Systems	2-13
2.2.2.2	Non-bonded Network Configuration for Oracle Exadata Database Machine X5-8 and X6-8	2-14

2.2.2.3	Non-bonded Network Configuration for Oracle Exadata Database Machine X4-8 Full Rack	2-15
2.2.2.4	Non-bonded Network Configuration for Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack	2-15
2.2.3	Additional Networks	2-16
2.3	Network Partitioning on Oracle Exadata Database Machine	2-17
2.3.1	VLAN Support on Customer-Facing Networks	2-17
2.3.2	Access VLAN Support with RoCE Network Fabric	2-18
2.3.3	Using Exadata Secure RDMA Fabric Isolation	2-19
2.3.4	Using InfiniBand Partitioning for Network Isolation with InfiniBand Network Fabric	2-21
2.4	Configuring a Separate Network for ILOM	2-22
2.5	Default IP Addresses	2-22
2.6	Port Assignments When Using a Firewall	2-23

3 Using Oracle Exadata Deployment Assistant

3.1	Overview of Oracle Exadata Deployment Assistant	3-1
3.1.1	Considerations and Requirement for Using OEDA	3-2
3.2	OEDA Browser-based User Interface	3-4
3.2.1	Using the Browser-based Version of Oracle Exadata Deployment Assistant	3-6
3.2.2	Select Hardware Page	3-7
3.2.3	Choose Operating System Page	3-13
3.2.4	Rack Networks Page	3-14
3.2.5	Users and Groups	3-19
3.2.6	Define Clusters	3-23
3.2.7	Diskgroups	3-25
3.2.8	Database Home	3-27
3.2.9	Database	3-28
3.2.10	Cluster Networks	3-30
3.2.11	Alerting	3-39
3.2.12	Tape Library Connectivity	3-41
3.2.13	Comments	3-42

4 OEDA Command Line Interface

4.1	About the OEDA Command Line Interface	4-1
4.2	Starting the OEDACLI Utility	4-3
4.2.1	Displaying Help Output for OEDACLI	4-4
4.2.2	Displaying Help Output for an OEDACLI Command	4-5
4.3	OEDACLI Command Reference	4-6

4.3.1	OEDACLI Commands	4-7
4.3.1.1	LOAD FILE	4-7
4.3.1.2	SAVE ACTION	4-7
4.3.1.3	MERGE ACTIONS	4-8
4.3.1.4	DEPLOY ACTIONS	4-8
4.3.1.5	SAVE FILE	4-9
4.3.1.6	SAVE FILES	4-9
4.3.1.7	SET SECURITY	4-9
4.3.1.8	SET SU	4-10
4.3.1.9	LIST STEPS	4-10
4.3.1.10	LIST XMLACTIONS	4-12
4.3.2	ACFSVOLUME	4-12
4.3.2.1	ADD ACFSVOLUME	4-12
4.3.2.2	ALTER ACFSVOLUME	4-13
4.3.2.3	DELETE ACFSVOLUME	4-13
4.3.2.4	LIST ACFSVOLUME	4-14
4.3.3	CELL	4-15
4.3.3.1	CLONE CELL	4-15
4.3.3.2	CLONE NEWCELL	4-16
4.3.3.3	DELETE CELL	4-18
4.3.3.4	LIST CELLS	4-18
4.3.4	CLUSTER	4-18
4.3.4.1	ALTER CLUSTER	4-18
4.3.4.2	CREATE CLUSTER	4-20
4.3.4.3	DESTROY CLUSTER	4-21
4.3.4.4	LIST CLUSTERS	4-21
4.3.4.5	UPGRADE CLUSTER	4-22
4.3.5	COMPUTE	4-23
4.3.5.1	CLONE COMPUTE	4-23
4.3.5.2	DELETE COMPUTE	4-26
4.3.6	DATABASE	4-27
4.3.6.1	ADD DATABASE	4-27
4.3.6.2	ALTER DATABASE	4-30
4.3.6.3	DELETE DATABASE	4-31
4.3.6.4	LIST DATABASES	4-33
4.3.7	DATABASE HOME	4-33
4.3.7.1	ADD DATABASEHOME	4-33
4.3.7.2	ALTER DATABASEHOME	4-36
4.3.7.3	DELETE DATABASEHOME	4-37
4.3.7.4	LIST DATABASEHOMES	4-38
4.3.8	DISKGROUP	4-39

4.3.8.1	ADD DISKGROUP	4-39
4.3.8.2	ALTER DISKGROUP	4-40
4.3.8.3	ALTER DISKGROUPS	4-43
4.3.8.4	DELETE DISKGROUP	4-44
4.3.8.5	LIST DISKGROUPS	4-45
4.3.9	ES	4-45
4.3.9.1	ALTER ES	4-45
4.3.9.2	DISCOVER ES	4-46
4.3.10	GUEST	4-47
4.3.10.1	CLONE GUEST	4-47
4.3.10.2	DELETE GUEST	4-49
4.3.11	ILOM	4-50
4.3.11.1	ALTER ILOM	4-50
4.3.11.2	LIST ILOMS	4-50
4.3.12	MACHINE	4-51
4.3.12.1	ALTER MACHINE	4-51
4.3.12.2	LIST MACHINES	4-52
4.3.13	NETWORK	4-52
4.3.13.1	ADD NETWORK	4-53
4.3.13.2	ALTER NETWORK	4-54
4.3.13.3	DELETE NETWORK	4-56
4.3.13.4	LIST NETWORKS	4-57
4.3.14	SCAN	4-57
4.3.14.1	ADD SCAN	4-57
4.3.14.2	ALTER SCAN	4-58
4.3.14.3	DELETE SCAN	4-59
4.3.14.4	LIST SCANS	4-59
4.3.15	SSHKEYS	4-60
4.3.15.1	DEPLOY SSHKEYS	4-60
4.3.15.2	GENERATE SSHKEYS	4-61
4.3.15.3	SET SSHKEYS	4-61
4.3.16	SWITCH	4-62
4.3.16.1	ALTER SWITCH	4-62
4.3.16.2	LIST SWITCHES	4-62
4.3.17	VIP	4-63
4.3.17.1	ADD VIP	4-63
4.3.17.2	ALTER VIP	4-64
4.3.17.3	DELETE VIP	4-64
4.3.17.4	LIST VIPS	4-65

5 Installing Oracle Exadata Database Machine or Oracle Exadata Storage Expansion Rack at the Site

5.1	Reviewing Safety Guidelines	5-1
5.2	Unpacking Oracle Exadata Rack	5-3
5.2.1	Contents of the Shipping Kit for X7 and Later Racks	5-4
5.2.2	Contents of the Shipping Kit for X6 and Earlier Racks	5-4
5.2.3	Removing Oracle Exadata Rack from the Shipping Crate	5-5
5.3	Placing Oracle Exadata Rack in Its Allocated Space	5-6
5.3.1	Moving Oracle Exadata Rack	5-7
5.3.2	Stabilizing Oracle Exadata Rack	5-8
5.3.2.1	Stabilize Oracle Exadata Rack with Leveling Feet	5-8
5.3.2.2	Stabilize Oracle Exadata Rack X7 and Later with Leveling Feet	5-10
5.3.3	Attaching a Ground Cable (Optional)	5-12
5.4	Acclimating the Oracle Exadata Database Machine	5-13
5.5	Powering on the System the First Time	5-14
5.5.1	Inspecting the Machine After it is in Place	5-14
5.5.2	Connecting Power Cords	5-14
5.5.3	Powering on Oracle Exadata Rack	5-15

6 Configuring Oracle Exadata Database Machine

6.1	Verifying the Network Configuration Prior to Configuring the Rack	6-2
6.2	Configuring the RDMA Network Fabric Switch	6-3
6.2.1	Configuring Sun Datacenter InfiniBand Switch 36 Switch	6-3
6.2.2	Configuring the Cisco Nexus C9336C-FX2 Switch	6-7
6.2.3	Configuring the Cisco Nexus C9336C-FX2 Switches to Enable Exadata Secure RDMA Fabric Isolation	6-12
6.3	Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack	6-12
6.4	Configuring the Cisco Ethernet Switch	6-14
6.4.1	Configuring the Cisco Catalyst 4948 Ethernet Switch	6-15
6.4.2	Configuring the Cisco Nexus 93108-1G or 9348 Ethernet Switch	6-22
6.4.2.1	Performing the Initial Switch Configuration for the Cisco Nexus 93108-1G or 9348 Ethernet Switch	6-22
6.4.2.2	Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch	6-26
6.5	Configuring the Power Distribution Units	6-27
6.6	Configuring the Threshold Settings for the Power Distribution Units	6-29
6.6.1	PDU Thresholds for Oracle Exadata Database Machine X5-2 and Later	6-30
6.6.2	PDU Thresholds for Oracle Exadata Database Machine X4-2	6-30
6.6.2.1	PDU Thresholds for Oracle Exadata Database Machine X4-2 Full Rack	6-31

6.6.2.2	PDU Thresholds for Oracle Exadata Database Machine X4-2 Half Rack	6-32
6.6.2.3	PDU Thresholds for Oracle Exadata Database Machine X4-2 Quarter Rack	6-34
6.6.3	PDU Thresholds for Oracle Exadata Database Machine X3-2	6-36
6.6.3.1	PDU Thresholds for Oracle Exadata Database Machine X3-2 Full Rack	6-36
6.6.3.2	PDU Thresholds for Oracle Exadata Database Machine X3-2 Half Rack	6-38
6.6.3.3	PDU Thresholds for Oracle Exadata Database Machine X3-2 Quarter Rack	6-40
6.6.4	PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers)	6-42
6.6.4.1	PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack	6-43
6.6.4.2	PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack	6-45
6.6.4.3	PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack	6-47
6.6.5	PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers)	6-49
6.6.5.1	PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack	6-49
6.6.5.2	PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack	6-51
6.6.5.3	PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack	6-53
6.6.6	PDU Thresholds for Oracle Exadata Database Machine X4-8 and Later	6-55
6.6.7	PDU Thresholds for Oracle Exadata Database Machine X3-8 Full Rack	6-55
6.6.8	PDU Thresholds for Oracle Exadata Database Machine X2-8 Full Rack	6-58
6.6.9	PDU Thresholds for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers	6-60
6.6.9.1	PDU Thresholds for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers	6-60
6.6.9.2	PDU Thresholds for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers	6-62
6.6.9.3	PDU Thresholds for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers	6-64
6.7	Checking Exadata Storage Servers	6-66
6.8	Checking Oracle Exadata Database Servers	6-70
6.9	Performing Additional Checks and Configuration	6-74
6.10	Verifying the RoCE Network Fabric Configuration	6-75
6.11	Verifying the InfiniBand Network Fabric Network	6-75
6.12	Imaging a New System	6-78
6.12.1	Using ISO Images to Image a New System	6-78

6.12.2	Using Elastic Configuration to Image a New System	6-81
6.12.3	Using ISO to Image a New System	6-83
6.12.4	Using PXE to Image a New System	6-89
6.12.4.1	Using PXE to Image a New System Using Release 12.1.2.2.0 or Later	6-90
6.12.4.2	Using PXE to Image a New System for Releases Earlier than 12.1.2.2.0	6-94
6.12.5	Using USB to Image a New System	6-98
6.12.6	Updates to Imaging Parameters	6-103
6.13	Performing Initial Elastic Configuration of Oracle Exadata Database Machine	6-104
6.14	Adding Additional Elastic Nodes to an Existing Rack	6-110
6.15	Using the OEDA setuprootssh Utility	6-112
6.16	Loading the Configuration Information and Installing the Software	6-112
6.16.1	Configuring Oracle Exadata Database Machine Using OEDA	6-113
6.16.2	Using Oracle Exadata Deployment Assistant in Conjunction with Key-Based Authentication	6-116
6.16.3	Using Oracle Exadata Deployment Assistant on Systems with Non-Default root Passwords	6-117
6.16.4	Configuring Oracle Database and Oracle ASM Instances for Oracle Exadata Database Machine Manually	6-118
6.16.4.1	Configuring the Compatible Parameter for a Database Instance	6-120
6.16.4.2	Configuring Initialization Parameters for an Oracle ASM Instance	6-120
6.16.4.3	Using the Same DB_UNIQUE_NAME for Multiple Database Instances	6-121
6.17	Using the OEDA changePassword Utility	6-122
6.18	Installing Oracle Enterprise Manager Cloud Control	6-122
6.19	Adding a VM Cluster to Oracle Exadata Database Machine Using OEDA	6-123

A Site Checklists

A.1	System Components Checklist	A-2
A.2	Data Center Room Checklist	A-3
A.3	Data Center Environment Checklist	A-4
A.4	Access Route Checklist	A-6
A.5	Facility Power Checklist	A-8
A.6	Safety Checklist	A-10
A.7	Logistics Checklist	A-11
A.8	Network Configuration Checklist	A-14
A.9	Auto Service Request Checklist	A-16
A.10	Re-racking Checklists	A-17
A.11	PDU Checklist for Re-racking	A-20
A.12	Recycling the Racks	A-21

B Database High Availability Checklist

Preface

This guide describes Oracle Exadata Database Machine for online transaction processing (OLTP) and enterprise data warehousing. It includes information about site planning and configuration, as well as physical, electrical, and environmental specifications.

- [Audience](#)
- [Documentation Accessibility](#)
- [Related Documentation](#)
- [Conventions](#)

Audience

This guide is intended for Oracle Exadata Database Machine customers and those responsible for data center site planning, configuration, and maintenance of 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.

Related Documentation

In addition to the Oracle Exadata Database Machine documentation set, the following guides contain hardware information for Oracle Exadata Database Machine:

- *Oracle Exadata Database Machine System Overview*
- *Oracle Exadata Database Machine Security Guide*
- *Oracle Engineered System Safety and Compliance Guide, Compliance Model No.: ESY27*
- *Oracle Engineered Systems Extending and Multi-Rack Cabling Guide*
- *Oracle Exadata Database Machine Maintenance Guide*

- *Oracle Exadata System Software User's Guide*
- *Oracle Exadata X8-2 Database Server Service Manual* at http://docs.oracle.com/cd/E93359_01/html/E93386/index.html
- *Oracle Exadata Storage Server X8-2 EF,HC, and XT Service Manual* at https://docs.oracle.com/cd/E93361_01/html/E93395/index.html
- *Oracle Exadata Database Server X8-8 Service Manual* at http://docs.oracle.com/cd/E93360_01/html/E96097/index.html
- *Oracle Server X7-2 Service Manual* at http://docs.oracle.com/cd/E72435_01/html/E72445/index.html
- *Oracle Server X7-2L Service Manual* at http://docs.oracle.com/cd/E72463_01/html/E72474/index.html
- *Oracle Server X7-8 Service Manual* at http://docs.oracle.com/cd/E71925_01/html/E71936/index.html
- *Oracle Server X6-2 Service Manual* at http://docs.oracle.com/cd/E62159_01/html/E62171/index.html
- *Oracle Server X6-2L Service Manual* at http://docs.oracle.com/cd/E62172_01/html/E62184/index.html
- *Oracle Server X5-2 Service Manual* at http://docs.oracle.com/cd/E41059_01/html/E48312/napsm.html
- *Oracle Server X5-2L Service Manual* at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.html#scrolltoc
- *Sun Server X4-8 Service Manual* at http://docs.oracle.com/cd/E40591_01/html/E40317/index.html
- *Sun Server X4-2 Service Manual* at http://docs.oracle.com/cd/E36975_01/html/E38045/gentextid-14757.html#scrolltoc
- *Sun Server X4-2L Service Manual* at http://docs.oracle.com/cd/E36974_01/html/E38145/gentextid-14728.html#scrolltoc
- *Sun Server X3-2 (formerly Sun Fire X4170 M3) Service Manual* at http://docs.oracle.com/cd/E22368_01/html/E27242/gentextid-14840.html#scrolltoc
- *Sun Server X3-2L (formerly Sun Fire X4270 M3) Service Manual* at http://docs.oracle.com/cd/E23393_01/html/E27229/gentextid-14804.html#scrolltoc
- *Sun Server X2-8 (formerly Sun Fire X4800 M2) Service Manual* at http://docs.oracle.com/cd/E20815_01/html/E20819/index.html
- *Sun Fire X4800 Server Service Manual* at <http://docs.oracle.com/cd/E19140-01/html/821-0282/index.html>
- *Sun Fire X4270 M2 Server Service Manual* at <http://docs.oracle.com/cd/E19245-01/E21671/index.html>
- *Sun Fire X4170 M2 Server Service Manual* at <http://docs.oracle.com/cd/E19762-01/E22369-02/index.html>
- *Sun Fire X4170, X4270, and X4275 Servers Service Manual* at <http://docs.oracle.com/cd/E19477-01/820-5830-13/index.html>
- *Sun Datacenter InfiniBand Switch 36 Firmware Version 2.1 Documentation* at http://docs.oracle.com/cd/E36265_01/index.html

- Sun Datacenter InfiniBand Switch 36 Firmware Version 2.2 Documentation at http://docs.oracle.com/cd/E76424_01/index.html
- Sun Flash Accelerator F20 PCIe Card User's Guide at <http://docs.oracle.com/cd/E19682-01/E21358/index.html>
- Sun Flash Accelerator F40 PCIe Card User's Guide at http://docs.oracle.com/cd/E29748_01/html/E29741/index.html
- Sun Flash Accelerator F80 PCIe Card User's Guide at http://docs.oracle.com/cd/E41278_01/html/E41251/index.html
- Oracle Flash Accelerator F160 PCIe Card User Guide at http://docs.oracle.com/cd/E54943_01/html/E54947/index.html
- Oracle Flash Accelerator F320 PCIe Card User Guide at http://docs.oracle.com/cd/E65386_01/html/E65387/index.html
- Oracle Flash Accelerator F640 PCIe Card User Guide at https://docs.oracle.com/cd/E87231_01/html/E87233/index.html
- Sun Storage 6 Gb SAS PCIe RAID HBA Documentation at <http://docs.oracle.com/cd/E19221-01/>
- Oracle Storage 12 Gb/s SAS PCIe RAID HBA, Internal Documentation Library at http://docs.oracle.com/cd/E52363_01/index.html
- Oracle Integrated Lights Out Manager (ILOM) Documentation at <http://www.oracle.com/goto/ilom/docs>
- "Cisco Catalyst 4948E and 4948E-F Ethernet Switches Data Sheet" at https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-4948e-ethernet-switch/data_sheet_c78-598933.html
- "Cisco Nexus 9300-EX and 9300-FX Platform Switches Data Sheet" at <https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736651.html>

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.
\$ prompt	The dollar sign (\$) prompt indicates a command run as the <code>oracle</code> user.
# prompt	The pound (#) prompt indicates a command that is run as the <code>root</code> user.

1

Site Requirements for Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack

This chapter describes the site requirements for Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack.

Note:

For ease of reading, the name "Oracle Exadata Rack" is used when information refers to both Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack.

- [General Environmental Requirements](#)
- [Space Requirements](#)

The space requirements for Oracle Exadata Database Machine depend on the hardware rack that is supplied with the system.
- [Receiving, Unpacking, and Access Route Requirements](#)

Before your Oracle Exadata Rack arrives, ensure that the receiving area is large enough for the package.
- [Maintenance Access Requirements](#)

The maintenance area must be large enough for Oracle Exadata Rack, and have the required access space.
- [Flooring Requirements](#)

Oracle Exadata Rack may be installed on raised floor or solid floor environments.
- [Electrical Power Requirements](#)

The Oracle Exadata Rack can operate effectively over a wide range of voltages and frequencies. However, each rack must have a reliable power source.
- [Temperature and Humidity Requirements](#)

Excessive internal temperatures may result in full or partial shut down of Oracle Exadata Database Machine.
- [Ventilation and Cooling Requirements](#)

Always provide adequate space in front and behind the rack to allow for proper ventilation.
- [Network Connection and IP Address Requirements for Oracle Exadata Database Machine](#)
- [Network Connection and IP Address Requirements for Oracle Exadata Storage Expansion Rack](#)

Prior to installation, network cables must be run from your existing network infrastructure to the installation site.

- [Ensuring That the Site is Ready](#)

Related Topics

- [Site Checklists](#)

1.1 General Environmental Requirements

The following sections describe the general environmental requirements for Oracle Exadata Racks.

- [General Environmental Requirements for Oracle Exadata Database Machine X6 and Later](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X5-2](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X5-8](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X4-2](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X4-8 with Exadata Storage Server X5-2 Servers](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X4-8 Full Rack](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X3-2](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X3-2 Servers](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X2-2](#)
- [General Environmental Requirements for Oracle Exadata Database Machine X2-8 Full Rack](#)
- [General Environmental Requirements for Oracle Exadata Storage Expansion Rack X6-2 and Later](#)
- [General Environmental Requirements for Oracle Exadata Storage Expansion Rack X5-2](#)
- [General Environmental Requirements for Oracle Exadata Storage Expansion Rack X4-2](#)
- [General Environmental Requirements for Oracle Exadata Storage Expansion Rack X3-2](#)
- [General Environmental Requirements for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers](#)
- [General Environmental Requirements for Single Servers](#)

1.1.1 General Environmental Requirements for Oracle Exadata Database Machine X6 and Later

For Oracle Exadata Database Machine X6 and later models, use the Oracle Exadata Configuration Assistant (OECA) to determine environmental requirements, such as

weight, acoustic level, power, cooling, and airflow. Note that OECA also supports Oracle Exadata Database Machine X8M, which uses RoCE Network Fabric.

Related Topics

- [Oracle Exadata Configuration Assistant \(OECA\) Downloads](#)

1.1.2 General Environmental Requirements for Oracle Exadata Database Machine X5-2

The environmental requirements for Oracle Exadata Database Machine X5-2 depend on the size of the system. The following table shows the general environmental requirements for Oracle Exadata Database Machine X5-2.

Table 1-1 Environmental Requirements for Oracle Exadata Database Machine X5-2

Environmental Component	Oracle Exadata Database Machine X5-2 Full Rack	Oracle Exadata Database Machine X5-2 Half Rack	Oracle Exadata Database Machine X5-2 Quarter Rack	Oracle Exadata Database Machine X5-2 Eighth Rack
Weight See Also: Flooring Requirements	<ul style="list-style-type: none"> • Extreme Flash drives: 804.9 kg (1774 lbs) • High capacity drives: 874.8 kg (1928 lbs) 	<ul style="list-style-type: none"> • Extreme Flash drives: 525.9 kg (1160 lbs) • High capacity drives: 560.9 kg (1236 lbs) 	<ul style="list-style-type: none"> • Extreme Flash drives: 366.3 kg (808 lbs) • High capacity drives: 381.2 kg (841 lbs) 	<ul style="list-style-type: none"> • Extreme Flash drives: 366.3 kg (808 lbs) • High capacity drives: 381.2 kg (841 lbs)
Acoustic levels	86 db(A)	85 db(A)	83 db(A)	83 db(A)
Power See Also: Electrical Power Requirements	<ul style="list-style-type: none"> • Extreme Flash drives Maximum: 15.7 kVA Typical: 11.0 kVA (varies based on application load) • High capacity drives Maximum: 15.6 kVA Typical: 10.9 kVA (varies based on application load) 	<ul style="list-style-type: none"> • Extreme Flash drives Maximum: 8.5 kVA Typical: 5.9 kVA (varies based on application load) • High capacity drives Maximum: 8.4 kVA Typical: 5.9 kVA (varies based on application load) 	<ul style="list-style-type: none"> • Extreme Flash drives Maximum: 4.2 kVA Typical: 3.0 kVA (varies based on application load) • High capacity drives Maximum: 4.2 kVA Typical: 2.9kVA (varies based on application load) 	<ul style="list-style-type: none"> • Extreme Flash drives Maximum: 3.5 kVA Typical: 2.4 kVA (varies based on application load) • High capacity drives Maximum: 34 kVA Typical: 2.4 kVA (varies based on application load)

Table 1-1 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X5-2

Environmental Component	Oracle Exadata Database Machine X5-2 Full Rack	Oracle Exadata Database Machine X5-2 Half Rack	Oracle Exadata Database Machine X5-2 Quarter Rack	Oracle Exadata Database Machine X5-2 Eighth Rack
<p>Cooling</p> <p>See Also: Temperature and Humidity Requirements, and Ventilation and Cooling Requirements</p>	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 52,578 BTU/hour (55,469 kJ/hour) Typical: 36,804 BTU/hour (38,829 kJ/hour) High capacity drives Maximum: 52,052 BTU/hour (54,915 kJ/hour) Typical: 36,437 BTU/hour (38,441 kJ/hour) 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 28,396 BTU/hour (29,958 kJ/hour) Typical: 19,877 BTU/hour (20,970 kJ/hour) High capacity drives Maximum: 28,133 BTU/hour (29,680 kJ/hour) Typical: 19,693 BTU/hour (20,776 kJ/hour) 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 14,195 BTU/hour (14,975 kJ/hour) Typical: 9,936 BTU/hour (10,483 kJ/hour) High capacity drives Maximum: 14,082 BTU/hour (14,856 kJ/hour) Typical: 9,857 BTU/hour (10,399 kJ/hour) 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 11,674 BTU/hour (12,317 kJ/hour) Typical: 8,172 BTU/hour (8,622 kJ/hour) High capacity drives Maximum: 11,530 BTU/hour (12,164 kJ/hour) Typical: 8,071 BTU/hour (8,515 kJ/hour)
<p>Air flow front-to-back (subject to actual data center environment)</p> <p>See Also: Temperature and Humidity Requirements, and Ventilation and Cooling Requirements</p>	<ul style="list-style-type: none"> Extreme Flash drives Maximum: Approximately 2,434 CFM Typical: Approximately 1,704 CFM High capacity drives Maximum: Approximately 2,410 CFM Typical: Approximately 1,687 CFM 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: Approximately 1,315 CFM Typical: Approximately 920 CFM High capacity drives Maximum: Approximately 1,302 CFM Typical: Approximately 912 CFM 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: Approximately 657 CFM Typical: Approximately 460 CFM High capacity drives Maximum: Approximately 652 CFM Typical: Approximately 456 CFM 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: Approximately 540 CFM Typical: Approximately 378 CFM High capacity drives Maximum: Approximately 534 CFM Typical: Approximately 374 CFM
<p>IP Addresses</p> <p>See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"</p>	<p>68 for Ethernet network, assuming single cluster</p> <p>Up to 36 for RDMA Network Fabric</p>	<p>38 for Ethernet network, assuming single cluster</p> <p>Up to 18 for RDMA Network Fabric</p>	<p>22 for Ethernet network, assuming single cluster</p> <p>Up to 8 for RDMA Network Fabric</p>	<p>22 for Ethernet network, assuming single cluster</p> <p>Up to 8 for RDMA Network Fabric</p>

Table 1-1 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X5-2

Environmental Component	Oracle Exadata Database Machine X5-2 Full Rack	Oracle Exadata Database Machine X5-2 Half Rack	Oracle Exadata Database Machine X5-2 Quarter Rack	Oracle Exadata Database Machine X5-2 Eighth Rack
Network drops See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	Minimum of 11 network drops	Minimum of 7 network drops	Minimum of 5 network drops	Minimum of 5 network drops
External connectivity See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	18 x 1 GbE/10 GbE Ethernet ports 16 x 10 GbE Ethernet ports At least 12 RDMA Network Fabric ports	12 x 1 GbE/10 GbE Ethernet ports 8 x 10 GbE Ethernet ports At least 12 RDMA Network Fabric ports	6 x 1 GbE/10 GbE Ethernet ports 4 x 10 GbE Ethernet ports At least 12 RDMA Network Fabric ports	6 x 1 GbE/10 GbE Ethernet ports 4 x 10 GbE Ethernet ports At least 12 RDMA Network Fabric ports

1.1.3 General Environmental Requirements for Oracle Exadata Database Machine X5-8

Table 1-2 shows the general environmental requirements for Oracle Exadata Database Machine X5-8. Other sections in this chapter provide detailed information.

Table 1-2 Environmental Requirements for Oracle Exadata Database Machine X5-8

Environmental Component	Oracle Exadata Database Machine X5-8 Full Rack	Oracle Exadata Database Machine X5-8 Half Rack
Weight See Also: "Flooring Requirements"	Extreme Flash drives: 1826.5 lbs (828.1 kg) High capacity drives: 1980.5 lbs (898.3 kg)	Extreme Flash drives: 1160.1 lbs (526.2 kg) High capacity drives: 1193.1 lbs (541.2 kg)
Acoustic levels	8.68	8.58
Power See Also: "Electrical Power Requirements"	Extreme Flash drives: <ul style="list-style-type: none"> Max: 16.4 kW (16.7 kVA) Typical: 11.5 kW (11.7 kVA) High capacity drives: <ul style="list-style-type: none"> Max: 16.2 kW (16.6 kVA) Typical: 11.4 kW (11.6 kVA) 	Extreme Flash drives: <ul style="list-style-type: none"> Max: 9.5 kW (9.7 kVA) Typical: 6.7 kW (6.8 kVA) High capacity drives: <ul style="list-style-type: none"> Max: 9.5 kW (9.7 kVA) Typical: 6.6 kW (6.8 kVA)

Table 1-2 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X5-8

Environmental Component	Oracle Exadata Database Machine X5-8 Full Rack	Oracle Exadata Database Machine X5-8 Half Rack
Cooling See Also: "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	Extreme Flash drives: <ul style="list-style-type: none"> Max: 55,935 BTU/hour (59,012 kJ/hour) Typical: 39,155 BTU/hour (41,308 kJ/hour) High capacity drives: <ul style="list-style-type: none"> Max: 55,410 BTU/hour (58,457 kJ/hour) Typical: 38,787 BTU/hour (40,920 kJ/hour) 	Extreme Flash drives: <ul style="list-style-type: none"> Max: 32,436 BTU/hour (34,220 kJ/hour) Typical: 22,705 BTU/hour (23,954 kJ/hour) High capacity drives: <ul style="list-style-type: none"> Max: 32,323 BTU/hour (34,101 kJ/hour) Typical: 22,626 BTU/hour (23,871 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	Extreme Flash drives: <ul style="list-style-type: none"> Max: 2590 CFM Typical: 1813 CFM High capacity drives: <ul style="list-style-type: none"> Max: 2565 CFM Typical: 1796 CFM 	Extreme Flash drives: <ul style="list-style-type: none"> Max: 1502 CFM Typical: 1051 CFM High capacity drives: <ul style="list-style-type: none"> Max: 1496 CFM Typical: 1048 CFM

1.1.4 General Environmental Requirements for Oracle Exadata Database Machine X4-2

The environmental requirements for Oracle Exadata Database Machine X4-2 depend on the size of the system. [Table 1-3](#) shows the general environmental requirements for Oracle Exadata Database Machine X4-2. The other sections in this chapter provide detailed information.

Table 1-3 Environmental Requirements for Oracle Exadata Database Machine X4-2

Environmental Component	Oracle Exadata Database Machine X4-2 Full Rack	Oracle Exadata Database Machine X4-2 Half Rack	Oracle Exadata Database Machine X4-2 Quarter Rack	Oracle Exadata Database Machine X4-2 Eighth Rack
Weight See Also: "Flooring Requirements"	852.8 kg (1880 lbs)	525.3 kg (1158 lbs)	365.6 kg (806 lbs)	365.6 kg (806 lbs)
Acoustic levels	93 db(A)	89 db(A)	85 db(A)	85 db(A)

Table 1-3 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X4-2

Environmental Component	Oracle Exadata Database Machine X4-2 Full Rack	Oracle Exadata Database Machine X4-2 Half Rack	Oracle Exadata Database Machine X4-2 Quarter Rack	Oracle Exadata Database Machine X4-2 Eighth Rack
Power See Also: "Electrical Power Requirements"	<ul style="list-style-type: none"> High performance drives Maximum: 12.4 kVA Typical: 8.7 kVA (varies based on application load) High capacity drives Maximum: 11.4 kVA Typical: 8.1 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 6.8 kVA Typical: 4.8 kVA (varies based on application load) High capacity drives Maximum: 6.3 kVA Typical: 4.4 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 3.4 kVA Typical: 2.5 kVA (varies based on application load) High capacity drives Maximum: 3.2 kVA Typical: 2.3 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 3.1 kVA Typical: 2.2 kVA (varies based on application load) High capacity drives Maximum: 2.9 kVA Typical: 2.1 kVA (varies based on application load)
Cooling See Also: "Temperature and Humidity Requirements" , and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> High performance drives Maximum: 41,300 BTU/hour (43,600 kJ/hour) Typical: 29,000 BTU/hour (30,600 kJ/hour) High capacity drives Maximum: 38,300 BTU/hour (40,400 kJ/hour) Typical: 27,000 BTU/hour (28,500 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 22,500 BTU/hour (23,750 kJ/hour) Typical: 16,000 BTU/hour (16,900 kJ/hour) High capacity drives Maximum: 21,200 BTU/hour (22,400 kJ/hour) Typical: 14,700 BTU/hour (15,500 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 11,300 BTU/hour (11,900 kJ/hour) Typical: 8,200 BTU/hour (8,700 kJ/hour) High capacity drives Maximum: 10,600 BTU/hour (11,200 kJ/hour) Typical: 7,500 BTU/hour (7,900 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 10,200 BTU/hour (10,700 kJ/hour) Typical: 7,100 BTU/hour (7,500 kJ/hour) High capacity drives Maximum: 9,500 BTU/hour (10,000 kJ/hour) Typical: 6,600 BTU/hour (7,000 kJ/hour)

Table 1-3 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X4-2

Environmental Component	Oracle Exadata Database Machine X4-2 Full Rack	Oracle Exadata Database Machine X4-2 Half Rack	Oracle Exadata Database Machine X4-2 Quarter Rack	Oracle Exadata Database Machine X4-2 Eighth Rack
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 1,900 CFM Typical: Approximately 1,350 CFM High capacity drives Maximum: Approximately 1,780 CFM Typical: Approximately 1,250 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 1,050 CFM Typical: Approximately 750 CFM High capacity drives Maximum: Approximately 980 CFM Typical: Approximately 680 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 520 CFM Typical: Approximately 365 CFM High capacity drives Maximum: Approximately 490 CFM Typical: Approximately 350 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 470 CFM Typical: Approximately 325 CFM High capacity drives Maximum: Approximately 440 CFM Typical: Approximately 310 CFM
IP Addresses See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	68 for Ethernet network, assuming single cluster Up to 36 for InfiniBand network	38 for Ethernet network, assuming single cluster Up to 18 for InfiniBand network	22 for Ethernet network, assuming single cluster Up to 8 for InfiniBand network	22 for Ethernet network, assuming single cluster Up to 8 for InfiniBand network
Network drops See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	Minimum of 11 network drops	Minimum of 7 network drops	Minimum of 5 network drops	Minimum of 5 network drops
External connectivity See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	18 x 1 GbE/10 GbE Ethernet ports 16 x 10 GbE Ethernet ports At least 12 InfiniBand ports	12 x 1 GbE/10 GbE Ethernet ports 8 x 10 GbE Ethernet ports At least 12 InfiniBand ports	6 x 1 GbE/10 GbE Ethernet ports 4 x 10 GbE Ethernet ports At least 12 InfiniBand ports	6 x 1 GbE/10 GbE Ethernet ports 4 x 10 GbE Ethernet ports At least 12 InfiniBand ports

1.1.5 General Environmental Requirements for Oracle Exadata Database Machine X4-8 with Exadata Storage Server X5-2 Servers

Table 1-4 shows the general environmental requirements for Oracle Exadata Database Machine X4-8 with Exadata Storage Server X5-2 Servers. The other sections in this chapter provide detailed information.

Table 1-4 Environmental Requirements for Oracle Exadata Database Machine X4-8 Racks with Exadata Storage Server X5-2 Servers

Environmental Component	Oracle Exadata Database Machine X4-8 Full Rack with Exadata Storage Server X5-2 Servers	Oracle Exadata Database Machine X4-8 Half Rack with Exadata Storage Server X5-2 Servers
Weight See Also: "Flooring Requirements"	<ul style="list-style-type: none"> Extreme Flash drives: 828.1 kg (1826 lbs) High capacity drives: 898.3 kg (1980 lbs) 	<ul style="list-style-type: none"> Extreme Flash drives: 526.2 kg (1160 lbs) High capacity drives: 541.2 kg (1193 lbs)
Acoustic levels	86 db(A)	85 db(A)
Power See Also: "Electrical Power Requirements"	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 17.1 kVA Typical: 12.0 kVA (varies based on application load) High capacity drives Maximum: 17.0 kVA Typical: 11.9 kVA (varies based on application load) 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 10.1 kVA Typical: 7.1 kVA (varies based on application load) High capacity drives Maximum: 10.1 kVA Typical: 7.0 kVA (varies based on application load)
Cooling See Also: "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 57,300 BTU/hour (60,452 kJ/hour) Typical: 40,110 BTU/hour (42,316 kJ/hour) High capacity drives Maximum: 56,775 BTU/hour (59,897 kJ/hour) Typical: 39,742 BTU/hour (41,928 kJ/hour) 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 33,801 BTU/hour (35,660 kJ/hour) Typical: 23,660 BTU/hour (24,962 kJ/hour) High capacity drives Maximum: 33,688 BTU/hour (35,541 kJ/hour) Typical: 23,582 BTU/hour (24,879 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> Extreme Flash drives Maximum: Approximately 2,653 CFM Typical: Approximately 1,857 CFM High capacity drives Maximum: Approximately 2,628 CFM Typical: Approximately 1,840 CFM 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: Approximately 1,565 CFM Typical: Approximately 1,095 CFM High capacity drives Maximum: Approximately 1,560 CFM Typical: Approximately 1,092 CFM

Table 1-4 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X4-8 Racks with Exadata Storage Server X5-2 Servers

Environmental Component	Oracle Exadata Database Machine X4-8 Full Rack with Exadata Storage Server X5-2 Servers	Oracle Exadata Database Machine X4-8 Half Rack with Exadata Storage Server X5-2 Servers
IP Addresses See Also: " Network Connection and IP Address Requirements for Oracle Exadata Database Machine "	44 for Ethernet network, assuming single cluster Up to 44 for InfiniBand network	22 for Ethernet network, assuming single cluster Up to 22 for InfiniBand network
Network drops See Also: " Network Connection and IP Address Requirements for Oracle Exadata Database Machine "	Minimum of 5 network drops	Minimum of 5 network drops
External connectivity See Also: " Network Connection and IP Address Requirements for Oracle Exadata Database Machine "	16 x 1 GbE Ethernet ports 16 x 10 GbE Ethernet SFP+ ports At least 12 InfiniBand ports	16 x 1 GbE Ethernet ports 16 x 10 GbE Ethernet SFP+ ports At least 12 InfiniBand ports

1.1.6 General Environmental Requirements for Oracle Exadata Database Machine X4-8 Full Rack

[Table 1-5](#) shows the general environmental requirements for Oracle Exadata Database Machine X4-8 Full Rack. The other sections in this chapter provide detailed information.

Table 1-5 Environmental Requirements for Oracle Exadata Database Machine X4-8 Full Rack

Environmental Component	Oracle Exadata Database Machine X4-8 Full Rack
Weight See Also: " Flooring Requirements "	867.3 kg (1912 lbs)
Acoustic levels	93 db(A)
Power See Also: " Electrical Power Requirements "	<ul style="list-style-type: none"> High performance drives Maximum: 15.3 kVA Typical: 10.7 kVA (varies based on application load) High capacity drives Maximum: 14.7 kVA Typical: 10.3 kVA (varies based on application load)

Table 1-5 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X4-8 Full Rack

Environmental Component	Oracle Exadata Database Machine X4-8 Full Rack
Cooling See Also: "Temperature and Humidity Requirements" , and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> • High performance drives Maximum: 51,200 BTU/hour (54,000 kJ/hour) Typical: 36,400 BTU/hour (38,400 kJ/hour) • High capacity drives Maximum: 49,110 BTU/hour (52,000 kJ/hour) Typical: 34,400 BTU/hour (36,300 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements" , and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> • High performance drives Maximum: Approximately 2,410 CFM Typical: Approximately 1,700 CFM • High capacity drives Maximum: Approximately 2,280 CFM Typical: Approximately 1,600 CFM
IP Addresses See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	44 for Ethernet network, assuming single cluster Up to 36 for InfiniBand network
Network drops See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	Minimum of 5 network drops
External connectivity See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	16 x 1 GbE Ethernet ports 16 x 10 GbE Ethernet SFP+ ports 12 InfiniBand ports

1.1.7 General Environmental Requirements for Oracle Exadata Database Machine X3-2

The environmental requirements for Oracle Exadata Database Machine X3-2 depend on the size of the system. [Table 1-6](#) shows the general environmental requirements for Oracle Exadata Database Machine X3-2. The other sections in this chapter provide detailed information.

Table 1-6 Environmental Requirements for Oracle Exadata Database Machine X3-2

Environmental Component	Oracle Exadata Database Machine X3-2 Full Rack	Oracle Exadata Database Machine X3-2 Half Rack	Oracle Exadata Database Machine X3-2 Quarter Rack	Oracle Exadata Database Machine Eighth Rack
Weight See Also: "Flooring Requirements"	871.4 kg (1921 lbs)	543.9 kg (1199 lbs)	368.8 kg (813 lbs)	368.8 kg (813 lbs)
Acoustic levels	87 db(A)	84 db(A)	81 db(A)	81 db(A)
Power See Also: "Electrical Power Requirements"	<ul style="list-style-type: none"> High performance drives Maximum: 12.1 kVA Typical: 8.6 kVA (varies based on application load) High capacity drives Maximum: 11.1 kVA Typical: 7.8 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 6.6 kVA Typical: 4.7 kVA (varies based on application load) High capacity drives Maximum: 6.1 kVA Typical: 4.3 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 3.3 kVA Typical: 2.4 kVA (varies based on application load) High capacity drives Maximum: 3.1 kVA Typical: 2.2 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 3.0 kVA Typical: 2.1 kVA (varies based on application load) High capacity drives Maximum: 2.8 kVA Typical: 2.0 kVA (varies based on application load)
Cooling See Also: "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> High performance drives Maximum: 40,600 BTU/hour (42,800 kJ/hour) Typical: 28,700 BTU/hour (30,300 kJ/hour) High capacity drives Maximum: 37,200 BTU/hour (39,250 kJ/hour) Typical: 26,000 BTU/hour (27,400 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 22,200 BTU/hour (23,400 kJ/hour) Typical: 15,700 BTU/hour (16,600 kJ/hour) High capacity drives Maximum: 20,500 BTU/hour (21,600 kJ/hour) Typical: 14,300 BTU/hour (15,100 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 10,900 BTU/hour (11,500 kJ/hour) Typical: 7,850 BTU/hour (8,300 kJ/hour) High capacity drives Maximum: 10,200 BTU/hour (10,800 kJ/hour) Typical: 7,200 BTU/hour (7,600 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 9,900 BTU/hour (10,500 kJ/hour) Typical: 6,800 BTU/hour (7,200 kJ/hour) High capacity drives Maximum: 9,200 BTU/hour (9,700 kJ/hour) Typical: 6,500 BTU/hour (6,850 kJ/hour)

Table 1-6 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X3-2

Environmental Component	Oracle Exadata Database Machine X3-2 Full Rack	Oracle Exadata Database Machine X3-2 Half Rack	Oracle Exadata Database Machine X3-2 Quarter Rack	Oracle Exadata Database Machine Eighth Rack
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 1,900 CFM Typical: Approximately 1,350 CFM High capacity drives Maximum: Approximately 1,700 CFM Typical: Approximately 1,200 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 1,050 CFM Typical: Approximately 750 CFM High capacity drives Maximum: Approximately 950 CFM Typical: Approximately 670 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 500 CFM Typical: Approximately 375 CFM High capacity drives Maximum: Approximately 470 CFM Typical: Approximately 330 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 460 CFM Typical: Approximately 325 CFM High capacity drives Maximum: Approximately 425 CFM Typical: Approximately 300 CFM
IP Addresses See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	69 for Ethernet network, assuming single cluster 22 for InfiniBand network	39 for Ethernet network, assuming single cluster (39 IP addresses are needed for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers)) 11 for InfiniBand network	22 for Ethernet network, assuming single cluster 5 for InfiniBand network	22 for Ethernet network, assuming single cluster 5 for InfiniBand network
Network drops See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	Minimum of 11 network drops	Minimum of 7 network drops	Minimum of 5 network drops	Minimum of 5 network drops

Table 1-6 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X3-2

Environmental Component	Oracle Exadata Database Machine X3-2 Full Rack	Oracle Exadata Database Machine X3-2 Half Rack	Oracle Exadata Database Machine X3-2 Quarter Rack	Oracle Exadata Database Machine Eighth Rack
External connectivity See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	24 x 1 GbE/10 GbE Ethernet ports 16 x 10 GbE Ethernet ports At least 12 InfiniBand ports	12 x 1 GbE/10 GbE Ethernet ports 8 x 10 GbE Ethernet ports At least 12 InfiniBand ports	6 x 1 GbE/10 GbE Ethernet ports 4 x 10 GbE Ethernet ports At least 12 InfiniBand ports	6 x 1 GbE/10 GbE Ethernet ports 4 x 10 GbE Ethernet ports At least 12 InfiniBand ports

1.1.8 General Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers

[Table 1-7](#) shows the general environmental requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers. The other sections in this chapter provide detailed information.

Table 1-7 Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers

Environmental Component	Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers
Weight See Also: "Flooring Requirements"	911.27 kg (2009 lbs)
Acoustic levels	84 db(A)
Power See Also: "Electrical Power Requirements"	<ul style="list-style-type: none"> • High performance drives Maximum: 15.3 kVA Typical: 10.7 kVA (varies based on application load) • High capacity drives Maximum: 14.5 kVA Typical: 10.2 kVA (varies based on application load)

Table 1-7 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers

Environmental Component	Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers
Cooling See Also: "Temperature and Humidity Requirements" , and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> • High performance drives Maximum: 51,200 BTU/hour (54,000 kJ/hour) Typical: 35,900 BTU/hour (37,900 kJ/hour) • High capacity drives Maximum: 48,500 BTU/hour (51,200 kJ/hour) Typical: 34,200 BTU/hour (36,100 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements" , and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> • High performance drives Maximum: Approximately 2,400 CFM Typical: Approximately 1,700 CFM • High capacity drives Maximum: Approximately 2,250 CFM Typical: Approximately 1,600 CFM
IP Addresses See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	45 for Ethernet network, assuming single cluster Up to 36 for InfiniBand network
Network drops See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	Minimum of 5 network drops
External connectivity See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	16 x 1 GbE Ethernet ports 16 x 10 GbE Ethernet SFP+ ports 12 InfiniBand ports

1.1.9 General Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X3-2 Servers

[Table 1-8](#) shows the general environmental requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X3-2 Servers. The other sections in this chapter provide detailed information.

Table 1-8 Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X3-2 Servers

Environmental Component	Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X3-2 Servers
Weight See Also: "Flooring Requirements"	929.9 kg (2050 lbs)
Acoustic levels	84 db(A)
Power See Also: "Electrical Power Requirements"	<ul style="list-style-type: none"> • High performance drives Maximum: 15.3 kVA Typical: 10.7 kVA (varies based on application load) • High capacity drives Maximum: 14.3 kVA Typical: 10.0 kVA (varies based on application load)
Cooling See Also: "Temperature and Humidity Requirements" , and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> • High performance drives Maximum: 51,200 BTU/hour (54,000 kJ/hour) Typical: 35,900 BTU/hour (37,900 kJ/hour) • High capacity drives Maximum: 47,800 BTU/hour (50,400 kJ/hour) Typical: 33,500 BTU/hour (35,300 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements" , and "Ventilation and Cooling Requirements"	<ul style="list-style-type: none"> • High performance drives Maximum: Approximately 2,400 CFM Typical: Approximately 1,700 CFM • High capacity drives Maximum: Approximately 2,200 CFM Typical: Approximately 1,550 CFM
IP Addresses See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	45 for Ethernet network, assuming single cluster 22 for InfiniBand network
Network drops See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	Minimum of 5 network drops
External connectivity See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	16 x 1 GbE Ethernet ports 16 x 10 GbE Ethernet SFP+ ports 12 InfiniBand ports

1.1.10 General Environmental Requirements for Oracle Exadata Database Machine X2-2

The environmental requirements for Oracle Exadata Database Machine X2-2 depend on the size of the system. [Table 1-9](#) shows the general environmental requirements for Oracle Exadata Database Machine X2-2. The other sections in this chapter provide detailed information.

Table 1-9 Environmental Requirements for Oracle Exadata Database Machine X2-2

Environmental Component	Oracle Exadata Database Machine X2-2 Full Rack	Oracle Exadata Database Machine X2-2 Half Rack	Oracle Exadata Database Machine X2-2 Quarter Rack
Weight See Also: " Flooring Requirements "	966.6 kg (2131 lbs)	602.8 kg (1329 lbs)	409.1kg (902 lbs)
Acoustic levels	89 db(A)	86 db(A)	83 db(A)
Power See Also: " Electrical Power Requirements "	Maximum: 14.0 kW (14.3 kVA) Typical: 9.8 kW (10.0 kVA) (varies based on application load)	Maximum: 7.2 kW (7.3 kVA) Typical: 5.1 kW (5.2 kVA) (varies based on application load)	Maximum: 3.6 kW (3.7 kVA) Typical: 2.7 kW (2.75kVA) (varies based on application load)
Cooling See Also: " Temperature and Humidity Requirements ", and " Ventilation and Cooling Requirements "	Maximum: 47,800 BTU/hour (50,400 kJ/hour) Typical: 33,400 BTU/hour (35,300 kJ/hour)	Maximum: 26,400 BTU/hour (25,950 kJ/hour) Typical: 17,400 BTU/hour (35,300 kJ/hour)	Maximum: 12,300 BTU/hour (13,000 kJ/hour) Typical: 9,200 BTU/hour (9,700 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: " Temperature and Humidity Requirements ", and " Ventilation and Cooling Requirements "	Maximum: Approximately 2,200 CFM Typical: Approximately 1,560 CFM	Maximum: Approximately 1,130 CFM Typical: Approximately 840 CFM	Maximum: Approximately 550 CFM Typical: Approximately 410 CFM
IP Addresses See Also: " Network Connection and IP Address Requirements for Oracle Exadata Database Machine "	70 for Ethernet network, assuming single cluster 22 for InfiniBand network	40 for Ethernet network, assuming single cluster (39 IP addresses are needed for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers)) 11 for InfiniBand network	23 for Ethernet network, assuming single cluster 5 for InfiniBand network

Table 1-9 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X2-2

Environmental Component	Oracle Exadata Database Machine X2-2 Full Rack	Oracle Exadata Database Machine X2-2 Half Rack	Oracle Exadata Database Machine X2-2 Quarter Rack
Network drops See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	Minimum of 12 network drops	Minimum of 8 network drops	Minimum of 6 network drops
External connectivity See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	24 x 1 GbE Ethernet ports 16 x 10 GbE Ethernet ports (valid for M2 servers only) At least 12 InfiniBand ports	12 x 1 GbE Ethernet ports 8 x 10 GbE Ethernet ports (valid for M2 servers only) At least 12 InfiniBand ports	6 x 1 GbE Ethernet ports 4 x 10 GbE Ethernet ports (valid for M2 servers only) At least 12 InfiniBand ports

1.1.11 General Environmental Requirements for Oracle Exadata Database Machine X2-8 Full Rack

Table 1-10 shows the general environmental requirements for Oracle Exadata Database Machine X2-8 Full Rack. The other sections in this chapter provide detailed information.

Table 1-10 Environmental Requirements for Oracle Exadata Database Machine X2-8 Full Rack

Environmental Component	Oracle Exadata Database Machine X2-8 Full Rack (Sun Fire X4800 Oracle Database Servers)	Oracle Exadata Database Machine X2-8 Full Rack (Sun Server X2-8 Oracle Database Servers)
Weight See Also: "Flooring Requirements"	943.5 kg (2080 lbs)	980.7 kg (2162 lbs)
Acoustic levels	85 db(A)	85 db(A)
Power See Also: "Electrical Power Requirements"	Maximum: 17.0 kW (17.4 kVA) Typical: 11.9 kW (12.2 kVA) (varies based on application load)	Maximum: 17.7 kW (18.1 kVA) Typical: 12.4 kW (112.7 kVA) (varies based on application load)
Cooling See Also: "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	Maximum: 58,050 BTU/hour (61,200 kJ/hour) Typical: 40,630 BTU/hour (42,840 kJ/hour)	Maximum: 60,350 BTU/hour (63,630 kJ/hour) Typical: 42,280 BTU/hour (44,580 kJ/hour)

Table 1-10 (Cont.) Environmental Requirements for Oracle Exadata Database Machine X2-8 Full Rack

Environmental Component	Oracle Exadata Database Machine X2-8 Full Rack (Sun Fire X4800 Oracle Database Servers)	Oracle Exadata Database Machine X2-8 Full Rack (Sun Server X2-8 Oracle Database Servers)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	Maximum: Approximately 2,690 CFM Typical: Approximately 1,880 CFM	Maximum: Approximately 2,781CFM Typical: Approximately 1,950 CFM
IP Addresses See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	45 for Ethernet network, assuming single cluster 22 for InfiniBand network	45 for Ethernet network, assuming single cluster 22 for InfiniBand network
Network drops See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	Minimum of 5 network drops	Minimum of 5 network drops
External connectivity See Also: "Network Connection and IP Address Requirements for Oracle Exadata Database Machine"	16 x 1 GbE Ethernet ports 16 x 10 GbE Ethernet SFP+ ports 12 InfiniBand ports	16 x 1 GbE Ethernet ports 16 x 10 GbE Ethernet SFP+ ports 12 InfiniBand ports

1.1.12 General Environmental Requirements for Oracle Exadata Storage Expansion Rack X6-2 and Later

For Oracle Exadata Storage Expansion Rack X6-2 and later models, use the Oracle Exadata Configuration Assistant (OECA) to determine environmental requirements, such as weight, acoustic level, power, cooling, and airflow. Note that OECA also supports Oracle Exadata Storage Expansion Rack X8M-2, which uses RoCE Network Fabric.

Related Topics

- [Oracle Exadata Configuration Assistant \(OECA\) Downloads](#)

1.1.13 General Environmental Requirements for Oracle Exadata Storage Expansion Rack X5-2

The environmental requirements for Oracle Exadata Storage Expansion Rack X5-2 depend on the size of the system. [Table 1-11](#) shows the general environmental requirements for each type of Oracle Exadata Storage Expansion Rack X5-2.

Table 1-11 Environmental Requirements for Oracle Exadata Storage Expansion Rack X5-2

Environmental Component	Oracle Exadata Storage Expansion X5-2 Full Rack	Oracle Exadata Storage Expansion X5-2 Half Rack	Oracle Exadata Storage Expansion X5-2 Quarter Rack
Weight	<ul style="list-style-type: none"> Extreme Flash drives: 821 kg (1810 lbs) High capacity drives: 915.8 kg (2019 lbs) 	<ul style="list-style-type: none"> Extreme Flash drives: 546.6 kg (1205 lbs) High capacity drives: 591.5 kg (1304 lbs) 	<ul style="list-style-type: none"> Extreme Flash drives: 364.5 kg (804 lbs) High capacity drives: 384.4 kg (848 lbs)
Acoustic levels	82 db(A)	79 db(A)	76 db(A)
Power	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 13.1 kVA Typical: 9.2 kVA (varies based on application load) High capacity drives Maximum: 12.8 kVA Typical: 8.9 kVA (varies based on application load) 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 5.9 kVA Typical: 4.2 kVA (varies based on application load) High capacity drives Maximum: 6.9 kVA Typical: 4.8 kVA (varies based on application load) 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 3.7 kVA Typical: 2.6 kVA (varies based on application load) High capacity drives Maximum: 3.7 kVA Typical: 2.6 kVA (varies based on application load)
Cooling	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 43,765 BTU/hour (46,170 kJ/hour) Typical: 30,635 BTU/hour (32,320 kJ/hour) High capacity drives Maximum: 42,670 BTU/hour (45,020 kJ/hour) Typical: 29,870 BTU/hour (31,515 kJ/hour) 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 19,880 BTU/hour (20,970 kJ/hour) Typical: 13,915 BTU/hour (14,680 kJ/hour) High capacity drives Maximum: 22,950 BTU/hour (24,210 kJ/hour) Typical: 16,065 BTU/hour (16,950 kJ/hour) 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: 12,362 BTU/hour (13,042 kJ/hour) Typical: 8,654 BTU/hour (9,129 kJ/hour) High capacity drives Maximum: 12,212 BTU/hour (12,884 kJ/hour) Typical: 8,548 BTU/hour (9,019 kJ/hour)

Table 1-11 (Cont.) Environmental Requirements for Oracle Exadata Storage Expansion Rack X5-2

Environmental Component	Oracle Exadata Storage Expansion X5-2 Full Rack	Oracle Exadata Storage Expansion X5-2 Half Rack	Oracle Exadata Storage Expansion X5-2 Quarter Rack
Air flow front-to-back (subject to actual data center environment)	<ul style="list-style-type: none"> Extreme Flash drives Maximum: Approximately 2,030 CFM Typical: Approximately 1,420 CFM High capacity drives Maximum: Approximately 1,975 CFM Typical: Approximately 1,385 CFM 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: Approximately 920 CFM Typical: Approximately 645 CFM High capacity drives Maximum: Approximately 1,065 CFM Typical: Approximately 745 CFM 	<ul style="list-style-type: none"> Extreme Flash drives Maximum: Approximately 565 CFM Typical: Approximately 396 CFM High capacity drives Maximum: Approximately 572 CFM Typical: Approximately 401 CFM
IP Addresses	44 for Ethernet network, assuming single cluster 38 for InfiniBand network	24 for Ethernet network, assuming single cluster 18 for InfiniBand network	13 for Ethernet network, assuming single cluster 8 for InfiniBand network
Network drops	Minimum of 1 network drop	Minimum of 1 network drop	Minimum of 1 network drop
External connectivity	12 InfiniBand ports	12 InfiniBand ports	12 InfiniBand ports

1.1.14 General Environmental Requirements for Oracle Exadata Storage Expansion Rack X4-2

The environmental requirements for Oracle Exadata Storage Expansion Rack X4-2 depend on the size of the system. [Table 1-12](#) shows the general environmental requirements for each type of Oracle Exadata Storage Expansion Rack X4-2.

Table 1-12 Environmental Requirements for Oracle Exadata Storage Expansion Rack X4-2

Environmental Component	Oracle Exadata Storage Expansion X4-2 Full Rack	Oracle Exadata Storage Expansion X4-2 Half Rack	Oracle Exadata Storage Expansion X4-2 Quarter Rack
Weight	824.2 kg (1817 lbs)	513.0 kg (1131 lbs)	346.5 kg (764 lbs)
Acoustic levels	82 db(A)	79 db(A)	76 db(A)

Table 1-12 (Cont.) Environmental Requirements for Oracle Exadata Storage Expansion Rack X4-2

Environmental Component	Oracle Exadata Storage Expansion X4-2 Full Rack	Oracle Exadata Storage Expansion X4-2 Half Rack	Oracle Exadata Storage Expansion X4-2 Quarter Rack
Power	<ul style="list-style-type: none"> High performance drives Maximum: 10.2 kVA Typical: 7.1 kVA (varies based on application load) High capacity drives Maximum: 9.1 kVA Typical: 6.3 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 5.7 kVA Typical: 4.1 kVA (varies based on application load) High capacity drives Maximum: 5.2 kVA Typical: 3.7 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 2.9 kVA Typical: 2.0 kVA (varies based on application load) High capacity drives Maximum: 2.7 kVA Typical: 2.0 kVA (varies based on application load)
Cooling	<ul style="list-style-type: none"> High performance drives Maximum: 34,100 BTU/hour (36,000 kJ/hour) Typical: 23,900 BTU/hour (25,200 kJ/hour) High capacity drives Maximum: 30,400 BTU/hour (32,100 kJ/hour) Typical: 21,200 BTU/hour (22,400 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 19,100 BTU/hour (20,150 kJ/hour) Typical: 13,650 BTU/hour (14,400 kJ/hour) High capacity drives Maximum: 17,400 BTU/hour (18,400 kJ/hour) Typical: 12,300 BTU/hour (13,000 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 9,500 BTU/hour (10,000 kJ/hour) Typical: 6,700 BTU/hour (7,000 kJ/hour) High capacity drives Maximum: 8,900 BTU/hour (9,400 kJ/hour) Typical: 6,500 BTU/hour (6,900 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 1,600 CFM Typical: Approximately 1,100 CFM High capacity drives Maximum: Approximately 1,410 CFM Typical: Approximately 980 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 900 CFM Typical: Approximately 650 CFM High capacity drives Maximum: Approximately 810 CFM Typical: Approximately 570 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 440 CFM Typical: Approximately 310 CFM High capacity drives Maximum: Approximately 410 CFM Typical: Approximately 300 CFM

Table 1-12 (Cont.) Environmental Requirements for Oracle Exadata Storage Expansion Rack X4-2

Environmental Component	Oracle Exadata Storage Expansion X4-2 Full Rack	Oracle Exadata Storage Expansion X4-2 Half Rack	Oracle Exadata Storage Expansion X4-2 Quarter Rack
IP Addresses	42 for Ethernet network, assuming single cluster 36 for InfiniBand network	24 for Ethernet network, assuming single cluster 18 for InfiniBand network	13 for Ethernet network, assuming single cluster 8 for InfiniBand network
Network drops	Minimum of 1 network drop	Minimum of 1 network drop	Minimum of 1 network drop
External connectivity	12 InfiniBand ports	12 InfiniBand ports	12 InfiniBand ports

1.1.15 General Environmental Requirements for Oracle Exadata Storage Expansion Rack X3-2

The environmental requirements for Oracle Exadata Storage Expansion Rack X3-2 depend on the size of the system. [Table 1-13](#) shows the general environmental requirements for each type of Oracle Exadata Storage Expansion Rack X3-2.

Table 1-13 Environmental Requirements for Oracle Exadata Storage Expansion Rack X3-2

Environmental Component	Oracle Exadata Storage Expansion X3-2 Full Rack	Oracle Exadata Storage Expansion X3-2 Half Rack	Oracle Exadata Storage Expansion X3-2 Quarter Rack
Weight	827.4 kg (1824 lbs)	516.2 kg (1138 lbs)	349.7 kg (771 lbs)
Acoustic levels	82 db(A)	79 db(A)	76 db(A)
Power	<ul style="list-style-type: none"> High performance drives Maximum: 10.2 kVA Typical: 7.1 kVA (varies based on application load) High capacity drives Maximum: 8.9 kVA Typical: 6.2 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 5.7 kVA Typical: 4.1 kVA (varies based on application load) High capacity drives Maximum: 5.1 kVA Typical: 3.6 kVA (varies based on application load) 	<ul style="list-style-type: none"> High performance drives Maximum: 2.9 kVA Typical: 2.0 kVA (varies based on application load) High capacity drives Maximum: 2.6 kVA Typical: 1.9 kVA (varies based on application load)

Table 1-13 (Cont.) Environmental Requirements for Oracle Exadata Storage Expansion Rack X3-2

Environmental Component	Oracle Exadata Storage Expansion X3-2 Full Rack	Oracle Exadata Storage Expansion X3-2 Half Rack	Oracle Exadata Storage Expansion X3-2 Quarter Rack
Cooling	<ul style="list-style-type: none"> High performance drives Maximum: 34,100 BTU/hour (36,000 kJ/hour) Typical: 23,900 BTU/hour (25,200 kJ/hour) High capacity drives Maximum: 29,700 BTU/hour (31,300 kJ/hour) Typical: 20,800 BTU/hour (21,950 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 19,100 BTU/hour (20,150 kJ/hour) Typical: 13,650 BTU/hour (14,400 kJ/hour) High capacity drives Maximum: 17,100 BTU/hour (18,000 kJ/hour) Typical: 11,950 BTU/hour (12,600 kJ/hour) 	<ul style="list-style-type: none"> High performance drives Maximum: 9,500 BTU/hour (10,000 kJ/hour) Typical: 6,700 BTU/hour (7,000 kJ/hour) High capacity drives Maximum: 8,500 BTU/hour (9,000 kJ/hour) Typical: 6,150 BTU/hour (6,500 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 1,600 CFM Typical: Approximately 1,100 CFM High capacity drives Maximum: Approximately 1,375 CFM Typical: Approximately 960 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 900 CFM Typical: Approximately 650 CFM High capacity drives Maximum: Approximately 790 CFM Typical: Approximately 550 CFM 	<ul style="list-style-type: none"> High performance drives Maximum: Approximately 440 CFM Typical: Approximately 310 CFM High capacity drives Maximum: Approximately 400 CFM Typical: Approximately 285 CFM
IP Addresses	42 for Ethernet network, assuming single cluster 18 for InfiniBand network	24 for Ethernet network, assuming single cluster 9 for InfiniBand network	13 for Ethernet network, assuming single cluster 4 for InfiniBand network
Network drops	Minimum of 1 network drop	Minimum of 1 network drop	Minimum of 1 network drop
External connectivity	12 InfiniBand ports	12 InfiniBand ports	12 InfiniBand ports

1.1.16 General Environmental Requirements for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers

The environmental requirements for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers depend on the size of the system. [Table 1-14](#) shows the general environmental requirements for each type of Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers.

Table 1-14 Environmental Requirements for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers

Environmental Component	Oracle Exadata Storage Expansion Full Rack	Oracle Exadata Storage Expansion Half Rack	Oracle Exadata Storage Expansion Quarter Rack
Weight	917.6 kg (2023 lbs)	578.3 kg (1275 lbs)	396.8 kg (875 lbs)
Acoustic levels	83 db(A)	82 db(A)	81 db(A)
Power	Maximum: 12.6 kW (12.9 kVA) Typical: 8.8 kW (9.0 kVA) (varies based on application load)	Maximum: 6.9 kW (7.1 kVA) Typical: 4.8 kW (5.0 kVA) (varies based on application load)	Maximum: 3.4 kW (3.5 kVA) Typical: 2.4 kW (2.5kVA) (varies based on application load)
Cooling	Maximum: 43,000 BTU/hour (45,400 kJ/hour) Typical: 30,100 BTU/hour (31,800 kJ/hour)	Maximum: 23,600 BTU/hour (24,900 kJ/hour) Typical: 16,500 BTU/hour (17,400 kJ/hour)	Maximum: 11,600 BTU/hour (12,250 kJ/hour) Typical: 8,100 BTU/hour (8,600 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	Maximum: Approximately 1,980 CFM Typical: Approximately 1,390 CFM	Maximum: Approximately 1,090 CFM Typical: Approximately 760 CFM	Maximum: Approximately 530 CFM Typical: Approximately 375 CFM
IP Addresses	43 for Ethernet network, assuming single cluster 18 for InfiniBand network	25 for Ethernet network, assuming single cluster 9 for InfiniBand network	14 for Ethernet network, assuming single cluster 4 for InfiniBand network
Network drops	Minimum of 2 network drops	Minimum of 2 network drops	Minimum of 2 network drops
External connectivity	12 InfiniBand ports	12 InfiniBand ports	12 InfiniBand ports

1.1.17 General Environmental Requirements for Single Servers

The environmental requirements for single servers are shown in the following tables.

Table 1-15 Environmental Requirements for X8M-2 Single Servers

Environmental Component	Database Server X8M-2	Exadata Storage Server X8M-2 - High Capacity	Exadata Storage Server X8M-2 Eighth Rack - High Capacity	Exadata Storage Server X8M-2 - Extreme Flash	Exadata Storage Server X8M-2 XT Server
Weight	20.7 kg (45.6 lbs)	34.8 kg (76.7 lbs)	30.6 kg (67.5 lbs)	27.5 kg (60.6 lbs)	30.2 kg (66.7 lbs)
Acoustic levels	77 db	82 dB	82 dB	82 dB	82 dB
Power	Maximum: 0.7 kW (0.7 kVA) Typical: 0.5 kW (0.5 kVA) (varies based on application load)	Maximum: 0.8 kW (0.8 kVA) Typical: 0.5 kW (0.6 kVA) (varies based on application load)	Maximum: 0.6 kW (0.6 kVA) Typical: 0.4 kW (0.4 kVA) (varies based on application load)	Maximum: 0.8 kW (0.8 kVA) Typical: 0.6 kW (0.6 kVA) (varies based on application load)	Maximum: 0.5 kW (0.5 kVA) Typical: 0.3 kW (0.3 kVA) (varies based on application load)
Cooling	At max usage: 2,409 BTU/hour (2,541 kJ/hour) At typical usage: 1,686 BTU/hour (1,779 kJ/hour)	At max usage: 2,631 BTU/hour (2,775 kJ/hour) At typical usage: 1,842 BTU/hour (1,943 kJ/hour)	At max usage: 1,947 BTU/hour (2,054 kJ/hour) At typical usage: 1,363 BTU/hour (1,438 kJ/hour)	At max usage: 2,730 BTU/hour (2,880 kJ/hour) At typical usage: 1,911 BTU/hour (2,016 kJ/hour)	At max usage: 1570 BTU/hour (1,656 kJ/hour) At typical usage: 1,099 BTU/hour (1,159 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	At max usage: Approximately 112 CFM At typical usage: Approximately 78 CFM	At max usage: Approximately 122 CFM At typical usage: Approximately 85 CFM	At max usage: Approximately 90 CFM At typical usage: Approximately 63 CFM	At max usage: Approximately 126 CFM At typical usage: Approximately 88 CFM	At max usage: Approximately 73 CFM At typical usage: Approximately 51 CFM

Table 1-16 Environmental Requirements for X8-2 Single Servers

Environmental Component	Database Server X8-2	Exadata Storage Server X8-2 - High Capacity	Exadata Storage Server X8-2 Eighth Rack - High Capacity	Exadata Storage Server X8-2 - Extreme Flash	Exadata Storage Server X8-2 XT Server
Weight	20.7 kg (45.6 lbs)	34.8 kg (76.7 lbs)	30.6 kg (67.5 lbs)	27.5 kg (60.6 lbs)	30.2 kg (66.7 lbs)
Acoustic levels	77 db	82 dB	82 dB	82 dB	82 dB

Table 1-16 (Cont.) Environmental Requirements for X8-2 Single Servers

Environmental Component	Database Server X8-2	Exadata Storage Server X8-2 - High Capacity	Exadata Storage Server X8-2 Eighth Rack - High Capacity	Exadata Storage Server X8-2 - Extreme Flash	Exadata Storage Server X8-2 XT Server
Power	Maximum: 0.7 kW (0.7 kVA) Typical: 0.5 kW (0.5 kVA) (varies based on application load)	Maximum: 0.7 kW (0.7 kVA) Typical: 0.5 kW (0.5 kVA) (varies based on application load)	Maximum: 0.5 kW (0.5 kVA) Typical: 0.4 kW (0.4 kVA) (varies based on application load)	Maximum: 0.7 kW (0.7 kVA) Typical: 0.5 kW (0.5 kVA) (varies based on application load)	Maximum: 0.5 kW (0.5 kVA) Typical: 0.3 kW (0.3 kVA) (varies based on application load)
Cooling	At max usage: 2,409 BTU/hour (2,541 kJ/hour) At typical usage: 1,686 BTU/hour (1,779 kJ/hour)	At max usage: 2,317 BTU/hour (2,444 kJ/hour) At typical usage: 1,622 BTU/hour (1,711 kJ/hour)	At max usage: 1,714 BTU/hour (1,809 kJ/hour) At typical usage: 1,200 BTU/hour (1,266 kJ/hour)	At max usage: 2,283 BTU/hour (2,408 kJ/hour) At typical usage: 1,598 BTU/hour (1,686 kJ/hour)	At max usage: 1570 BTU/hour (1,656 kJ/hour) At typical usage: 1,099 BTU/hour (1,159 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	At max usage: Approximately 112 CFM At typical usage: Approximately 78 CFM	At max usage: Approximately 107 CFM At typical usage: Approximately 75 CFM	At max usage: Approximately 79 CFM At typical usage: Approximately 56 CFM	At max usage: Approximately 106 CFM At typical usage: Approximately 74 CFM	At max usage: Approximately 73 CFM At typical usage: Approximately 51 CFM

Table 1-17 Environmental Requirements for X7-2 Single Servers

Environmental Component	Exadata Storage Server X7-2 - Extreme Flash	Exadata Storage Server X7-2 Eighth Rack - High Capacity	Exadata Storage Server X7-2 - High Capacity	Database Server X7-2
Weight	27.5 kg (61 lbs)	30.1 kg (66.3 lbs)	34.4 kg (76 lbs)	20.7 kg (46 lbs)
Acoustic levels	82 db(A)	79 db(A)	82 db(A)	77 db(A)
Power	Maximum: 0.6 kW (0.6 kVA) Typical: 0.4 kW (0.5 kVA) (varies based on application load)	Maximum: 0.5 kW (0.5 kVA) Typical: 0.3 kW (0.3 kVA) (varies based on application load)	Maximum: 0.6 kW (0.6 kVA) Typical: 0.4 kW (0.4 kVA) (varies based on application load)	Maximum: 0.8 kW (0.8 kVA) Typical: 0.5 kW (0.6 kVA) (varies based on application load)

Table 1-17 (Cont.) Environmental Requirements for X7-2 Single Servers

Environmental Component	Exadata Storage Server X7-2 - Extreme Flash	Exadata Storage Server X7-2 Eighth Rack - High Capacity	Exadata Storage Server X7-2 - High Capacity	Database Server X7-2
Cooling	Maximum: 2,167 BTU/hour (2,286 kJ/hour) Typical: 1,600 BTU/hour (1,378 kJ/hour)	Maximum: 1,588 BTU/hour (1,676 kJ/hour) Typical: 1,112 BTU/hour (1,173 kJ/hour)	Maximum: 2,146 BTU/hour (2,264 kJ/hour) Typical: 1,502 BTU/hour (1,585 kJ/hour)	Maximum: 2,631 BTU/hour (2,775 kJ/hour) Typical: 1,842 BTU/hour (1,943 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	Maximum: Approximately 100 CFM Typical: Approximately 70 CFM	Maximum: Approximately 74 CFM Typical: Approximately 51 CFM	Maximum: Approximately 99 CFM Typical: Approximately 70 CFM	Maximum: Approximately 122 CFM Typical: Approximately 85 CFM

Table 1-18 Environmental Requirements for X6-2 Single Servers

Environmental Component	Exadata Storage Server X6-2 - Extreme Flash	Exadata Storage Server X6-2 - High Capacity	Database Server X6-2
Weight	28.1 kg (62 lbs)	33.1 kg (73 lbs)	20.4 kg (45 lbs)
Acoustic levels	78 db(A)	78 db(A)	78 db(A)
Power	Maximum: 0.6 kVA Typical: 0.4 kVA (varies based on application load)	Maximum: 0.6 kVA Typical: 0.4 kVA (varies based on application load)	Maximum: 0.665 kW (0.678 kVA) Typical: 0.466 kW (0.475 kVA) (varies based on application load)
Cooling	Maximum: 1,866 BTU/hour (1,969 kJ/hour) Typical: 1,307 BTU/hour (1,378 kJ/hour)	Maximum: 2,006 BTU/hour (2,117 kJ/hour) Typical: 1,404 BTU/hour (1,482 kJ/hour)	Maximum: 2,269 BTU/hour (2,394 kJ/hour) Typical: 1,588 BTU/hour (1,676 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	Maximum: Approximately 86 CFM Typical: Approximately 60 CFM	Maximum: Approximately 93 CFM Typical: Approximately 65 CFM	Maximum: Approximately 105 CFM Typical: Approximately 74 CFM

Table 1-19 Environmental Requirements for X5-2 Single Servers

Environmental Component	Exadata Storage Server X5-2 - Extreme Flash	Exadata Storage Server X5-2 - High Capacity	Database Server X5-2
Weight	28.1 kg (62 lbs)	33.1 kg (73 lbs)	20.4 kg (45 lbs)
Acoustic levels	78 db(A)	78 db(A)	78 db(A)

Table 1-19 (Cont.) Environmental Requirements for X5-2 Single Servers

Environmental Component	Exadata Storage Server X5-2 - Extreme Flash	Exadata Storage Server X5-2 - High Capacity	Database Server X5-2
Power	Maximum: 0.6 kVA Typical: 0.4 kVA (varies based on application load)	Maximum: 0.6 kVA Typical: 0.4 kVA (varies based on application load)	Maximum: 0.7 kVA Typical: 0.5 kVA (varies based on application load)
Cooling	Maximum: 2,037 BTU/hour (2,149 kJ/hour) Typical: 1,426 BTU/hour (1,504 kJ/hour)	Maximum: 2,000 BTU/hour (2,109 kJ/hour) Typical: 1,400 BTU/hour (1,477 kJ/hour)	Maximum: 2,481 BTU/hour (2,617 kJ/hour) Typical: 1,736 BTU/hour (1,832 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	Maximum: Approximately 94 CFM Typical: Approximately 66 CFM	Maximum: Approximately 93 CFM Typical: Approximately 65 CFM	Maximum: Approximately 115 CFM Typical: Approximately 80 CFM

Table 1-20 Environmental Requirements for 8-Socket Servers, Model X4-8 and Later

Environmental Component	Database Server X7-8 and X8-8	Database Server X5-8 and X6-8	Database Server X4-8
Weight	90.2 kg (198.8 lbs)	93.4 kg (206 lbs)	93.4 kg (206 lbs)
Acoustic levels	87 dB	81 db(A)	81 db(A)
Power	Maximum: 3.6 kVA Typical: 2.5 kVA (varies based on application load)	Maximum: 3.5 kVA Typical: 2.4 kVA (varies based on application load)	Maximum: 3.7 kVA Typical: 2.6 kVA (varies based on application load)
Cooling	Maximum: 11,942 BTU/hour (12,599 kJ/hour) Typical: 8,360 BTU/hour (8,820 kJ/hour)	Maximum: 11,601 BTU/hour (12,239 kJ/hour) Typical: 8,121 BTU/hour (8,568 kJ/hour)	Maximum: 12,284 BTU/hour (12,959 kJ/hour) Typical: 8,599 BTU/hour (9,072 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	Maximum: Approximately 553 CFM Typical: Approximately 387 CFM	Maximum: Approximately 537 CFM Typical: Approximately 376 CFM	Maximum: Approximately 569 CFM Typical: Approximately 398 CFM

1.2 Space Requirements

The space requirements for Oracle Exadata Database Machine depend on the hardware rack that is supplied with the system.

- [Space Requirements for Oracle Exadata Database Machine X7-2 and Later Systems](#)

Oracle Exadata Database Machine X7-2 and later systems use a different rack than previous systems.

- [Space Requirements for Racks up to Oracle Exadata Database Machine X6](#)
All racks up to Oracle Exadata Database Machine X6-2 and X6-8 use the same hardware rack, and have the same space requirements.

1.2.1 Space Requirements for Oracle Exadata Database Machine X7-2 and Later Systems

Oracle Exadata Database Machine X7-2 and later systems use a different rack than previous systems.

The space requirements are as follows:

- Height: 2000 mm (78.74 inches)
- Width: 600 mm (23.62 inches)
- Depth: 1235 mm with doors attached, from handle to handle (48.6 inches)

The minimum ceiling height for the cabinet is 2914 mm (114.72 inches), measured from the true floor or raised floor, whichever is higher. This includes an additional 914 mm (36 inches) of space required above the rack height for maintenance access, as described in [Table 1-23](#). The space surrounding the cabinet must not restrict the movement of cool air between the air conditioner and the front of the systems within the cabinet, or the movement of hot air coming out of the rear of the cabinet.

1.2.2 Space Requirements for Racks up to Oracle Exadata Database Machine X6

All racks up to Oracle Exadata Database Machine X6-2 and X6-8 use the same hardware rack, and have the same space requirements.

The space requirements are as follows:

- Height: 1998 mm (78.66 inches)
- Width: 600 mm with side panels (23.62 inches)
- Depth: 1200 mm (47.24 inches)

The minimum ceiling height for the cabinet is 2912 mm (114.65 inches), measured from the true floor or raised floor, whichever is higher. This includes an additional 914 mm (36 inches) of space required above the rack height for maintenance access, as described in [Table 1-23](#). The space surrounding the cabinet must not restrict the movement of cool air between the air conditioner and the front of the systems within the cabinet, or the movement of hot air coming out of the rear of the cabinet.

1.3 Receiving, Unpacking, and Access Route Requirements

Before your Oracle Exadata Rack arrives, ensure that the receiving area is large enough for the package.

Use the following package dimensions for Oracle Exadata Rack:

- Shipping height: 2159 mm (85 inches)
- Shipping width: 1219 mm (48 inches)
- Shipping depth: 1575 mm (62 inches)

If your loading dock meets the height and ramp requirements for a standard freight carrier truck, then you can use a pallet jack to unload the rack. If the loading dock does not meet the requirements, then you must provide a standard forklift or other means to unload the rack. You can also request that the rack be shipped in a truck with a lift gate.

When Oracle Exadata Rack arrives, leave the rack in its shipping packaging until it arrives at its installation site. Use a conditioned space to remove the packaging material to reduce particles before entering the data center. The entire access route to the installation site should be free of raised-pattern flooring that can cause vibration.

Allow enough space for unpacking it from its shipping cartons. Ensure that there is enough clearance and clear pathways for moving Oracle Exadata Rack from the unpacking location to the installation location.

▲ Caution:

Prior to moving the rack, always make sure that all four leveling and stabilizing feet are raised and out of the way.

Table 1-21 Access Route Requirements

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

- [Rack Weights for Oracle Exadata Database Machine](#)
Review the net and shipping weights for your Oracle Exadata Rack to ensure it can be delivered safely.
- [Rack Weights for Oracle Exadata Storage Expansion Rack](#)

1.3.1 Rack Weights for Oracle Exadata Database Machine

Review the net and shipping weights for your Oracle Exadata Rack to ensure it can be delivered safely.

For weight information for Oracle Exadata Database Machine flexible configurations, use the Oracle Exadata Configuration Assistant, available on Oracle Technology Network.

- [Rack Weights for Oracle Exadata Database Machine 2-Socket Servers](#)
Use this table to determine the net and shipping weight of your Oracle Exadata Rack X*-2.
- [Rack Weights for Oracle Exadata Database Machine 8-Socket Servers](#)
Use this table to determine the net and shipping weight of your Oracle Exadata Rack X*-8.

Related Topics

- [Oracle Exadata Configuration Assistant \(OECA\) Downloads](#)

1.3.1.1 Rack Weights for Oracle Exadata Database Machine 2-Socket Servers

Use this table to determine the net and shipping weight of your Oracle Exadata Rack X*-2.

Type of Rack	Shipping Weight	Net Weight
Oracle Exadata Database Machine X8M-2 Full Rack	Extreme Flash: 971.2 kg (2,141.2 lbs) High Capacity: 1073.5 kg (2366.7 lbs)	Extreme Flash: 812.5 kg (1791.3 lbs) High Capacity: 914.8 kg (2016.8 lbs)
Oracle Exadata Database Machine X8M-2 Eighth Rack	Extreme Flash: 546.1 kg (1203.9 lbs) High Capacity: 557.8 kg (1229.7 lbs)	Extreme Flash: 387.4 kg (854 lbs) High Capacity: 399.1 kg (879.8 lbs)
Oracle Exadata Database Machine X8-2 Full Rack	Extreme Flash: 972.3 kg (2144 lbs) High Capacity: 1074.5 kg (2369 lbs)	Extreme Flash: 813.6 kg (1794 lbs) High Capacity: 915.8 kg (2019 lbs)
Oracle Exadata Database Machine X8-2 Eighth Rack	Extreme Flash: 547.2 kg (1206 lbs) High Capacity: 558.9 kg (1232 lbs)	Extreme Flash: 388.5 kg (856 lbs) High Capacity: 400.2 (882 lbs)
Oracle Exadata Database Machine X7-2 Full Rack	Extreme Flash: 974.0 kg (2147.2 lbs) High Capacity: 1071.1 kg (2361.4 lbs)	Extreme Flash: 815.1 kg (1797.2 lbs) High Capacity: 912.4 kg (2011.4 lbs)
Oracle Exadata Database Machine X7-2 Half Rack	Extreme Flash: 698.4 kg (1539.6 lbs) High Capacity: 746.9 kg (1646.7 lbs)	Extreme Flash: 539.6 kg (1189.6 lbs) High Capacity: 558.8 kg (1296.7 lbs)
Oracle Exadata Database Machine X7-2 Quarter Rack	Extreme Flash: 551.1 kg (1215.0 lbs) High Capacity: 571.9 kg (1260.9 lbs)	Extreme Flash: 392.4 kg (865.0 lbs) High Capacity: 413.2 kg (910.9 lbs)
Oracle Exadata Database Machine X7-2 Eighth Rack	Extreme Flash: 548.8 kg (1209.9 lbs) High Capacity: 558.9 kg (1232.2 lbs)	Extreme Flash: 390.1 kg (859.9 lbs) High Capacity: 400.1 kg (882.2 lbs)
Oracle Exadata Database Machine X6-2 Full Rack	Extreme Flash: 871.1 kg (1920.5 lbs) High Capacity: 941.0 kg (2074.5 lbs)	Extreme Flash: 794.5 kg (1751.5 lbs) High Capacity: 864.3 kg (1905.5 lbs)
Oracle Exadata Database Machine X6-2 Half Rack	Extreme Flash: 592.2 kg (1305.5 lbs) High Capacity: 627.1 kg (1382.5 lbs)	Extreme Flash: 515.5 kg (1136.5 lbs) High Capacity: 550.4 kg (1213.5 lbs)

Type of Rack	Shipping Weight	Net Weight
Oracle Exadata Database Machine X6-2 Quarter Rack	Extreme Flash: 432.5 kg (953.5 lbs) High Capacity: 447.5 kg (986.5 lbs)	Extreme Flash: 366.3 kg (807.5 lbs) High Capacity: 381.2 kg (840.5 lbs)
Oracle Exadata Database Machine X6-2 Eighth Rack	Extreme Flash: 432.5 kg (953.5 lbs) High Capacity: 433.8 kg (956.3 lbs)	Extreme Flash: 366.3 kg (807.5 lbs) High Capacity: 376.6 kg (810.3 lbs)
Oracle Exadata Database Machine X5-2 Full Rack	984.8 kg (2171 lbs)	901.3 kg (1987 lbs)
Oracle Exadata Database Machine X5-2 Half Rack	639.6 kg (1410 lbs)	587.4 kg (1295 lbs)
Oracle Exadata Database Machine X5-2 Quarter Rack	453.1 kg (999 lbs)	418.2 kg (922 lbs)
Oracle Exadata Database Machine X5-2 Eighth Rack	453.1 kg (999 lbs)	418.2 kg (922 lbs)
Oracle Exadata Database Machine X4-2 Full Rack	936 kg (2064 lbs)	852.8 kg (1880 lbs)
Oracle Exadata Database Machine X4-2 Half Rack	577 kg (1273 lbs)	525.3 kg (1158 lbs)
Oracle Exadata Database Machine X4-2 Quarter Rack	401 kg (883 lbs)	365.6 kg (806 lbs)
Oracle Exadata Database Machine X4-2 Eighth Rack	401 kg (883 lbs)	365.6 kg (806 lbs)
Oracle Exadata Database Machine X3-2 Full Rack	954.8 kg (2105 lbs)	871.4 kg (1921 lbs)
Oracle Exadata Database Machine X3-2 Half Rack	596.0 kg (1314 lbs)	543.9 kg (1199 lbs)
Oracle Exadata Database Machine X3-2 Quarter Rack	403.7 kg (890 lbs)	368.8 kg (813 lbs)
Oracle Exadata Database Machine X3-2 Eighth Rack	403.7 kg (890 lbs)	368.8 kg (813 lbs)
Oracle Exadata Database Machine X2-2 Full Rack	1050.1 kg (2315 lbs)	966.61 kg (2131 lbs)
Oracle Exadata Database Machine X2-2 Half Rack	684.1 kg (1508 lbs)	602.8 kg (1329 lbs)
Oracle Exadata Database Machine X2-2 Quarter Rack	487.6 kg (1075 lbs)	409.1 kg (902 lbs)

1.3.1.2 Rack Weights for Oracle Exadata Database Machine 8-Socket Servers

Use this table to determine the net and shipping weight of your Oracle Exadata Rack X*-8.

Type of Rack	Shipping Weight	Net Weight
Oracle Exadata Database Machine X8M-8 Full Rack	Extreme Flash: 827.4 kg (2174 lbs) High Capacity: 1,088.3 kg (2399.3 lbs)	Extreme Flash: 827.4 kg (1824.1 lbs) High Capacity: 929.6 kg (2049.5 lbs)
Oracle Exadata Database Machine X8M-8 Half Rack	Extreme Flash: 701 kg (1545.4 lbs) High Capacity: 722.9 kg (1593.7 lbs)	Extreme Flash: 542.3 kg (1195.5 lbs) High Capacity: 564.2 kg (1243.8 lbs)
Oracle Exadata Database Machine X8-8 Full Rack	Extreme Flash: 987.2 kg (2176 lbs) High Capacity: 1089.4 kg (2402 lbs)	Extreme Flash: 828.5 kg (1826.5 lbs) High Capacity: 930.7 kg (2051.9 lbs)
Oracle Exadata Database Machine X8-8 Half Rack	Extreme Flash: 702.1 kg (1547.9 lbs) High Capacity: 724 kg (1596.2 lbs)	Extreme Flash: 543.4 kg (1197.9 lbs) High Capacity: 565.3 kg (1246.2 lbs)
Oracle Exadata Database Machine X7-8 Full Rack	Extreme Flash drives: 988.8 kg (2180 lbs) High capacity drives: 1085.9 kg (2394 lbs)	Extreme Flash drives: 830.1 kg (1830 lbs) High capacity drives: 927.2 kg (2044.2 lbs)
Oracle Exadata Database Machine X7-8 Half Rack	Extreme Flash: 703.6 kg (1551.2 lbs) High Capacity: 724.5 kg (1597.2 lbs)	Extreme Flash: 544.9 kg (1201.4 lbs) High Capacity: 565.8 kg (1247.3 lbs)
Oracle Exadata Database Machine X5-8 or X6-8 Full Rack	Extreme Flash drives: 894.7 kg (1972.5 lbs) High capacity drives: 964.6 kg (2126.5 lbs)	Extreme Flash drives: 828.5 kg (1826.5 lbs) High capacity drives: 898.3 kg (1980.5 lbs)
Oracle Exadata Database Machine X6-8 Half Rack	Extreme Flash drives: 592.2 kg (1305.5 lbs) High capacity drives: 607.1 kg (1338.5 lbs)	Extreme Flash drives: 525.9 kg (1159.5 lbs) High capacity drives: 540.9 kg (1192.5 lbs)
Oracle Exadata Database Machine X5-8 Half Rack	Extreme Flash drives: 592.4 kg (1306.1 lbs) High capacity drives: 607.4 kg (1339.1 lbs)	Extreme Flash drives: 526.2 kg (1160.1 lbs) High capacity drives: 541.2 kg (1193.1 lbs)
Oracle Exadata Database Machine X4-8 Full Rack	High performance drives: 951 kg (2097 lbs) High capacity drives: 1019 kg (2247 lbs)	High performance drives: 867 kg (1912 lbs) High capacity drives: 936 kg (2063 lbs)
Oracle Exadata Database Machine X3-8 Full Racks with Exadata Storage Server X4-2L Servers	994.7 kg (2193 lbs)	911.3 kg (2009 lbs)
Oracle Exadata Database Machine X3-8 Full Racks	1013.3 kg (2234 lbs)	929.9 kg (2050 lbs)

Type of Rack	Shipping Weight	Net Weight
Oracle Exadata Database Machine X2-8 Full Rack with Sun Fire X4800 Oracle Database Servers	1026.9 kg (2264 lbs)	943.8 kg (2080 lbs)
Oracle Exadata Database Machine X2-8 Full Rack with Sun Server X2-8 Oracle Database Servers	1036.0 kg (2284 lbs)	980.7 kg (2162 lbs)

1.3.2 Rack Weights for Oracle Exadata Storage Expansion Rack

For weight information for Oracle Exadata Database Machine flexible configurations, use the Oracle Exadata Configuration Assistant, available on Oracle Technology Network: Oracle Exadata Configuration Assistant (OECA) Downloads

Table 1-22 Rack Weights for Oracle Exadata Storage Expansion Rack

Type of Rack	Shipping Weight	Net Weight
Oracle Exadata Storage Expansion Rack X8M-2 Full Rack	Extreme Flash: 954.9 kg (2105.2 lbs) High Capacity: 1093.6 kg (2411 lbs)	Extreme Flash: 796.2 kg (1755.3 lbs) High Capacity: 934.9 kg (2061.1 lbs)
Oracle Exadata Storage Expansion Rack X8M-2 Quarter Rack	Extreme Flash: 544.8 kg (1201.1 lbs) High Capacity: 574.1 kg (1265.7 lbs)	Extreme Flash: 386.1 kg (851.3 lbs) High Capacity: 415.4 kg (915.7 lbs)
Oracle Exadata Storage Expansion Rack X8-2 Full Rack	Extreme Flash: 956.5 kg (2108.7 lbs) High Capacity: 1095.3 kg (2415 lbs)	Extreme Flash: 797.8 kg (1758.9 lbs) High Capacity: 936.6 kg (2064.8 lbs)
Oracle Exadata Storage Expansion Rack X8-2 Quarter Rack	Extreme Flash: 546.5 kg (1205 lbs) High Capacity: 565.3 kg (1246 lbs)	Extreme Flash: 387.8 kg (854.9 lbs) High Capacity: 406.6 kg (896.4 lbs)
Oracle Exadata Storage Expansion Rack X7-2 Full Rack	Extreme Flash: 850.3 kg (1874.5 lbs) High Capacity: 945.1 kg (2083.5 lbs)	Extreme Flash: 784.0 kg (1728.5 lbs) High Capacity: 878.8 kg (1937.5 lbs)
Oracle Exadata Storage Expansion Rack X7-2 Half Rack	Extreme Flash: 569.0 kg (1254.5 lbs) High Capacity: 613.9 kg (1353.5 lbs)	Extreme Flash: 502.8 kg (1108.5 lbs) High Capacity: 547.7 kg (1207.5 lbs)
Oracle Exadata Storage Expansion Rack X7-2 Quarter Rack	Extreme Flash: 430.7 kg (949.5 lbs) High Capacity: 450.6 kg (993.5 lbs)	Extreme Flash: 364.5 kg (803.5 lbs) High Capacity: 384.4 kg (847.5 lbs)

Table 1-22 (Cont.) Rack Weights for Oracle Exadata Storage Expansion Rack

Type of Rack	Shipping Weight	Net Weight
Oracle Exadata Storage Expansion Rack X6-2 Full Rack	Extreme Flash: 850.3 kg (1874.5 lbs) High Capacity: 945.1 kg (2083.5 lbs)	Extreme Flash: 784.0 kg (1728.5 lbs) High Capacity: 878.8 kg (1937.5 lbs)
Oracle Exadata Storage Expansion Rack X6-2 Half Rack	Extreme Flash: 569.0 kg (1254.5 lbs) High Capacity: 613.9 kg (1353.5 lbs)	Extreme Flash: 502.8 kg (1108.5 lbs) High Capacity: 547.7 kg (1207.5 lbs)
Oracle Exadata Storage Expansion Rack X6-2 Quarter Rack	Extreme Flash: 430.7 kg (949.5 lbs) High Capacity: 450.6 kg (993.5 lbs)	Extreme Flash: 364.5 kg (803.5 lbs) High Capacity: 384.4 kg (847.5 lbs)
Oracle Exadata Storage Expansion Rack X5-2 Full Rack	994.7 kg (2193 lbs)	915.8 kg (2019 lbs)
Oracle Exadata Storage Expansion Rack X5-2 Half Rack	640.9 kg (1413 lbs)	591.5 kg (1304 lbs)
Oracle Exadata Storage Expansion Rack X5-2 Quarter Rack	444.5 kg (980 lbs)	411 kg (906.1 lbs)
Oracle Exadata Storage Expansion Rack X4-2 Full Rack	903.1 kg (1991 lbs)	824.2 kg (1817.1 lbs)
Oracle Exadata Storage Expansion Rack X4-2 Half Rack	562.4 kg (1239.9 lbs)	513.0 kg (1131 lbs)
Oracle Exadata Storage Expansion Rack X4-2 Quarter Rack	380.1 kg (838 lbs)	346.5 kg (764 lbs)
Oracle Exadata Storage Expansion Rack X3-2 Full Rack	906.3 kg (1998 lbs)	827.4 kg (1824.1 lbs)
Oracle Exadata Storage Expansion Rack X3-2 Half Rack	565.6 kg (1247 lbs)	516.2 kg (1138 lbs)
Oracle Exadata Storage Expansion Rack X3-2 Quarter Rack	383.3 kg (845 lbs)	349.7 kg (771 lbs)
Oracle Exadata Storage Expansion Rack Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers	1001.1 kg (2207 lbs)	917.6 kg (2023 lbs)

Table 1-22 (Cont.) Rack Weights for Oracle Exadata Storage Expansion Rack

Type of Rack	Shipping Weight	Net Weight
Oracle Exadata Storage Expansion Rack Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers	659.6 kg (1454.2 lbs)	578.3 kg (1275 lbs)
Oracle Exadata Storage Expansion Rack Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers	475.3 kg (1047.9 lbs)	396.8 kg (874.8 lbs)

1.4 Maintenance Access Requirements

The maintenance area must be large enough for Oracle Exadata Rack, and have the required access space.

The required space to remove the side panels is 675.64 mm (26.6 inches). Side panel access is not required for normal maintenance.

Front access space requirement is typically the length of the longest chassis supported in the rack, plus the distance the rails need to travel out to install and remove equipment from the front.

Rear access space requirement is the space needed for a person working behind the Oracle Exadata Rack.

Top access space requirement is for a person standing on a ladder to be able service components installed in the top of the rack.

Table 1-23 Maintenance Access Requirements for Oracle Exadata Rack (All Models)

Location	Maintenance Access Requirement
Rear maintenance	914 mm (36 inches)
Front maintenance	1232 (48.5 inches)
Top maintenance	914 mm (36 inches)

 **Note:**

- In raised floor data centers, an adjacent tile at the rear of the rack must be able to be opened to access any cabling going under the floor.
- Oracle Exadata Racks can be placed side-by-side when interconnecting the racks.

1.5 Flooring Requirements

Oracle Exadata Rack may be installed on raised floor or solid floor environments.

The site floor or the raised flooring must be able to support the total weight of Oracle Exadata Rack.

Table 1-24 Floor Load Requirements for Oracle Exadata Rack

Description	Requirement
Maximum allowable weight of installed rack equipment	952.5 kg (2100 lbs)
Maximum allowable weight of installed power distribution units	52.16 kg (115 lbs)
Maximum dynamic load (maximum allowable weight of installed equipment including PDUs)	1004.66 kg (2215 lbs)



Note:

For raised floors, a floor tile with a cutout at the rear of the Exadata rack is needed for routing power and network cables into the rack from under the raised floor. The cutout is located between the two rear leveling feet. See [Figure 5-3](#) and [Figure 5-5](#).

1.6 Electrical Power Requirements

The Oracle Exadata Rack can operate effectively over a wide range of voltages and frequencies. However, each rack must have a reliable power source.

Damage may occur if the ranges are exceeded. Electrical disturbances such as the following may damage Oracle Exadata Rack:

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

To protect Oracle Exadata Rack from such disturbances, you should have a dedicated power distribution system, power-conditioning equipment, as well as lightning arresters or power cables to protect from electrical storms.

- [PDU Power Requirements](#)
Each rack has two pre-installed power distribution units (PDUs). The PDUs accept different power sources. You must specify the type of PDU that is correct for your Oracle Exadata Rack and data center.
- [Facility Power Requirements](#)
To prevent catastrophic failures, design the input power sources to ensure adequate power is provided to the PDUs.

- [Circuit Breaker Requirements](#)
If computer equipment is subjected to repeated power interruptions and fluctuations, then it is susceptible to a higher rate of component failure.
- [Electrical Grounding Guidelines](#)
The cabinets for Oracle Exadata Rack are shipped with grounding-type power cords.

1.6.1 PDU Power Requirements

Each rack has two pre-installed power distribution units (PDUs). The PDUs accept different power sources. You must specify the type of PDU that is correct for your Oracle Exadata Rack and data center.

The following are the PDUs that Oracle Exadata Database Machine currently supports or has supported in the past:

- [Low-voltage 15 kVA Single Phase PDUs for North America, South America, Japan and Taiwan](#)
- [Low-voltage 15 kVA Three Phase PDUs for North America, South America, Japan and Taiwan](#)
- [High-voltage 15 kVA Single Phase PDUs for Europe, the Middle East and Africa \(EMEA\), and Asia Pacific \(APAC\), except for Japan and Taiwan](#)
- [High-voltage 15 kVA Three Phase for Europe, the Middle East and Africa \(EMEA\), and Asia Pacific \(APAC\), except for Japan and Taiwan](#)
- [Low-voltage 22 kVA Single Phase PDUs for North America, South America, Japan and Taiwan](#)
- [High-voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa \(EMEA\), and Asia Pacific \(APAC\), except for Japan and Taiwan](#)
- [Low-voltage 24 kVA Three Phase PDUs for North America, South America, Japan and Taiwan](#)
- [High-voltage 24 kVA Three Phase PDUs for Europe, the Middle East and Africa \(EMEA\), and Asia Pacific \(APAC\), except for Japan and Taiwan](#)

1.6.1.1 Low-voltage 15 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Table 1-25 Specifications for Low-voltage 15 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Specification	Required values
Voltage	200-240 VAC
Frequency	50/60 Hz
Current	24 A maximum per input
Power rating	15 kVA
Output current	72 A (3 x 24 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6

Table 1-25 (Cont.) Specifications for Low-voltage 15 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Specification	Required values
Group protection (UL489 2-pole circuit breaker)	20A

Materials Required

The following are needed to connect an Oracle Exadata rack to a low-voltage three-phase power source:

- 6 power cords for two PDUs, each rated for 30 amperes at 200-240 VAC
- 6 receptacles for 2 PDUs, each must be 15kVA, with three 30A/250V 2-pole/3-wire NEMA L6-30P plugs

Figure 1-1 Low-voltage Single Phase Power Connector for North America, South America, Japan and Taiwan



1.6.1.2 Low-voltage 15 kVA Three Phase PDUs for North America, South America, Japan and Taiwan

Table 1-26 Specifications for Low-voltage 15 kVA Three Phase PDUs for North America, South America, Japan and Taiwan

Specification	Required values, per PDU
Voltage	190-220 VAC
Frequency	50/60 Hz
Current	40 A maximum per phase
Power rating	14.4 kVA
Output current	69.3 A (3 x 23.1A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 2-pole circuit breaker)	20A
Data center receptacle	IEC309-3P4W-IP67 60A 250VAC 3ph (Hubbell equivalent is HBL460R9W)

Materials Required

The following are needed to connect an Oracle Exadata system to a low-voltage three-phase power source:

- 2 power cords for two PDUs, 60 amperes at 190-220 VAC three-phase
- 2 receptacles to connect the PDUs to 2 IEC 60309 60A 4-pin 250VAC 3ph IP67 data center receptacles

Figure 1-2 Low-voltage Three Phase Power Connector for North America, South America, Japan and Taiwan



1.6.1.3 High-voltage 15 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Table 1-27 Specifications for High-voltage 15 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Required values, per PDU
Voltage	220-240 VAC
Frequency	50/60 Hz
Current	24 A maximum per input
Power rating	15 kVA
Output current	72 A (3 x 24 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 2-pole circuit breaker)	20A
Data center receptacle	15kVA, IEC309-2P3W-IP44 32A 250VAC (Hubbell equivalent is HBL332R6W)

 **Note:**

The high-voltage 15kVA single phase PDU is no longer available. The high-voltage 22kVA is a compatible replacement with the same receptacle and breaker requirements

Figure 1-3 High-voltage Single Phase Power Connector



1.6.1.4 High-voltage 15 kVA Three Phase for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Table 1-28 Specifications for High-voltage 15 kVA Three Phase for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Required values, per PDU
Voltage	220/380-240/415 VAC 3-phase
Frequency	50/60 Hz
Current	25 A maximum per phase
Power rating	14.4 kVA
Output current	62.7 A (3 x 20.9 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 1-pole circuit breaker)	20A
Data center receptacle	15kVA, three phase, 5-pin, IEC 60309 32A, 5-pin 230/400V, three phase IP44

Materials Required

The following are needed to connect an Oracle Exadata system to a high-voltage three-phase power source:

- 2 power cords for two PDUs, 25 amperes at 220/380-240/415 VAC three-phase
- 2 receptacles to connect the PDUs to 2 IEC 60309 32A 5-pin 230/400VAC 3ph IP44 data center receptacles

Figure 1-4 High-voltage Three Phase Power Connector



1.6.1.5 Low-voltage 22 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Table 1-29 Specifications for Low-voltage 22 kVA Single Phase PDUs for North America, South America, Japan and Taiwan

Specifications	Required values, per PDU
Voltage	200-240 VAC
Frequency	50/60 Hz
Current	36.8 A maximum per input / 110.4 A Maximum per PDU
Power rating	22 kVA
Output current	110.4A (3x 36.8)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 2-pole circuit breaker)	20A
Data center receptacle	Hubbell CS8269 or CS8264

Materials Required

The following are needed to connect an Oracle Exadata system to a low-voltage single-phase power source:

- 3 power cords for two PDUs
- 6 receptacles to connect the PDUs

The following image shows the low-voltage, 22 kVA single phase PDU power connector for North America, South America, Japan and Taiwan. This power connector is different from the power connector for the low-voltage 15 kVA single phase PDU.

Figure 1-5 Low-voltage 22 kVA Single Phase Power Connector



1.6.1.6 High-voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Table 1-30 Specifications for High-voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specifications	Required values, per PDU
Voltage	220-240 VAC
Frequency	50/60 Hz
Current	32 A maximum per input
Power rating	22 kVA
Output current	96 A (3 x 32 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 1-pole circuit breaker)	20A
Data center receptacle	IEC309-2P3W-IP44 32A 250VAC (Hubbell equivalent is HBL332R6W)

Materials Required

 **Note:**

The high-voltage 15kVA single phase PDU is no longer available. The high-voltage 22kVA is a compatible replacement with the same receptacle and breaker requirements

The following are needed to connect an Oracle Exadata system to a high-voltage one-phase power source:

- 6 power cords for two PDUs, each rated for 25 amperes at 220/380-240/415 VAC 1-phase voltage
- 6 receptacles to connect the PDUs

Figure 1-6 High-voltage Single Phase Power Connector



1.6.1.7 Low-voltage 24 kVA Three Phase PDUs for North America, South America, Japan and Taiwan

Table 1-31 Specifications for Low-voltage 24 kVA Three Phase PDUs for North America, South America, Japan and Taiwan

Specifications	Required values, per PDU
Voltage	200-208 VAC 3-phase
Frequency	50/60 Hz
Current	34.6 A maximum per phase
Power rating	25 kVA
Output current	120 A (6 x 20 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 2-pole circuit breaker)	20A
Data center receptacle	IEC309-3P4W-IP67 60A 250VAC 3ph (Hubbell equivalent is HBL460R9W)

Materials Required

The following are needed to connect an Oracle Exadata system to a low-voltage three-phase power source:

- 4 power cords for two PDUs
- 4 receptacles to connect the PDUs

1.6.1.8 High-voltage 24 kVA Three Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Table 1-32 Specifications for High-voltage 24 kVA Three Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specifications	Required values, per PDU
Voltage	220/380-240/415 VAC 3-phase
Frequency	50/60 Hz
Current	18.1 A maximum per phase
Power rating	25 kVA
Output current	108.6A (6 x 18.1 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection (UL489 1-pole circuit breaker)	20A

Table 1-32 (Cont.) Specifications for High-voltage 24 kVA Three Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specifications	Required values, per PDU
Data center receptacle	IEC309-4P5W-IP44 32A 400VAC 3ph (Hubbell equivalent is HBL532R6W)

Materials Required

The following are needed to connect an Oracle Exadata system to a high-voltage three-phase power source:

- 4 power cords for two PDUs
- 4 receptacles to connect the PDUs

1.6.2 Facility Power Requirements

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

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

 **Note:**

Electrical work and installations must comply with applicable local, state, or national electrical codes. Contact your facilities manager or qualified electrician to determine what type of power is supplied to the building.

PDU power cords are 4 meters (13.12 feet) long, and 1 to 1.5 meters (3.3 to 4.9 feet) of the cord is routed within the rack cabinet. The installation site AC power receptacle must be within 2 meters (6.6 feet) of the rack.

1.6.3 Circuit Breaker Requirements

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

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

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

 **Note:**

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

1.6.4 Electrical Grounding Guidelines

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

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

1.7 Temperature and Humidity Requirements

Excessive internal temperatures may result in full or partial shut down of Oracle Exadata Database Machine.

Airflow through Oracle Exadata Rack is from front to back. Refer to [General Environmental Requirements](#) for information on cooling and airflow.

 **Note:**

Studies have shown that temperature increases of 10 degrees Celsius (15 degrees Fahrenheit) above 20 degrees Celsius (70 degrees Fahrenheit) reduce long-term electronics reliability by 50 percent.

The following table lists the temperature, humidity and altitude requirements for operating and non-operating machines.

Table 1-33 Temperature, Humidity and Altitude Requirements

Condition	Operating Requirement	Non-operating Requirement	Optimum
Temperature	5 to 32 degrees Celsius (41 to 89.6 degrees Fahrenheit)	-40 to 70 degrees Celsius (-40 to 158 degrees Fahrenheit).	For optimal rack cooling, data center temperatures from 21 to 23 degrees Celsius (70 to 74 degrees Fahrenheit)
Relative humidity	10 to 90 percent relative humidity, non-condensing	Up to 93 percent relative humidity.	For optimal data center rack cooling, 45 to 50 percent, non-condensing

Table 1-33 (Cont.) Temperature, Humidity and Altitude Requirements

Condition	Operating Requirement	Non-operating Requirement	Optimum
Altitude	3048 meters (10000 feet) maximum	12000 meters (40000 feet).	Ambient temperature is reduced by 1 degree Celsius per 300 m above 900 m altitude above sea level

Set conditions to the optimal temperature and humidity ranges to minimize the chance of downtime due to component failure. Operating Oracle Exadata Rack for extended periods at or near the operating range limits, or installing it in an environment where it remains at or near non-operating range limits could significantly increase hardware component failure.

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

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

- Optimal range helps protect computer systems from corrosion problems associated with high humidity levels.
- Optimal range provides the greatest operating time buffer in the event of air conditioner control failure.
- This range helps avoid failures or temporary malfunctions caused by intermittent interference from static discharges that may occur when relative humidity is too low.

Electrostatic discharge (ESD) is easily generated, and hard to dissipate in areas of low relative humidity, such as below 35 percent. ESD becomes critical when humidity drops below 30 percent. It is not difficult to maintain humidity in a data center because of the high-efficiency vapor barrier and low rate of air changes normally present.

1.8 Ventilation and Cooling Requirements

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

Do not obstruct the front or rear of the rack with equipment or objects that might prevent air from flowing through the rack. Rack-mountable servers and equipment typically draw cool air in through the front of the rack and let out warm air out the rear of the rack. There is no air flow requirement for the left and right sides due to front-to-back cooling.

If the rack is not completely filled with components, then cover the empty sections with filler panels. Gaps between components can adversely affect air flow and cooling within the rack.

Relative humidity is the percentage of the total water vapor that can exist in the air without condensing, and is inversely proportional to air temperature. Humidity goes down when the temperature rises, and goes up when the temperature drops. For example, air with a relative humidity of 45 percent at a temperature of 24 degrees Celsius (75 degrees Fahrenheit) has a relative humidity of 65 percent at a temperature of 18 degrees Celsius (64 degrees Fahrenheit). As the temperature drops, the relative humidity rises to more than 65 percent, and water droplets are formed.

Air conditioning facilities usually do not precisely monitor or control temperature and humidity throughout an entire computer room. Generally, monitoring is done at individual points corresponding to multiple exhaust vents in the main unit, and other units in the room. Special consideration should be paid to humidity when using underfloor ventilation. When underfloor ventilation is used, monitoring is done at each point close to an exhaust vent. Distribution of the temperature and humidity across the entire room is uneven.


Oracle Exadata Racks have been designed to function while mounted in a natural convection air flow. The following requirements must be followed to meet the environmental specification:

- Ensure there is adequate air flow through the server.
- Ensure the server has front-to-back cooling. The air inlet is at the front of the server, and the air is let out the rear.
- Allow a minimum clearance of 1219.2 mm (48 inches) at the front of the server, and 914 mm (36 inches) at the rear of the server for ventilation.

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

Type of Oracle Exadata Database Machine or Oracle Exadata Storage Expansion Rack	Recommended Number of Floor Tiles
Oracle Exadata full rack systems	4
Oracle Exadata half rack systems	3
Oracle Exadata quarter and eighth rack systems	1

Figure 1-7 shows a typical installation of the floor tiles for Oracle Exadata Database Machine Full Rack in a typical data center.

Figure 1-7 Typical Data Center Configuration for Perforated Floor Tiles

1.9 Network Connection and IP Address Requirements for Oracle Exadata Database Machine

The network cables must be run from your network equipment to the location of Oracle Exadata Database Machine. The requirements to connect to your existing Ethernet network infrastructure are described in the following sections:

- [Network Connection Requirements for Oracle Exadata Database Machine](#)
- [DNS Configuration for Oracle Exadata Database Machine](#)
- [IP Address Requirements for Oracle Exadata Database Machine](#)
Oracle Exadata Database Machine requires a large number of host names and IP addresses during initial configuration.
- [Calculating the IP Address Requirements when Using Oracle VMs](#)
The number of IP addresses depends on the physical deployment, as well as the number of virtual user domains.

1.9.1 Network Connection Requirements for Oracle Exadata Database Machine

Prior to installation, network cables must be run from your existing network infrastructure to the installation site. The requirements to connect Oracle Exadata Database Machine to your existing network infrastructure are as follows:

- Management network connection requirements for all Oracle Exadata Database Machine systems except for X2-n systems:

- 1 Ethernet connection for the Ethernet switch in the rack to the existing management network.
- Management network connection requirements for Oracle Exadata Database Machine X2-n:
 - 1 Ethernet connection for the Ethernet switch in the rack to the existing management network.
 - 1 Ethernet connection for the KVM switch in the rack to the existing management network.
- Client access network connection requirements
 - One Ethernet connection per database server to the client access network. The number depends on the system type (Oracle Exadata Database Machine Full Rack, Oracle Exadata Database Machine Half Rack, or Oracle Exadata Database Machine Quarter Rack), and whether network channel bonding is used. Use the following formula to determine the number of connections you need:

For bonded configurations, you need two network connections for each database node.

For non-bonded configurations, you need one network connection for each database node.
- Additional network NET2 connection requirements

One Ethernet connection per database server to connect NET2 to your existing network. The number depends on the system type, as follows:

For bonded configurations, you need two network connections for each database node.

For non-bonded configurations, you need one network connection for each database node.

 **Note:**

- This network is unavailable if network channel bonding is used for the client access network.
 - This network is optional. A cable is not required unless the network is configured during deployment.
- Additional network NET3 connection requirements

One Ethernet connection per database server to connect NET3 to your existing network. The number depends on the system type, as follows:

 - 8 connections for Oracle Exadata Database Machine 2-socket, Full Rack systems
 - 4 connections for Oracle Exadata Database Machine 2-socket, Half Rack systems
 - 2 connections for Oracle Exadata Database Machine 2-socket, Quarter Rack systems

- 2 connections for Oracle Exadata Database Machine 8-socket, Full Rack systems

**Note:**

This network is optional. A cable is not required unless the network is configured during deployment.

Oracle Exadata Database Machine X2-8 Full Rack has a total of 16 Ethernet ports per server. The additional networks for Oracle Exadata Database Machine X2-8 Full Rack are NET2 through NET15.

The following table shows examples of connection requirements for different sample environments. The table applies to two-socket Oracle Exadata Database Machine models X3-2 to X6-2.

Table 1-34 Examples of Network Configurations and Connection Requirements for Oracle Exadata Database Machine 2-Socket Systems X3-2 to X6-2

Scenario	Network Requirements
Core Requirements	<ul style="list-style-type: none"> • 1 for Ethernet switch • 2 for PDU monitoring
Bonded network configuration, and connected to a dedicated backup network	<ul style="list-style-type: none"> • 2 per database node for client access for bonded Ethernet interface on database servers • 1 per database node for backup network access for database servers <p>Total = 3 + (3 * number of database nodes)</p>
Non-bonded network configuration, and connected to a dedicated backup network and a dedicated disaster recovery network	<ul style="list-style-type: none"> • 1 for each database node for client access for the database servers • 1 per database node for backup network access for the database servers • 1 per database node for disaster recovery network access for the database servers <p>Total = 3 + (3 * number of database nodes)</p>
Non-bonded network configuration, and not connected to an additional network	<ul style="list-style-type: none"> • 1 per database node for database servers <p>Total = 3 + (1 * database node)</p>
Non-bonded network configuration, and connected to a dedicated backup network	<ul style="list-style-type: none"> • 1 for each database node for client access for the database servers • 1 per database node for backup network access for the database servers <p>Total = 3 + (2 * number of database nodes)</p>

The following table shows examples of connection requirements for different sample environments with Oracle Exadata Database Machine X2-2.

Table 1-35 Examples of Network Configurations and Connection Requirements for Oracle Exadata Database Machine X2-2

Scenario	Network Requirements
Oracle Exadata Database Machine X2-2 Full Rack using a bonded network configuration, and connected to a dedicated backup network	Requires 28 connections to your existing networks <ul style="list-style-type: none"> • 1 for Ethernet switch • 1 for KVM switch • 16 for database client access for bonded Ethernet interface on database servers • 8 for backup network access for database servers • 2 for monitoring the PDUs
Oracle Exadata Database Machine X2-2 Full Rack using a non-bonded network configuration, and connected to a dedicated backup network and a dedicated disaster recovery network	Requires 28 connections to your existing networks <ul style="list-style-type: none"> • 1 for Ethernet switch • 1 for KVM switch • 8 for database client access for the database servers • 8 for backup network access for the database servers • 8 for disaster recovery network access for the database servers • 2 for monitoring the PDUs
Oracle Exadata Database Machine X2-2 Half Rack using a non-bonded network configuration, and not connected to an additional network	Requires 8 connections to your existing networks <ul style="list-style-type: none"> • 1 for Ethernet switch • 1 for KVM switch • 4 for database servers • 2 for monitoring the PDUs
Oracle Exadata Database Machine X2-2 Quarter Rack using a non-bonded network configuration, and connected to a dedicated backup network	Requires 8 connections to your existing networks <ul style="list-style-type: none"> • 1 for Ethernet switch • 1 for KVM switch • 2 for database client access for the database servers • 2 for backup network access for the database servers • 2 for monitoring the PDUs

Related Topics

- [Understanding the Network Requirements for Oracle Exadata Database Machine](#)
Review the network requirements for Oracle Exadata Database Machine before installing or configuring the hardware.

1.9.2 DNS Configuration for Oracle Exadata Database Machine

Prior to receiving your Oracle Exadata Database Machine, use Oracle Exadata Database Machine Deployment Assistant. The assistant generates a file to be used when setting up the system. The host names and IP addresses specified in the assistant-generated file for the management network, client access network, and

additional networks should be registered in Domain Name System (DNS) prior to initial configuration. In addition, all public addresses, single client access name (SCAN) addresses, and VIP addresses should be registered in DNS prior to installation.

The assistant-generated file defines the SCAN as a single name with three IP addresses on the client access network. The three SCAN addresses provide service access for clients to Oracle Exadata Database Machine. Configure DNS for round robin resolution for the SCAN name to these three SCAN addresses.

All addresses registered in DNS must be configured for both forward resolution and reverse resolution. Reverse resolution must be forward confirmed (forward-confirmed reverse DNS) such that both the forward and reverse DNS entries match each other.

 **Note:**

Grid Naming Service (GNS) is not configured on Oracle Exadata Database Machine during initial configuration.

 **See Also:**

- ["Using Oracle Exadata Deployment Assistant"](#)
- *Oracle Grid Infrastructure Installation and Upgrade Guide for Linux* for additional information about SCAN addresses
- Your DNS vendor documentation for additional information about configuring round robin name resolution

1.9.3 IP Address Requirements for Oracle Exadata Database Machine

Oracle Exadata Database Machine requires a large number of host names and IP addresses during initial configuration.

The number of IP addresses required for a particular network, such as the management network, depends on the type of system. The network configuration, such as host names and IP addresses, used during installation is generated from information when completing the Oracle Exadata Deployment Assistant (OEDA).

Configure the new IP addresses in your existing networks only after you have run OEDA. All IP addresses in the assistant-generated file must be unassigned at the time of initial configuration. In addition, all IP addresses must be statically-assigned IP addresses, not dynamically-assigned (DHCP) addresses.

All RDMA Network Fabric IP addresses must be in the same subnet, with a minimum subnet mask of 255.255.240.0 (or /20). The subnet mask chosen should be wide enough to accommodate possible future expansion of the Oracle Exadata Database Machine and the internal network.

- [IP Address Requirements for Oracle Exadata Database Machine Systems with RoCE Network Fabric \(X8M\)](#)
Review the requirements for Oracle Exadata Database Machine systems that use RoCE Network Fabric (X8M).
- [IP Address Requirements for Oracle Exadata Database Machine X5 to X8 Systems with InfiniBand Network Fabric](#)
- [IP Address Requirements for Oracle Exadata Database Machine X4 Systems](#)
- [IP Address Requirements for Oracle Exadata Database Machine X3 Systems](#)
- [IP Address Requirements for Oracle Exadata Database Machine X2 Systems](#)

Related Topics

- [Calculating the IP Address Requirements when Using Oracle VMs](#)
The number of IP addresses depends on the physical deployment, as well as the number of virtual user domains.
- [Using Oracle Exadata Deployment Assistant](#)

1.9.3.1 IP Address Requirements for Oracle Exadata Database Machine Systems with RoCE Network Fabric (X8M)

Review the requirements for Oracle Exadata Database Machine systems that use RoCE Network Fabric (X8M).

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2, NET3, and so on) Requirements	RoCE Network Fabric Requirements
Base configuration for 2-socket systems (X8M-2) OR Eighth Rack configuration for 2-socket systems (X8M-2)	Minimum of 15 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 3 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 3 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 1 for each RoCE Network Fabric switch, and 1 for the Ethernet switch • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	For each additional network, a minimum of 2 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for the database servers, one per database server 	Minimum of 10 IP addresses: <ul style="list-style-type: none"> • 4 IP addresses, two per database server • 6 IP addresses, two per Exadata Storage Server
Flexible configuration for 2-socket systems (X8M-2)	15 IP addresses for the base configuration and: <ul style="list-style-type: none"> • One IP address per additional database server for administration • One IP address per additional Exadata Storage Server for administration • One IP address per additional database server for ILOM • One IP address per additional Exadata Storage Server for ILOM 	7 IP addresses for the base configuration and: <ul style="list-style-type: none"> • One IP address per additional database server for public access • One IP address per additional database server for Oracle RAC VIP addresses 	For each additional network, 2 IP addresses for the base configuration and: <ul style="list-style-type: none"> • One IP address per additional database server 	10 IP addresses for the base configuration and: <ul style="list-style-type: none"> • Two IP addresses per additional database server • Two IP addresses per additional Exadata Storage Server

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2, NET3, and so on) Requirements	RoCE Network Fabric Requirements
Full Rack configuration for 8-socket systems (X8M-8)	Minimum of 37 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 14 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 14 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 1 for each RoCE Network Fabric switch, and 1 for the Ethernet switch • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC 	Not applicable	44 IP addresses <ul style="list-style-type: none"> • 16 IP addresses, two per database server • 28 IP addresses, two per Exadata Storage Server

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2, NET3, and so on) Requirements	RoCE Network Fabric Requirements
Half Rack configuration for 8-socket systems (X8M-8)	Minimum of 15 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 3 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 3 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 1 for each RoCE Network Fabric switch, and 1 for the Ethernet switch • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC 	Not applicable	22 IP addresses <ul style="list-style-type: none"> • 16 IP addresses, two per database server • 6 IP addresses, two per Exadata Storage Server

1.9.3.2 IP Address Requirements for Oracle Exadata Database Machine X5 to X8 Systems with InfiniBand Network Fabric

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Full Rack configuration for 2-socket systems (X5-2 to X8-2)	Minimum of 49 IP addresses <ul style="list-style-type: none"> • 8 IP addresses for administration, one per database server • 14 IP addresses for administration, one per Exadata Storage Server • 8 IP addresses for ILOM, one per database server • 14 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 2 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 19 IP address <ul style="list-style-type: none"> • 8 IP addresses for public access, one per database server • 8 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC 	Minimum of 8 IP addresses for each additional network <ul style="list-style-type: none"> • 8 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	44 IP addresses <ul style="list-style-type: none"> • 16 IP addresses, two per database server • 28 IP addresses, two per Exadata Storage Server

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Half Rack configuration for 2-socket systems (X5-2 to X8-2)	Minimum of 27 IP addresses <ul style="list-style-type: none"> • 4 IP addresses for administration, one per database server • 7 IP addresses for administration, one per Exadata Storage Server • 4 IP addresses for ILOM, one per database server • 7 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 2 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 11 IP addresses <ul style="list-style-type: none"> • 4 IP addresses for public access, one per database server • 4 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	Minimum of 4 IP addresses for each additional network <ul style="list-style-type: none"> • 4 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	For X6-2 and X5-2 Half Rack: 22 IP addresses: <ul style="list-style-type: none"> • 8 IP addresses, two per database server • 14 IP addresses, two per Exadata Storage Server

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Quarter Rack or Eighth Rack configuration for 2-socket systems (X5-2 to X8-2)	Minimum of 15 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 3 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 3 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 2 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	Minimum of 2 IP addresses for each additional network <ul style="list-style-type: none"> • 2 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	For X5-2 and X6-2: 10 IP addresses: <ul style="list-style-type: none"> • 4 IP addresses, two per database server • 6 IP addresses, two per Exadata Storage Server

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Full Rack configuration for 8-socket systems (X5-8 to X8-8)	Minimum of 37 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 14 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 14 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 2 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC 		44 IP addresses <ul style="list-style-type: none"> • 16 IP addresses, two per database server • 28 IP addresses, two per Exadata Storage Server

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Half Rack configuration for 8-socket systems (X5-8 to X8-8)	Minimum of 15 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 3 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 3 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 2 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC 		22 IP addresses <ul style="list-style-type: none"> • 16 IP addresses, two per database server • 6 IP addresses, two per Exadata Storage Server

1.9.3.3 IP Address Requirements for Oracle Exadata Database Machine X4 Systems

Table 1-36 IP Address Requirements for Oracle Exadata Database Machine X4-2 and X4-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X4-2 Full Rack	Minimum of 49 IP addresses <ul style="list-style-type: none"> 8 IP addresses for administration, one per database server 14 IP addresses for administration, one per Exadata Storage Server 8 IP addresses for ILOM, one per database server 14 IP addresses for ILOM, one per Exadata Storage Server 3 IP addresses for switches: 2 for InfiniBand, and 1 for Ethernet 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 19 IP address <ul style="list-style-type: none"> 8 IP addresses for public access, one per database server 8 IP addresses for Oracle RAC VIP addresses, one per database server 3 IP addresses for SCAN addresses for Oracle RAC 	Minimum of 8 IP addresses for each additional network <ul style="list-style-type: none"> 8 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	44 IP addresses <ul style="list-style-type: none"> 16 IP addresses, two per database server 28 IP addresses, two per Exadata Storage Server

Table 1-36 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X4-2 and X4-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X4-2 Half Rack	Minimum of 27 IP addresses <ul style="list-style-type: none"> • 4 IP addresses for administration, one per database server • 7 IP addresses for administration, one per Exadata Storage Server • 4 IP addresses for ILOM, one per database server • 7 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 2 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 11 IP addresses <ul style="list-style-type: none"> • 4 IP addresses for public access, one per database server • 4 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	Minimum of 4 IP addresses for each additional network <ul style="list-style-type: none"> • 4 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	22 IP addresses <ul style="list-style-type: none"> • 8 IP addresses, two per database server • 14 IP addresses, two per Exadata Storage Server

Table 1-36 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X4-2 and X4-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X4-2 Quarter Rack and Oracle Exadata Database Machine X4-2 Eighth Rack	Minimum of 15 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 3 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 3 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 2 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	Minimum of 2 IP addresses for each additional network <ul style="list-style-type: none"> • 2 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	10 IP addresses <ul style="list-style-type: none"> • 4 IP addresses, two per database server • 6 IP addresses, two per Exadata Storage Server

Table 1-36 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X4-2 and X4-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X4-8 Full Rack	Minimum of 37 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 14 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 14 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches: 3 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	Minimum of 2 IP addresses for each additional network <ul style="list-style-type: none"> • 2 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	16 IP addresses <ul style="list-style-type: none"> • 2 IP addresses, one per database server • 14 IP addresses, one per Exadata Storage Server

1.9.3.4 IP Address Requirements for Oracle Exadata Database Machine X3 Systems

Table 1-37 IP Address Requirements for Oracle Exadata Database Machine X3-2 and Oracle Exadata Database Machine X3-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X3-2 Full Rack	Minimum of 50 IP addresses <ul style="list-style-type: none"> • 8 IP addresses for administration, one per database server • 14 IP addresses for administration, one per Exadata Storage Server • 8 IP addresses for ILOM, one per database server • 14 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches: 3 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 19 IP address <ul style="list-style-type: none"> • 8 IP addresses for public access, one per database server • 8 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC 	Minimum of 8 IP addresses for each additional network <ul style="list-style-type: none"> • 8 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	22 IP addresses <ul style="list-style-type: none"> • 8 IP addresses, one per database server • 14 IP addresses, one per Exadata Storage Server

Table 1-37 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X3-2 and Oracle Exadata Database Machine X3-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X3-2 Half Rack	Minimum of 28 IP addresses <ul style="list-style-type: none"> • 4 IP addresses for administration, one per database server • 7 IP addresses for administration, one per Exadata Storage Server • 4 IP addresses for ILOM, one per database server • 7 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches: 3 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 11 IP addresses <ul style="list-style-type: none"> • 4 IP addresses for public access, one per database server • 4 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	Minimum of 4 IP addresses for each additional network <ul style="list-style-type: none"> • 4 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	11 IP addresses <ul style="list-style-type: none"> • 4 IP addresses, one per database server • 7 IP addresses, one per Exadata Storage Server

Table 1-37 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X3-2 and Oracle Exadata Database Machine X3-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X3-2 Quarter Rack	Minimum of 15 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 3 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 3 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches: 2 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	Minimum of 2 IP addresses for each additional network <ul style="list-style-type: none"> • 2 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	5 IP addresses <ul style="list-style-type: none"> • 2 IP addresses, one per database server • 3 IP addresses, one per Exadata Storage Server

Table 1-37 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X3-2 and Oracle Exadata Database Machine X3-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers	Minimum of 38 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 14 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 14 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches: 3 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP address <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC 	Minimum of 2 IP addresses for each additional network <ul style="list-style-type: none"> • 2 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	36 IP addresses <ul style="list-style-type: none"> • 8 IP addresses, one per database server • 28 IP addresses, two per Exadata Storage Server

Table 1-37 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X3-2 and Oracle Exadata Database Machine X3-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X3-2 Servers or Exadata Storage Server with Sun Fire X4270 M2 Servers	Minimum of 38 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 14 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 14 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches: 3 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP address <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC 	Minimum of 2 IP addresses for each additional network <ul style="list-style-type: none"> • 2 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	22 IP addresses <ul style="list-style-type: none"> • 8 IP addresses, one per database server • 14 IP addresses, one per Exadata Storage Server

1.9.3.5 IP Address Requirements for Oracle Exadata Database Machine X2 Systems

Table 1-38 IP Address Requirements for Oracle Exadata Database Machine X2-2 and Oracle Exadata Database Machine X2-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X2-2 Full Rack	Minimum of 51 IP addresses <ul style="list-style-type: none"> 8 IP addresses for administration, one per database server 14 IP addresses for administration, one per Exadata Storage Server 8 IP addresses for ILOM, one per database server 14 IP addresses for ILOM, one per Exadata Storage Server 5 IP addresses for switches: 3 for InfiniBand, 1 for Ethernet, and 1 for KVM 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 19 IP address <ul style="list-style-type: none"> 8 IP addresses for public access, one per database server 8 IP addresses for Oracle RAC VIP addresses, one per database server 3 IP addresses for SCAN addresses for Oracle RAC 	Minimum of 8 IP addresses for each additional network <ul style="list-style-type: none"> 8 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	22 IP addresses <ul style="list-style-type: none"> 8 IP addresses, one per database server 14 IP addresses, one per Exadata Storage Server

Table 1-38 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X2-2 and Oracle Exadata Database Machine X2-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X2-2 Half Rack	<p>Minimum of 29 IP addresses (28 IP addresses are needed for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers))</p> <ul style="list-style-type: none"> • 4 IP addresses for administration, one per database server • 7 IP addresses for administration, one per Exadata Storage Server • 4 IP addresses for ILOM, one per database server • 7 IP addresses for ILOM, one per Exadata Storage Server • 5 IP addresses for switches: 3 for InfiniBand, 1 for Ethernet, and 1 for KVM <p>4 IP addresses are needed for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers): 2 for InfiniBand, 1 for Ethernet, and 1 for KVM</p> <ul style="list-style-type: none"> • 2 IP addresses for monitoring electric current of the PDUs 	<p>Minimum of 11 IP addresses</p> <ul style="list-style-type: none"> • 4 IP addresses for public access, one per database server • 4 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	<p>Minimum of 4 IP addresses for each additional network</p> <ul style="list-style-type: none"> • 4 IP addresses for the database servers, one per database server <p>If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.</p>	<p>11 IP addresses</p> <ul style="list-style-type: none"> • 4 IP addresses, one per database server • 7 IP addresses, one per Exadata Storage Server

Table 1-38 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X2-2 and Oracle Exadata Database Machine X2-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X2-2 Quarter Rack	Minimum of 16 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 3 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 3 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches: 2 for InfiniBand, 1 for Ethernet, and 1 for KVM • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC. 	Minimum of 2 IP addresses for each additional network <ul style="list-style-type: none"> • 2 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	5 IP addresses <ul style="list-style-type: none"> • 2 IP addresses, one per database server • 3 IP addresses, one per Exadata Storage Server

Table 1-38 (Cont.) IP Address Requirements for Oracle Exadata Database Machine X2-2 and Oracle Exadata Database Machine X2-8 Full Rack

Type of Oracle Exadata Database Machine	Management Network Requirements	Client Access Network Requirements	Additional Networks (NET2 and NET3) Requirements	InfiniBand Private Network Requirements
Oracle Exadata Database Machine X2-8 Full Rack	Minimum of 38 IP addresses <ul style="list-style-type: none"> • 2 IP addresses for administration, one per database server • 14 IP addresses for administration, one per Exadata Storage Server • 2 IP addresses for ILOM, one per database server • 14 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches: 3 for InfiniBand, and 1 for Ethernet • 2 IP addresses for monitoring electric current of the PDUs 	Minimum of 7 IP address <ul style="list-style-type: none"> • 2 IP addresses for public access, one per database server • 2 IP addresses for Oracle RAC VIP addresses, one per database server • 3 IP addresses for SCAN addresses for Oracle RAC 	Minimum of 2 IP addresses for each additional network <ul style="list-style-type: none"> • 2 IP addresses for the database servers, one per database server If using a bonded network configuration for the client access network, then NET2 cannot be used as an additional network.	22 IP addresses <ul style="list-style-type: none"> • 8 IP addresses, one per database server • 14 IP addresses, one per Exadata Storage Server

1.9.4 Calculating the IP Address Requirements when Using Oracle VMs

The number of IP addresses depends on the physical deployment, as well as the number of virtual user domains.

The following equations show how to calculate the number of IP addresses required for the physical deployment and Oracle VM deployment.

Physical deployment with a single cluster:

```

Number of management network IP addresses =
  (number of databases servers)*2 +      \
  (number of storage servers)*2 +        \
  (number of PDUs) +                      \
  (number of IB switches) +              \
  (number of Ethernet switches)

```

Number of client access network IP addresses =
(number of database servers)*2 + 3

Number of InfiniBand network IP addresses =
(number of database servers)*2 + \\
(number of storage servers)*2

Number of additional network IP addresses =
(number of database servers) * (number of additional networks)

Oracle VM deployment:

Number of management network IP addresses =
(number of databases servers)*2 + \\
(number of Oracle VM user domains) + \\
(number of storage servers)*2 + \\
(number of PDUs) + \\
(number of IB switches) + \\
(number of Ethernet switches) \\
\\

Number of client access network IP addresses =
(number of Oracle VM user domains)*2 + \\
(number of Oracle RAC VM clusters)*3

Number of InfiniBand network IP addresses =
(number of Oracle VM user domains)*2 + \\
(number of storage servers)*2

Number of additional network IP addresses =
(number of database servers) * (number of Oracle VM user domains)

The following table shows the network requirements by component. The IP address numbers are per component, not rack.

Table 1-39 IP Address Requirements by Component

Component	Management Network	Client Access Network	Additional Networks (NET2 and NET3)	RDMA Network Fabric
Database server	Physical deployment <ul style="list-style-type: none"> 1 IP address for administration 1 IP address for ILOM Oracle VM deployment <ul style="list-style-type: none"> 1 IP address for administration of management domain 1 IP address for ILOM 1 IP address for each user domain for administration 	Physical deployment <ul style="list-style-type: none"> 1 IP address for public access 1 IP address for Oracle RAC VIP 3 IP addresses for each Oracle RAC cluster SCAN addresses Oracle VM deployment <ul style="list-style-type: none"> 1 IP address for each user domain public access 1 IP address for each user domain Oracle RAC VIP 3 IP addresses for each Oracle RAC VM SCAN addresses 	Physical deployment <ul style="list-style-type: none"> 1 IP address for each additional configured network Oracle VM deployment <ul style="list-style-type: none"> 1 IP address for each user domain for each additional configured network 	Physical deployment <ul style="list-style-type: none"> 2 IP addresses when using active-active bonding, or 1 IP address when not using active-active bonding Oracle VM deployment <ul style="list-style-type: none"> 2 IP addresses for each user domain when using active-active bonding, or 1 IP address for each user domain when not using active-active bonding
Exadata Storage Server	<ul style="list-style-type: none"> 1 IP address for administration 1 IP address for ILOM 	None	None	2 IP addresses when using active-active bonding, or 1 IP address when not using active-active bonding.
Switch for RoCE Network Fabric (X8M systems)	1 IP address	None	None	None
Switch for InfiniBand Network Fabric	1 IP address	None	None	None
Ethernet switch	1 IP address	None	None	None
Spine switch	1 IP address	None	None	None
PDU	1 IP address	None	None	None

The following table shows an example of the IP addresses requirements when using Oracle VMs in Oracle Exadata Database Machine X5 or later models.

Table 1-40 Examples of IP Address Requirements when Using Oracle VMs

Deployment	IP Address Requirements
Oracle Exadata Database Machine Quarter Rack with the following configuration: <ul style="list-style-type: none"> • Oracle VM deployment with two 2-node Oracle RAC VM clusters (2 clusters and 4 user domains total) • No additional networks • Active-active bonding on the RDMA Network Fabric 	<ul style="list-style-type: none"> • 19 IP addresses for the management network • 14 IP addresses for the client access network • No IP addresses for the additional networks • 14 IP addresses for the RDMA Network Fabric
Oracle Exadata Database Machine Full Rack with the following configuration: <ul style="list-style-type: none"> • Oracle VM deployment with one 8-node Oracle RAC VM cluster, and five 2-node Oracle RAC VM clusters (6 clusters and 18 user domains total) • No additional networks • Active-active bonding on the RDMA Network Fabric 	<ul style="list-style-type: none"> • 67 IP addresses for the management network • 54 IP addresses for the client access network • No IP addresses for the additional networks • 64 IP addresses for the RDMA Network Fabric
Oracle Exadata Database Machine Elastic Configuration with the following configuration: <ul style="list-style-type: none"> • 3 database servers, 9 Exadata Storage Servers, and 2 RDMA Network Fabric switches (InfiniBand or RoCE). • Oracle VM deployment with two 3-node Oracle RAC VM clusters (2 clusters and 6 user domains total) • Dedicated backup network • Active-active bonding on the RDMA Network Fabric 	<ul style="list-style-type: none"> • 37 IP addresses for the management network • 18 IP addresses for the client access network • 6 IP addresses for the additional networks • 30 IP addresses for the RDMA Network Fabric

1.10 Network Connection and IP Address Requirements for Oracle Exadata Storage Expansion Rack

Prior to installation, network cables must be run from your existing network infrastructure to the installation site.

The requirements to connect Oracle Exadata Storage Expansion Rack to Oracle Exadata Database Machine are as follows:

- Management network connection requirements for Oracle Exadata Storage Expansion Rack X4-2 and later:
 - 1 Ethernet connection from the ILOM to each Exadata Storage Server. For example, an Oracle Exadata Storage Expansion Full Rack has 18 connections, and an Oracle Exadata Storage Expansion Quarter Rack has 4 connections.
 - 1 Ethernet connection for each RDMA Network Fabric switch.

- 1 Ethernet connection for each Exadata Storage Server. For example, an Oracle Exadata Storage Expansion Full Rack has 18 connections, and an Oracle Exadata Storage Expansion Quarter Rack has 4 connections.
- 1 Ethernet connection for the Ethernet switch in the rack to the existing management network.
- 1 Ethernet connection for each of the two PDUs. Network connectivity to the PDUs is only required if the electric current is monitored remotely.
- Management network connection requirements for Oracle Exadata Storage Expansion Rack X3-2
 - 1 Ethernet connection from the ILOM to each Exadata Storage Server. For Oracle Exadata Storage Expansion Full Rack, there are 18 connections, for Oracle Exadata Storage Expansion Half Rack there are 9 connections, and for Oracle Exadata Storage Expansion Quarter Rack there are 4 connections.
 - 1 Ethernet connection for each Sun Datacenter InfiniBand Switch 36.
 - 1 Ethernet connection for each Exadata Storage Server. For Oracle Exadata Storage Expansion Full Rack, there are 18 connections, for Oracle Exadata Storage Expansion Half Rack there are 9 connections, and for Oracle Exadata Storage Expansion Quarter Rack there are 4 connections.
 - 1 Ethernet connection for the Ethernet switch in the rack to the existing management network.
 - 1 Ethernet connection for each of the two PDUs. Network connectivity to the PDUs is only required if the electric current is monitored remotely.
- Management network connection requirements for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers
 - 1 Ethernet connection from the ILOM to each Exadata Storage Server. For Oracle Exadata Storage Expansion Full Rack, there are 18 connections, for Oracle Exadata Storage Expansion Half Rack there are 9 connections, and for Oracle Exadata Storage Expansion Quarter Rack there are 4 connections.
 - 1 Ethernet connection for each Sun Datacenter InfiniBand Switch 36.
 - 1 Ethernet connection for each Exadata Storage Server. For Oracle Exadata Storage Expansion Full Rack, there are 18 connections, for Oracle Exadata Storage Expansion Half Rack there are 9 connections, and for Oracle Exadata Storage Expansion Quarter Rack there are 4 connections.
 - 1 Ethernet connection for the Ethernet switch in the rack to the existing management network.
 - 1 Ethernet connection for the KVM switch in the rack to the existing management network. This connection is only for Oracle Exadata Database Machine X2-2 racks.
 - 1 Ethernet connection for each of the two PDUs. Network connectivity to the PDUs is only required if the electric current is monitored remotely.

The following table lists the IP address requirements for Oracle Exadata Storage Expansion Racks:

Table 1-41 IP Address Requirements for Oracle Exadata Storage Expansion Rack

Type of Oracle Exadata Storage Expansion Rack	Network Requirements
Oracle Exadata Storage Expansion X4-2 Full Rack and later, including X8M systems	<p>Up to a total of 78 IP addresses (42 for the management network, and up to 36 for the RDMA Network Fabric)</p> <ul style="list-style-type: none"> • 18 IP addresses for administration, one per Exadata Storage Server • 18 IP addresses for ILOM, one per Exadata Storage Server • Up to 4 IP addresses for switches (2 for RDMA Network Fabric, 1 for spine switch, and 1 for Ethernet) • 2 IP addresses for monitoring electric current of the PDUs • Up to 36 IP addresses for the RDMA Network Fabric
Oracle Exadata Storage Expansion X4-2 Half Rack and later, including X8M systems	<p>Up to a total of 42 IP addresses (24 for the management network, and up to 18 for the RDMA Network Fabric)</p> <ul style="list-style-type: none"> • 9 IP addresses for administration, one per Exadata Storage Server • 9 IP addresses for ILOM, one per Exadata Storage Server • Up to 4 IP addresses for switches (2 for RDMA Network Fabric, 1 for spine switch, and 1 for Ethernet) • 2 IP addresses for monitoring electric current of the PDUs • Up to 18 IP addresses for the RDMA Network Fabric
Oracle Exadata Storage Expansion X4-2 Quarter Rack and later, including X8M systems	<p>Up to a total of 21 IP addresses (13 for the management network, and up to 8 for the RDMA Network Fabric)</p> <ul style="list-style-type: none"> • 4 IP addresses for administration, one per Exadata Storage Server • 4 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches (2 for RDMA Network Fabric, 1 for spine switch, and 1 for Ethernet) • 2 IP addresses for monitoring electric current of the PDUs • Up to 8 IP addresses for the RDMA Network Fabric

Table 1-41 (Cont.) IP Address Requirements for Oracle Exadata Storage Expansion Rack

Type of Oracle Exadata Storage Expansion Rack	Network Requirements
Oracle Exadata Storage Expansion X3-2 Full Rack	<p>Total of 60 IP addresses (42 for the management network, and 18 for the private InfiniBand network)</p> <ul style="list-style-type: none"> • 18 IP addresses for administration, one per Exadata Storage Server • 18 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches (3 for InfiniBand, and 1 for Ethernet) • 2 IP addresses for monitoring electric current of the PDUs • 18 IP addresses for the private InfiniBand network
Oracle Exadata Storage Expansion X3-2 Half Rack	<p>Total of 33 IP addresses (24 for the management network, and 9 for the private InfiniBand network)</p> <ul style="list-style-type: none"> • 9 IP addresses for administration, one per Exadata Storage Server • 9 IP addresses for ILOM, one per Exadata Storage Server • 4 IP addresses for switches (3 for InfiniBand, and 1 for Ethernet) • 2 IP addresses for monitoring electric current of the PDUs • 9 IP addresses for the private InfiniBand network
Oracle Exadata Storage Expansion X3-2 Quarter Rack	<p>Total of 17 IP addresses (13 for the management network, and 4 for the InfiniBand network)</p> <ul style="list-style-type: none"> • 4 IP addresses for administration, one per Exadata Storage Server • 4 IP addresses for ILOM, one per Exadata Storage Server • 3 IP addresses for switches (2 for InfiniBand, and 1 for Ethernet) • 2 IP addresses for monitoring electric current of the PDUs • 4 IP addresses for the private InfiniBand network

1.11 Ensuring That the Site is Ready

Before Oracle Exadata Rack is delivered to the site, the following tasks must be performed to ensure the site is ready:

Task 1 Review Site Requirements

Review the site requirements in this chapter and use the checklists in [Site Checklists](#) to understand the requirements for Oracle Exadata Rack.

Task 2 Run the Oracle Exadata Deployment Assistant

Run the Oracle Exadata Deployment Assistant to create the necessary files.

Task 3 Configure the Network

Configure the network using the files from Oracle Exadata Deployment Assistant. This includes registering the networks in the Domain Name System (DNS), assigning IP addresses, and configuring data center switches and firewalls.

Task 4 Prepare the Site Based on Requirements

Prepare the site based on the requirements, such as, install the network cables and power supplies, prior to the arrival of Oracle Exadata Rack:

1. Review the safety guidelines.
2. Note problems or peculiarities at the site that require special equipment. For example, ensure the doors are tall enough and wide enough for Oracle Exadata Rack.
3. Verify that the installation site flooring has a strength rating to withstand the combined weight of Oracle Exadata Rack and any other installed equipment.
4. Install network cables for Oracle Exadata Rack.
5. Install all necessary electrical equipment, and ensure that sufficient power is provided for Oracle Exadata Rack.
6. Ensure that the installation site provides adequate air conditioning.
7. Operate the air conditioning system for 48 hours to bring the room temperature to the appropriate level.
8. Clean and vacuum the area thoroughly in preparation for installation.

2

Understanding the Network Requirements for Oracle Exadata Database Machine

Review the network requirements for Oracle Exadata Database Machine before installing or configuring the hardware.

Note:

For ease of reading, the name "Oracle Exadata Rack" is used when information refers to both Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack.

- [Overview of Network Requirements](#)
In addition to the database and storage servers, Oracle Exadata Database Machine includes equipment to connect the system to your network. The network connections allow clients to connect to the database servers and also enables remote system administration.
- [Network Channel Bonding Support](#)
- [Network Partitioning on Oracle Exadata Database Machine](#)
- [Configuring a Separate Network for ILOM](#)
When configuring or re-imaging an Oracle Exadata Rack, you can use Oracle Exadata Deployment Assistant (OEDA) to configure a separate network for Integrated Lights Out Manager (ILOM).
- [Default IP Addresses](#)
Starting with Oracle Exadata System Software release 12.1.2.1.0, the NET0 IP addresses are assigned dynamically by the elastic configuration procedure during the first start of the system.
- [Port Assignments When Using a Firewall](#)

2.1 Overview of Network Requirements

In addition to the database and storage servers, Oracle Exadata Database Machine includes equipment to connect the system to your network. The network connections allow clients to connect to the database servers and also enables remote system administration.

Use the information in this section in conjunction with Oracle Exadata Deployment Assistant (OEDA) to configure your Oracle Exadata Database Machine environment.

To deploy Oracle Exadata Database Machine ensure that you meet the minimum network requirements. Oracle Exadata Database Machine requires a minimum of three networks, and there are interfaces available for additional networks. Each network must be on a distinct and separate subnet from the others. Oracle

recommends one common additional network for backup, which you can configure with OEDA. The network descriptions are as follows:

- **Management network:** This required network connects to your existing management network, and is used for administrative work for all components of Oracle Exadata Database Machine. It connects the servers, Integrated Lights Out Manager (ILOM) interfaces, and RDMA Network Fabric switches to the Ethernet switch in the rack. One uplink is required from the Ethernet switch to your management network.

 **Notes:**

- Separate uplinks to the management network are also recommended for remote monitoring of each power distribution unit (PDU). This configuration enables you to easily differentiate between system outages caused by PDU failure as opposed to failure of the Ethernet switch.
- A properly secured configuration requires full isolation of the management network from all other networks.
- Starting with Oracle Exadata System Software release 19.1.0, the Exadata database and storage servers can be on a separate network from the ILOM interfaces.

Each database server and storage server has two network interfaces for management. One provides management access to the operating system through the embedded NET0 Ethernet port on Oracle Exadata Database Machine Two-Socket systems, or the NEM0/NET1 Ethernet port on Oracle Exadata Database Machine Eight-Socket systems. The other network interface provides access to the ILOM through the NET MGT Ethernet port. Oracle Exadata Database Machines are delivered with the NET MGT and NET0 or NEM0/NET0 ports connected to the Ethernet switch in the rack. The NET0 or NEM0/NET0 port on the database servers should not be used for client or application network traffic. Cabling or configuration changes to these interfaces on Oracle Exadata Storage Servers is not permitted.

- **Client access network:** This required network connects the database servers to your existing client network and is used for client access to the database servers. Applications access the database through this network using Single Client Access Name (SCAN) and Oracle RAC Virtual IP (VIP) addresses. Database servers support channel bonding to provide higher bandwidth or availability for client connections to the database. Oracle recommends channel bonding for the client access network.

The following list describes the channel bonding configurations:

- Oracle Exadata Database Machine Two-Socket systems use two of the embedded ports.

On Oracle Exadata Database Machine X7-2, X8-2, and X8M-2 systems, the two 10 GbE copper ports or two 10/25GbE fiber ports are used. On earlier model systems, the embedded NET1 and NET2 ports are used.
- Oracle Exadata Database Machine Eight-Socket systems, X4-8 or later, use two ports, with each port on a different PCIe card.

- Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack use ports from each Network Express Module (NEM).

For the Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack database servers, using ports on the same Network Express Modules, such as NEM0/NET1 and NEM0/NET2, works but does not provide for high availability.

 **Note:**

If your system will use more than 10 VIP addresses in the cluster and you have VLAN configured for the Oracle Clusterware client network, then you must use 3 digit VLAN ids. Do not use 4 digit VLAN ids because the VLAN name can exceed the 15 character operating system interface name limit.

- **Additional networks (optional):** Database servers can connect to additional existing networks using the available open ports not used by the management and client networks.
- **RDMA Network Fabric (Private network):** This network connects the database servers and storage servers. Oracle Database uses this network for Oracle RAC cluster interconnect traffic and for accessing data on Oracle Exadata Storage Servers. This non-routable network is fully contained in Oracle Exadata Database Machine, and does not connect to your existing network. This network is automatically configured during installation.

Starting with Oracle Exadata System Software release 19.3 and Oracle Exadata Database Machine X8M, RDMA over Converged Ethernet (RoCE) is used for the RDMA Network Fabric. Previously, InfiniBand is used for the RDMA Network Fabric. RoCE Network Fabric uses different switches and cables from those used by InfiniBand Network Fabric.

 **Note:**

All networks must be on distinct and separate subnets from each other. For better security, the management network should be isolated from all other networks.

Oracle Exadata Database Machine database server networking requirements are identical across all models with the following exceptions:

- Oracle Exadata Database Machine X7-2, X8-2, and X8M-2 embedded ports are 10 GbE (copper) or 10/25 GbE (fiber).
- Oracle Exadata Database Machine X6-2, Oracle Exadata Database Machine X5-2, Oracle Exadata Database Machine X4-2, and Oracle Exadata Database Machine X3-2 embedded ports are 1 GbE/10 GbE.
- Oracle Exadata Database Machine X2-2 embedded ports are 1 GbE.
- Only Oracle Exadata Database Machine X2-2 has a keyboard, video and mouse (KVM) switch.

Additionally, for all database server network connections using SFP28 ports:

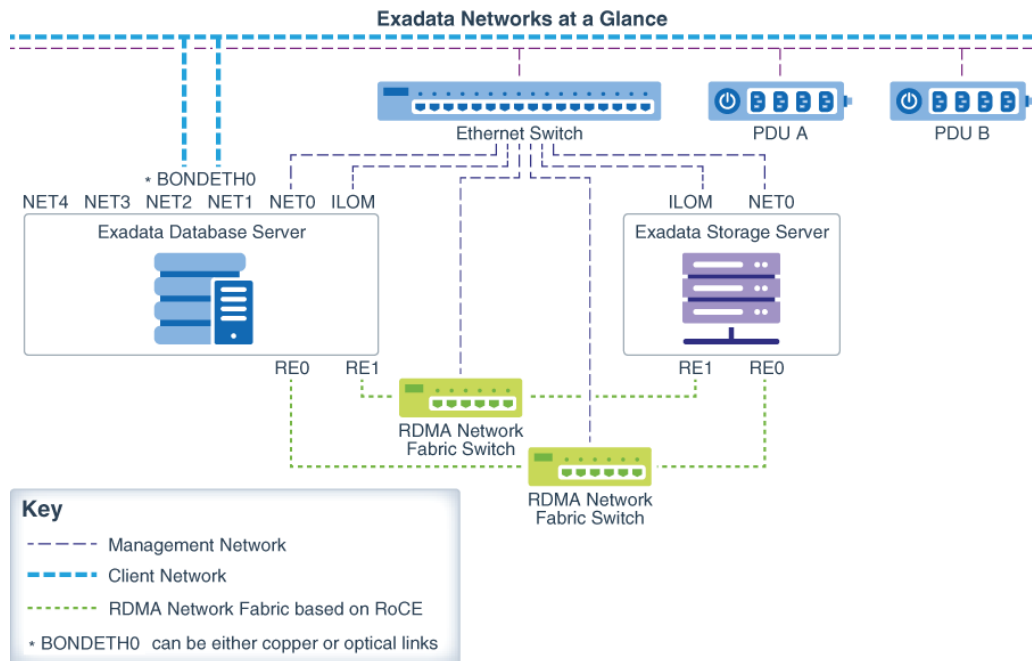
- SFP+ transceivers are required to support 10 Gigabit Ethernet (10 GbE).
- SFP28 transceivers are required to support 25 Gigabit Ethernet (25 GbE).

At start time on an 8-socket system, each database server generates a `cellaffinity.ora` configuration file. The `cellaffinity.ora` file resides in the `/etc/oracle/cell/network-config` directory, and must be readable by the Oracle Database software.

The `cellaffinity.ora` file contains a mapping between the NUMA node numbers and the IP address of the network interface card closest to each node. Oracle Database uses the file to select the closest network interface card when communicating with the storage servers, thereby optimizing performance. This file is only generated and used on an 8-socket system. On a 2-socket system, there is no performance to be gained in this manner, and no `cellaffinity.ora` file. The file is not intended to be directly edited with a text editor.

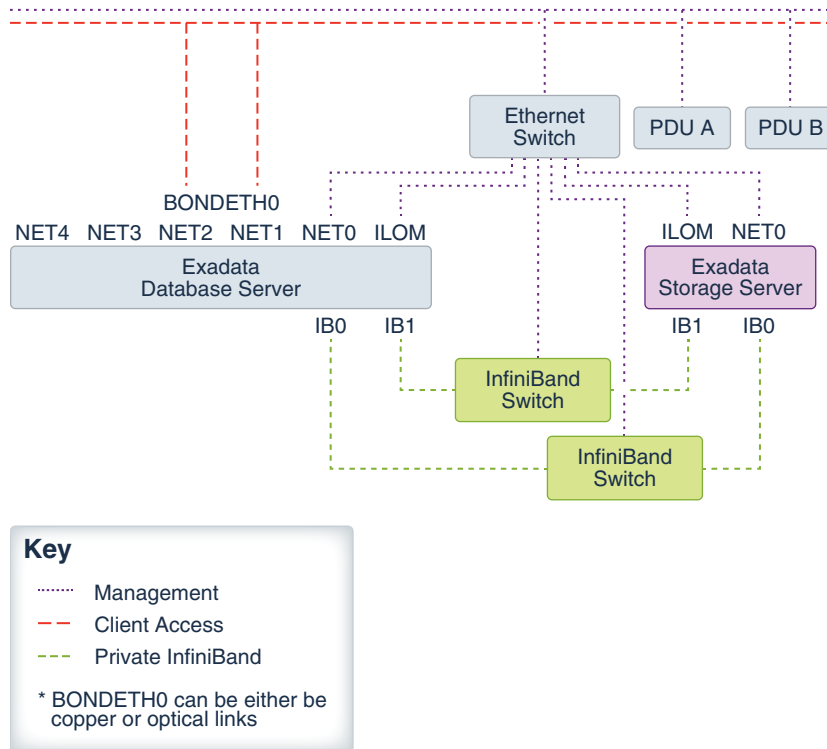
The following image shows the default network diagram for Oracle Exadata Database Machine X8M with RoCE Network Fabric. Use the files from OEDA to verify the cable connections when cabling Oracle Exadata Database Machine to the existing network.

Figure 2-1 Network Diagram for Oracle Exadata Database Machine X8M



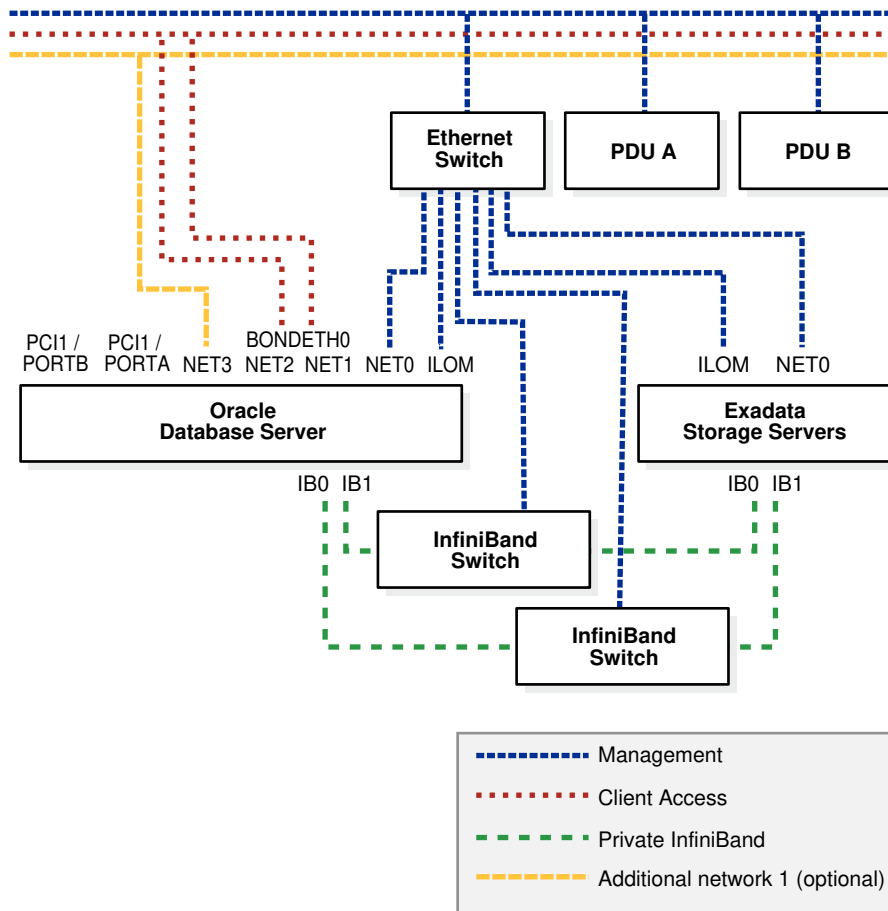
The following image shows the default network diagram for Oracle Exadata Database Machine X7-2 and X8-2 with InfiniBand Network Fabric. Use the files from OEDA to verify the cable connections when cabling Oracle Exadata Database Machine to the existing network.

Figure 2-2 Network Diagram for Oracle Exadata Database Machine X7-2 and X8-2



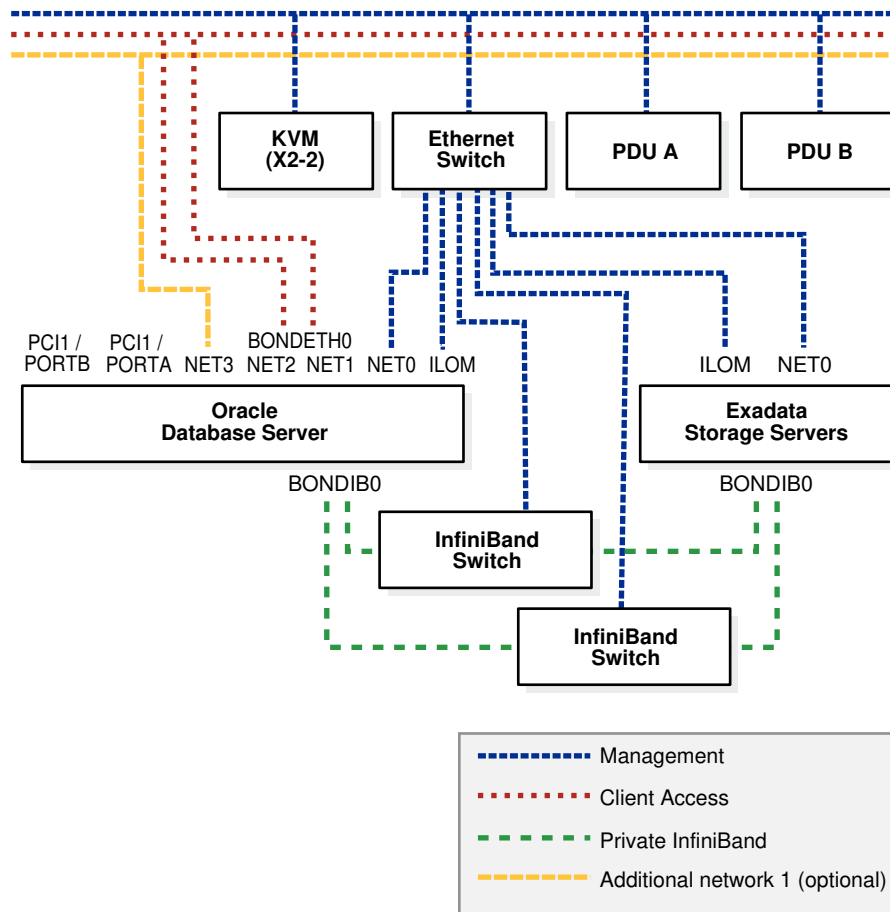
The following image shows the default network diagram for Oracle Exadata Database Machine X4-2, X5-2, and X6-2 models. Use the files from OEDA to verify the cable connections when cabling Oracle Exadata Database Machine to the existing network.

Figure 2-3 Network Diagram for Oracle Exadata Database Machine X4-2, X5-2, and X6-2 Models



The following image shows the default network diagram for Oracle Exadata Database Machine X3-2 and Oracle Exadata Database Machine X2-2. Use the files from OEDA to verify the cable connections when cabling Oracle Exadata Database Machine to the existing network.

Figure 2-4 Network Diagram for Oracle Exadata Database Machine X3-2 and Oracle Exadata Database Machine X2-2



Notes:

- Oracle Exadata Database Machine X3-8 Full Rack, and Oracle Exadata Database Machine X2-8 Full Rack have BONDIB0, BONDIB1, and so on for the RDMA Network Fabric.
- The keyboard, video, and mouse (KVM) switch is only available in Oracle Exadata Database Machine X2-2 racks.
- Additional configuration, such as defining multiple virtual local area networks (VLANs) or enabling routing, may be required for the switch to operate properly in your environment, and is not included in the installation service. If additional configuration is needed, then the network administrator performs the necessary configuration during Oracle Exadata Database Machine installation.

Related Topics

- [Using Oracle Exadata Deployment Assistant](#)

- [Hardware Components of Oracle Exadata Database Machine](#)

2.2 Network Channel Bonding Support

The network port on a database server used for the client access network can be bonded with another network on the database server to provide higher bandwidth or availability for client connections to the database. In a bonded network configuration, Oracle RAC SCAN and VIP addresses are defined on the BONDETH0 interface on database servers.

- [Bonded Network Configurations](#)
A bonded network configuration is used when there is a need to aggregate multiple network interfaces together for client connections to the database.
- [Non-bonded Network Configuration](#)
A non-bonded network configuration is used when there is no need or ability to implement network bonding.
- [Additional Networks](#)

2.2.1 Bonded Network Configurations

A bonded network configuration is used when there is a need to aggregate multiple network interfaces together for client connections to the database.

The bonded network configuration is the default configuration.

The following sections describe the bonded network configuration for Oracle Exadata Database Machine:

- [Bonded Network Configuration for Oracle Exadata Database Machine Two-Socket Systems](#)
- [Bonded Network Configuration for Oracle Exadata Database Machine X7-8, X8-8, and X8M-8](#)
- [Bonded Network Configuration for Oracle Exadata Database Machine X4-8, X5-8, and X6-8](#)
- [Bonded Network Configuration for Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack](#)

2.2.1.1 Bonded Network Configuration for Oracle Exadata Database Machine Two-Socket Systems

When using a bonded network configuration for database client access on Oracle Exadata Database Machine two-socket systems, the database servers are configured as follows:

- Port NET0: Provides access to the operating system using the management network through the `eth0` interface.
- Port NET1: Provides access to the database using the client access network, and is bonded with NET2 port to create the `bondeth0` interface.
- Port NET2: Provides access to the database using the client access network, and is bonded with NET1 port to create the `bondeth0` interface.

- Port NET MGT: Provides access to Integrated Lights Out Manager (iLOM) using the management network.

When channel bonding is configured for the client access network during initial configuration, the Linux bonding module is configured for active-backup mode (`mode=1`). If a different bonding policy is preferred, then you may reconfigure the bonding module after initial configuration. For configuration details, refer to the "Linux Ethernet Bonding Driver HOWTO" at <https://www.kernel.org/doc/Documentation/networking/bonding.txt>.

The following table shows the port connections for Oracle Exadata Database Machine two-socket systems database servers when using channel bonding. The client access network and the backup network must use different ports, in that one network must use the embedded GbE ports and the other network must use the PCIe ports.

Table 2-1 Channel Bonding Connections for Oracle Exadata Database Machine X7-2, X8-2, and X8M-2

Network	Using 1 GbE Ports	Using 10/25 GbE Ports
Management (1 GbE NET0) The administration network connection uses the 1 GbE port even when using 10 GbE ports for the other networks.	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch 	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch
Client (<code>bondeth0</code>)	N/A	<ul style="list-style-type: none"> • Ports: 10/25 GbE PCI Ports 1 and 2 • Connect to backup 10/25 GbE switches
Backup (<code>bondeth1</code>)	N/A	<ul style="list-style-type: none"> • Ports: 10/25 GbE NET2 and NET3 • Connect to backup 10/25 GbE switches

Table 2-2 Channel Bonding Connections for Other Oracle Exadata Database Machine Two-Socket Systems

Network	Using 1 GbE Ports	Using 10 GbE Ports
Management (1 GbE NET0) The administration network connection uses the 1 GbE port even when using 10 GbE ports for the other networks.	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch 	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch
Client (<code>bondeth0</code>)	<ul style="list-style-type: none"> • Ports: 1 GbE NET1 and 1 GbE NET2 • Connect to client access 1 GbE switches 	<ul style="list-style-type: none"> • Ports: 10 GbE PCI1/PORTA and 10 GbE PCI1/PORTB • Connect to client access 10 GbE switches
Backup (<code>bondeth1</code>)	<ul style="list-style-type: none"> • Ports: 10 GbE PCI1/PORTA and 10 GbE PCI1/PORTB • Connect to backup 10 GbE switches 	<ul style="list-style-type: none"> • Ports: 1 GbE NET2 and 1 GbE NET3 • Connect to backup 1 GbE switches



Note:

When using a bonded network configuration, you must provide a network switch capable of supporting the chosen bonding mode. For example, if mode 4 is configured (IEEE 802.3ad Dynamic link aggregation), then you must supply and configure the network switch capable of supporting this bonding mode. Requirements for specific bonding policies are documented in the "Linux Ethernet Bonding Driver HOWTO" at <https://www.kernel.org/doc/Documentation/networking/bonding.txt>

2.2.1.2 Bonded Network Configuration for Oracle Exadata Database Machine X7-8, X8-8, and X8M-8

When using a bonded network configuration for database client access on Oracle Exadata Database Machine X7-8, X8-8, and X8M-8, the database servers are configured as follows:

- Port NET0: Provides access to the operating system using the management network through the `eth0` interface.
- P1 at 3 PCIe 4: Provides access to the database using the client access network, and is bonded with port P1 at 13 PCIe 12, to create the `bondeth0` interface.
- P2 at 13 PCIe 6: Provides access to the database using the client access network, and is bonded with port P2 at 3 PCIe 14 to create the `bondeth1` interface.
- Port NET MGT: Provides access to Integrated Lights Out Manager (ILOM) using the management network.

The following table shows the port connections for Oracle Exadata Database Machine X7-8, X8-8, and X8M-8 database servers when using channel bonding.

Table 2-3 Channel Bonding Connections for Oracle Exadata Database Machine X7-8, X8-8, and X8M-8

Network	Using 1 GbE Ports	Using 25 GbE Ports
Management (NET0) The administration network connection uses the 1 GbE port even when using 10 GbE ports for the other networks.	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch 	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch
Client (<code>bondeth0</code>)	<ul style="list-style-type: none"> • SMOD0 and SMOD1's Net2 • Connect to client access 1 GbE switches 	<ul style="list-style-type: none"> • Port1 in PCIe 4 and 12 cards • Connect to client access 25 GbE switches
Backup (<code>bondeth1</code>)	<ul style="list-style-type: none"> • SMOD0 and SMOD1's Net3 • Connect to backup 1 GbE switches 	<ul style="list-style-type: none"> • Port1 in PCIe 6 and 14 cards • Connect to backup 25 GbE switches

2.2.1.3 Bonded Network Configuration for Oracle Exadata Database Machine X4-8, X5-8, and X6-8

When using a bonded network configuration for database client access on Oracle Exadata Database Machine X4-8, X5-8, and X6-8, the database servers are configured as follows:

- Port NET0: Provides access to the operating system using the management network through the ETH0 interface.
- P0 at 3 PCIe3: Provides access to the database using the client access network, and is bonded with port P0 at 13 PCIe3, to create the BONDETH0 interface.
- P0 at 13 PCIe3: Provides access to the database using the client access network, and is bonded with port P0 at 3 PCIe3 to create the BONDETH1 interface.
- Port NET MGT: Provides access to ILOM using the management network.

The following table shows the port connections for Oracle Exadata Database Machine X4-8, X5-8, and X6-8 database servers when using channel bonding.

Table 2-4 Channel Bonding Connections for Oracle Exadata Database Machine X4-8, X5-8, and X6-8

Network	Using 1 GbE Ports	Using 10 GbE Ports
Management (NET0) The administration network connection uses the 1 GbE port even when using 10 GbE ports for the other networks.	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch 	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch
Client (BONDETH0)	<ul style="list-style-type: none"> • 1 GbE Ports: P0 at 3 PCIe3, and at 13 PCIe3 • Connect to client access 1 GbE switches 	<ul style="list-style-type: none"> • 10 GbE Ports: 2 PCIe3 PORTA, and 4 PCIe 3 PORTA • Connect to client access 10 GbE switches
Backup (BONDETH1)	<ul style="list-style-type: none"> • 10 GbE ports: 2 PCIe3 PORTA, and 4 PCIe3 PORTA • Connect to backup 10 GbE switches 	<ul style="list-style-type: none"> • 1 GbE ports: P0 at 3 PCIe3, and at 13 PCIe3 • Connect to backup 1 GbE switches

2.2.1.4 Bonded Network Configuration for Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack

When using a bonded network configuration for database client access on Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack, the database servers are configured as follows:

- Port NEM0/NET0: Provides access to the operating system using the management network through the ETH0 interface.
- Port NEM0/NET1: Provides access to the database using the client access network, and is bonded with port NEM1/NET3 to create the BONDETH0 interface.

- Port NEM1/NET3: Provides access to the database using the client access network, and is bonded with port NEM0/NET1 to create the BONDETH0 interface.
- Port NET MGT: Provides access to ILOM using the management network.

Table 2-5 shows the port connections for Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack database servers when using channel bonding. The NEM0 ports are on the left, and the NEM1 ports are on the right.

Table 2-5 Channel Bonding Connections for Oracle Exadata Database Machine X2-8 and X3-8 Full Rack

Network	Using 1 GbE Ports	Using 10 GbE Ports
Management (NET0)	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch 	<ul style="list-style-type: none"> • Port: 1 GbE NET0 • Connect to Cisco switch
Client (BONDETH0)	<ul style="list-style-type: none"> • Ports: 1 GbE NET1 and 1 GbE NET2 • Connect to client access 1 GbE switches 	<ul style="list-style-type: none"> • Ports: 10 GbE NEM0/PORTA and 10 GbE NEM1/PORTB • Connect to client access 10 GbE switches
Backup (BONDETH1)	<ul style="list-style-type: none"> • Ports: 10 GbE NEM0/PORTA and 10 GbE NEM1/PORTB • Connect to backup 10 GbE switches 	<ul style="list-style-type: none"> • Ports: 1 GbE NET1 and 1 GbE NET2 • Connect to backup 1 GbE switches

When channel bonding is configured for the client access network during initial configuration, the Linux bonding module is configured for active-backup mode (`mode=1`). If a different bonding policy is preferred, then you may reconfigure the bonding module after initial configuration. For configuration details, refer to the Linux Ethernet Bonding Driver HOWTO in the `/usr/share/doc/kernel-doc-2.6.18/Documentation/networking/bonding.txt` directory on any Oracle Exadata Database Machine database server.



Note:

When using a bonded network configuration, you must provide a network switch capable of supporting the chosen bonding mode. For example, if mode 4 is configured (IEEE 802.3ad Dynamic link aggregation), then you must supply and configure the network switch capable of supporting this bonding mode. Requirements for specific bonding policies are documented in the Linux Ethernet Bonding Driver HOWTO.

2.2.2 Non-bonded Network Configuration

A non-bonded network configuration is used when there is no need or ability to implement network bonding.

 **Note:**

Non Bonded networks are not supported for Oracle Exadata Database Machine X7-2 and later systems.

The following sections describe the non-bonded network configuration for Oracle Exadata Database Machine:

- [Non-bonded Network Configuration for Oracle Exadata Database Machine Two-Socket Systems](#)
- [Non-bonded Network Configuration for Oracle Exadata Database Machine X5-8 and X6-8](#)
- [Non-bonded Network Configuration for Oracle Exadata Database Machine X4-8 Full Rack](#)
- [Non-bonded Network Configuration for Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack](#)

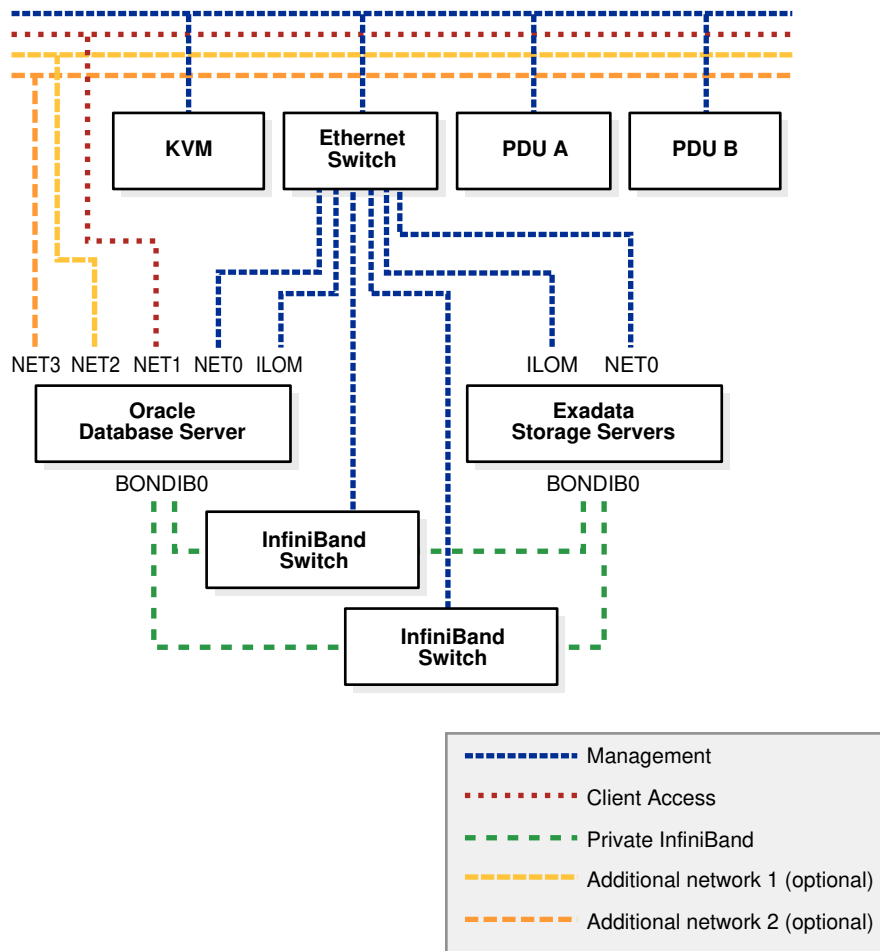
2.2.2.1 Non-bonded Network Configuration for Oracle Exadata Database Machine Two-Socket Systems

When using a non-bonded network configuration for Oracle Exadata Database Machine Two-Socket Systems, the database servers are configured as follows:

- Port NET0: Provides access to the operating system using the management network through the ETH0 interface.
- Port NET1: Provides access to the database using the client access network through the ETH1 interface. The Oracle Real Application Clusters (Oracle RAC) SCAN and VIP addresses are defined on this interface.
- Port NET MGT: Provides access to the Integrated Lights Out Manager (ILOM) using the management network.

The following diagram shows the network diagram for Oracle Exadata Database Machine Two-Socket Systems when not using bonding.

Figure 2-5 Network Diagram for Oracle Exadata Database Machine Two-Socket Systems When Not Using Bonding



2.2.2.2 Non-bonded Network Configuration for Oracle Exadata Database Machine X5-8 and X6-8

When using a non-bonded network configuration for Oracle Exadata Database Machine X5-8 and X6-8, the database servers are configured as follows:

- Port NET0: Provides access to the operating system using the management network through the ETH0 interface.
- Port NET1: Provides access to the database using the client access network through the ETH1 interface. The Oracle Real Application Clusters (Oracle RAC) SCAN and VIP addresses are defined on this interface.
- Port NET MGT: Provides access to the ILOM using the management network.

 **Note:**

Oracle Exadata Database Machine X5-8 and X6-8 have BONDIB0, BONDIB1, and so on for the InfiniBand private network.

2.2.2.3 Non-bonded Network Configuration for Oracle Exadata Database Machine X4-8 Full Rack

When using a non-bonded network configuration for Oracle Exadata Database Machine X4-8 Full Rack, the database servers are configured as follows:

- Port NET0: Provides access to the operating system using the management network through the ETH0 interface.
- Port NET1: Provides access to the database using the client access network through the ETH1 interface. The Oracle RAC SCAN and VIP addresses are defined on this interface.
- Port NET MGT: Provides access to the ILOM using the management network.

 **Note:**

Oracle Exadata Database Machine X4-8 Full Rack has BONDIB0, BONDIB1, and so on for the InfiniBand private network.

2.2.2.4 Non-bonded Network Configuration for Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack

When using a non-bonded network configuration for Oracle Exadata Database Machine X4-8 Full Rack, Oracle Exadata Database Machine X3-8 Full Rack and Oracle Exadata Database Machine X2-8 Full Rack, the database servers are configured as follows:

- Port NEM0/NET0: Provides access to the operating system using the management network through the ETH0 interface. This network must use the 1 GbE port.
- Port NEM0/NET1: Provides access to the database using the client access network through the ETH1 interface. The Oracle RAC SCAN and VIP addresses are defined on this interface.
- Port NET MGT: Provides access to the Integrated Lights Out Manager (ILOM) using the management network.

 **Note:**

Oracle Exadata Database Machine X4-8 Full Rack, Oracle Exadata Database Machine X3-8 Full Rack, and Oracle Exadata Database Machine X2-8 Full Rack have BONDIB0, BONDIB1, and so on for the InfiniBand private network.

2.2.3 Additional Networks

- On Oracle Exadata Database Machine X7-2 and later:
 - Bonded network configuration— There are two open ports (10/25GbE PCI ports) on each database server.
- On Oracle Exadata Database Machine two-socket systems (X2-2 to X6-2):
 - Non-bonded network configuration— There are four open ports (NET2, NET3, and two additional 10 GbE ports) on each database server
 - Bonded network configuration— There are three open ports (NET3 and two additional 10 GbE ports) on each database server
- On Oracle Exadata Database Machine X7-8 and later:
 - Bonded network configuration— There are 13 open ports (5 1/10 GbE ports and 8 open 10/25 GbE PCI ports) on each database server
- On Oracle Exadata Database Machine X5-8 and X6-8:
 - Non-bonded network configuration— There are 18 open ports (10 1GbE ports and 8 10GbE ports) on each database server
 - Bonded network configuration— There are 15 open ports (7 1 GbE ports and 8 open 10 GbE ports) on each database server
- On Oracle Exadata Database Machine X4-8 Full Rack:
 - Non-bonded network configuration— There are 16 open ports (8 1GbE ports and 8 10GbE ports) on each database server
 - Bonded network configuration— There are 15 open ports (7 1 GbE ports and 8 open 10 GbE ports) on each database server
- On Oracle Exadata Database Machine X3-8 Full Rack and X2-8 Full Rack:
 - Non-bonded network configuration— There are 14 open ports (NEM0/NET2, NEM0/NET3, NEM1/NET0, NEM1/NET1, NEM1/NET2, NEM1/NET3, and eight additional 10 GbE ports) on each database server
 - Bonded network configuration— There are 13 open ports (NEM0/NET2, NEM0/NET3, NEM1/NET0, NEM1/NET1, NEM1/NET2, and eight additional 10 GbE ports) on each database server

Use of these additional network ports is optional, and the configuration varies depending on your requirements and available network infrastructure. For example, these interfaces may be used for connectivity to a dedicated tape backup network or connectivity to a disaster recovery network.

2.3 Network Partitioning on Oracle Exadata Database Machine

Oracle Exadata Database Machine supports network partitioning using a variety of mechanisms.

- [VLAN Support on Customer-Facing Networks](#)
Oracle Exadata Database Machine can use VLANs to implement network partitioning in conjunction with the client, backup, administration, and ILOM networks.
- [Access VLAN Support with RoCE Network Fabric](#)
Oracle Exadata Database Machine can use Access VLAN settings to implement server-level isolation across the RoCE Network Fabric.
- [Using Exadata Secure RDMA Fabric Isolation](#)
Starting with Oracle Exadata System Software release 20.1.0, you can configure the RoCE Network Fabric to enable Exadata Secure RDMA Fabric Isolation.
- [Using InfiniBand Partitioning for Network Isolation with InfiniBand Network Fabric](#)
An InfiniBand partition defines a group of InfiniBand nodes or members that are allowed to communicate with one another.

2.3.1 VLAN Support on Customer-Facing Networks

Oracle Exadata Database Machine can use VLANs to implement network partitioning in conjunction with the client, backup, administration, and ILOM networks.

By default, the network switches are minimally configured, without VLAN tagging. If VLAN tagging is to be used, then it can be configured by the customer during the initial deployment. Customers can also configure VLAN tagging after the initial deployment. This applies to both physical and virtual machine (VM) deployments.

 **Notes:**

- Oracle Exadata Deployment Assistant (OEDA) supports VLAN tagging for both physical and VM deployments.
- Network VLAN tagging is supported for Oracle Real Application Clusters (Oracle RAC) on the public network.
- Client and backup VLAN networks must be bonded. The administration network is never bonded.
- If the backup network is on a tagged VLAN network, the client network must also be on a separate tagged VLAN network.
- The backup and client networks can share the same network cables.
- VLAN tagging on the client and backup networks is supported with IPv4 and IPv6 on all hardware models. For IPv6 support on Oracle Database version 12.1.0.2, and later, patch 22289350 is also required.
- VM deployments do not support IPv6 VLANs.
- VLAN tagging on the administration network is only supported with IPv4 addresses on X3-2 and above for two-socket servers, and X4-8 and above for eight-socket servers.
- If the client network uses VLAN tagging and your system uses more than 10 Oracle Clusterware virtual IP (VIP) addresses, then you must use 3-digit VLAN IDs. Do not use 4-digit VLAN IDs because the VLAN name can exceed the operating system interface name limit, which is the 15 characters.

Related Topics

- [Implementing InfiniBand Partitioning across Oracle VM Oracle RAC Clusters on Oracle Exadata](#)
- [Enabling 802.1Q VLAN Tagging in Exadata Database Machine over client networks \(My Oracle Support Doc ID 1423676.1\)](#)
- [Implementing Tagged VLAN Interfaces in Oracle VM Environments on Exadata \(My Oracle Support Doc ID 2018550.1\)](#)


2.3.2 Access VLAN Support with RoCE Network Fabric

Oracle Exadata Database Machine can use Access VLAN settings to implement server-level isolation across the RoCE Network Fabric.

By default, Oracle Exadata Database Machine uses Access VLAN ID 3888 for all RoCE Network Fabric private network traffic, on the server `re0` and `re1` interfaces. This setting enables all database servers and storage servers to communicate freely with each other, and is suitable for many system configurations. However, you can change the Access VLAN ID to a non-default value to implement server-level isolation.

You can use this capability to create isolated groups of servers in an Oracle Exadata Database Machine X8M system. For example, in a Half Rack X8M-2 system you might want to create two isolated server groups:

- Database servers 1 and 2, and storage servers 1, 2, and 3 using VLAN ID 3888
- Database servers 3 and 4, and storage servers 4, 5, 6, and 7 using VLAN ID 3889



With this configuration:

- Database servers 1 and 2 can only access storage servers 1, 2, and 3. But, they cannot access storage servers 4, 5, 6 or 7.
- Database servers 3 and 4 can only access storage servers 4, 5, 6 and 7. But, they cannot access storage servers 1, 2, and 3.
- Oracle Linux KVM guests on database servers 1 and 2 can communicate with each other, but cannot communicate with guests on database servers 3 and 4.
- Oracle Linux KVM guests on database servers 3 and 4 can communicate with each other, but cannot communicate with guests on database servers 1 and 2.

Related Topics

- Using Access VLANs with RoCE Network Fabric

2.3.3 Using Exadata Secure RDMA Fabric Isolation

Starting with Oracle Exadata System Software release 20.1.0, you can configure the RoCE Network Fabric to enable Exadata Secure RDMA Fabric Isolation.

Exadata Secure RDMA Fabric Isolation enables strict network isolation for Oracle Real Application Clusters (Oracle RAC) clusters on Oracle Exadata Database Machine systems that use RDMA over Converged Ethernet (RoCE).

Secure Fabric provides critical infrastructure for secure consolidation of multiple tenants on Oracle Exadata Database Machine, where each tenant resides in a dedicated virtual machine (VM) cluster. Using this feature ensures that:

- Database servers in separate clusters cannot communicate with each other. They are completely isolated from each other on the network.

- Database servers in multiple clusters can share all of the storage server resources. However, even though the different clusters share the same storage network, no cross-cluster network traffic is possible.


Exadata Secure RDMA Fabric Isolation uses RoCE VLANs to ensure that a VM cluster cannot see network packets from another VM cluster. Secure Fabric uses a double VLAN tagging system, where one tag identifies the network partition and the other tag specifies the membership level of the server in the partition. Within each network partition, a partition member with full membership can communicate with all other partition members, including other full and limited members. Partition members with limited membership cannot communicate with other limited membership partition members. However, a partition member with limited membership can communicate with other full membership partition members.

With Secure Fabric, each database cluster uses a dedicated network partition and VLAN ID for cluster networking between the database servers, which supports Oracle RAC inter-node messaging. In this partition, all of the database servers are full members. They can communicate freely within the partition but cannot communicate with database servers in other partitions.

Another partition, with a separate VLAN ID, supports the storage network partition. The storage servers are full members in the storage network partition, and every database server VM is also a limited member. By using the storage network partition:

- Each database server can communicate with all of the storage servers.
- Each storage server can communicate with all of the database servers that they support.
- Storage servers can communicate directly with each other to perform cell-to-cell operations.

Figure 2-6 Secure Fabric Network Partitions



To support the cluster network partition and the storage network partition, each database server VM is plumbed with 4 virtual interfaces:

- `clre0` and `clre1` support the cluster network partition.
- `stre0` and `stre1` support the storage network partition.

Corresponding `stre0` and `stre1` interfaces are also plumbed on each storage server.

On each server, the RoCE network interface card acts like a switch on the hypervisor, which performs VLAN tag enforcement. Since this is done at the KVM host level, cluster isolation cannot be bypassed by any software exploits or misconfiguration on the database server VMs.

In this release, you can only enable Secure Fabric as part of the initial system deployment using Oracle Exadata Deployment Assistant (OEDA). You cannot enable Secure Fabric on an existing X8M system without wiping the system and re-deploying it using OEDA. When enabled, Secure Fabric applies to all servers and clusters that share the same RoCE Network Fabric.

To use Secure Fabric you must:

1. Configure the RoCE Network Fabric switch hardware to enable Secure Fabric. After you complete the switch configuration, the leaf switch ports become trunk ports, which can carry network traffic with multiple VLAN IDs.

The switch configuration must occur before initial system deployment using Oracle Exadata Deployment Assistant (OEDA). See [Configuring the Cisco Nexus C9336C-FX2 Switches to Enable Exadata Secure RDMA Fabric Isolation](#).

2. As part of initial system deployment using Oracle Exadata Deployment Assistant (OEDA), select the option to enable Secure Fabric and specify VLAN IDs for all of the network partitions. See [Cluster Networks](#), in [Using Oracle Exadata Deployment Assistant](#).

2.3.4 Using InfiniBand Partitioning for Network Isolation with InfiniBand Network Fabric

An InfiniBand partition defines a group of InfiniBand nodes or members that are allowed to communicate with one another.

InfiniBand partitioning enables network separation between different clusters on systems with InfiniBand Network Fabric.

InfiniBand partitions are created and managed by the master subnet manager. Each partition is identified by a unique partition key, and partition members use the key for communication within the partition. Members within a partition can only communicate among themselves.

With Oracle Exadata Database Machine, each database cluster uses a dedicated network partition for cluster networking between the database servers. All of the database servers can communicate freely within the partition, other but cannot communicate with database servers in other partitions. Another partition enables communication between each database cluster and the storage servers. By using this partition, database servers can communicate with all of the storage servers, storage servers can communicate with all of the database servers that they support, and storage servers can communicate directly with each other to perform cell-to-cell operations.

You can use InfiniBand partitioning on physical or virtual machine (VM) deployments. For details see [Configuring InfiniBand Partitioning](#).

2.4 Configuring a Separate Network for ILOM

When configuring or re-imaging an Oracle Exadata Rack, you can use Oracle Exadata Deployment Assistant (OEDA) to configure a separate network for Integrated Lights Out Manager (ILOM).

Before Oracle Exadata System Software release 19.1.0, the Exadata servers and ILOM interfaces must have network access to each other for certain features, such as alert notification. Starting with Oracle Exadata System Software release 19.1.0, this network dependency is removed while maintaining all of the previously supported features. Now, you can configure ILOM interfaces on a completely separate network.

1. Use the OEDA Web Configurator to configure the settings for your Oracle Exadata Rack.
2. On the Rack Networks page, check the box **Separate ILOM Network** to configure a separate network for ILOM.

The screenshot shows the Oracle Engineered System Configurator interface. On the left is a navigation menu with categories: Hardware, Operating System, Rack Networks (selected), Users, Clusters, Diskgroups, Database Home, Database, and Cluster Networks. The main content area is titled 'ORACLE Engineered System Configurator' and contains the following configuration options:

- Default gateway for database servers
- Gateway * (text input)
- Subnet Mask (dropdown menu showing 255.255.255.0)
- Separate ILOM Network
- ILOM section:
 - Gateway * (text input)
 - Subnet Mask (dropdown menu showing 255.255.255.0)
 - Start IP Address * (text input)
 - Domain Name * (text input)
- Modify Masks button

A new section appears on the page, with the heading of ILOM. In this area you specify the information required to configure the ILOM network.

3. Provide the required values for the ILOM network:
 - Gateway
 - Start IP Address
 - Subnet Mask
 - Domain Name
4. Optional: If you do not want to use the default names and values for the ILOM network, click the **Modify Masks** button and customize the network settings.

2.5 Default IP Addresses

Starting with Oracle Exadata System Software release 12.1.2.1.0, the NET0 IP addresses are assigned dynamically by the elastic configuration procedure during the first start of the system.

The default NET0 IP addresses are in the 172.16.2.1 to 172.16.7.254 range. In earlier releases, Oracle Exadata Database Machine had default IP addresses set at the factory, and the range of IP addresses was 192.168.1.1 to 192.168.1.203.

 **Note:**

Prior to connecting Oracle Exadata Database Machine to the network, ensure these IP addresses do not conflict with other addresses on the network. Use the `checkip.sh` script generated by Oracle Exadata Deployment Assistant (OEDA) to check for conflicts. You run the `checkip.sh` script on the network after the DNS entries for the Oracle Exadata Database Machine have been created, but before the Oracle Exadata Database Machine is configured or connected to the network. Oracle recommends running the script to avoid configuration delays, even if a check was performed as part of planning process before the machine was delivered. See [Verifying the Network Configuration Prior to Configuring the Rack](#).

If you ran OEDA on a Microsoft Windows system, then the generated script is `checkip.bat`.

2.6 Port Assignments When Using a Firewall

The following table lists the ports used by services on Oracle Exadata Database Machine. When network communication requires access through a firewall, review the list and open the necessary ports. All ports are on the management network, except for port 1521 which is on the client network.

Table 2-6 Open Ports for Firewall

Source	Target	Protocol	Port	Network	Application
Any	Database servers, storage cells, and ILOMs for database servers, Exadata Storage Servers, and RDMA Network Fabric switches	SSH over TCP	22	Management	SSH
Any	KVM	SSH over TCP	22	Management	SSH for serial sessions to MPUIQ-SRL module
Exadata Storage Servers	SMTP e-mail server	SMTP	25	Management	SMTP (Simple Mail Transfer Protocol)

Table 2-6 (Cont.) Open Ports for Firewall

Source	Target	Protocol	Port	Network	Application
Any	ILOMs for database servers, Exadata Storage Servers, and RDMA Network Fabric switches	HTTP	80	Management	Web (user configurable)
Any	PDU	HTTP	80	Management	Browser interface
Any	KVM	HTTP	80	Management	Avocent video viewer download for Java applet
Any	rpcbind	TCP	111	Management	rpcbind
Database management	NTP servers	NTP over UDP	123	Client	Outgoing Network Time Protocol (NTP)
Database servers, Exadata Storage Servers, ILOMs, RDMA Network Fabric switches, and Cisco switch	NTP servers	NTP over UDP	123	Management	Outgoing NTP
Any	ILOMs for database servers, Exadata Storage Servers, and RDMA Network Fabric switches	SNMP over UDP	161	Management	SNMP (Simple Network Management Protocol) (user configurable)
Any	PDU	SNMP over UDP	161	Management	SNMP (user configurable)
Any	KVM	SNMP over UDP	161	Management	SNMP (user configurable)

Table 2-6 (Cont.) Open Ports for Firewall

Source	Target	Protocol	Port	Network	Application
Exadata Storage Servers	SNMP subscriber such as Oracle Enterprise Manager Cloud Control or an SNMP manager	SNMP	162	Management	SNMP version 1 (SNMPv1) outgoing traps (user-configurable)
Database servers, Exadata Storage Servers, and ILOMs for database servers, Exadata Storage Servers, and RDMA Network Fabric switches	ASR Manager	SNMP	162	Management	Telemetry messages sent to ASR Manager
ILOMs for database servers, Exadata Storage Servers, and RDMA Network Fabric switches (Excluding X8M racks)	Any	IPMI over UDP	162	Management	Outgoing IPMI (Intelligent Platform Management Interface) Platform Event Trap (PET)
PDU	SNMP trap receivers	SNMP over UDP	162	Management	Outgoing SNMPv2 traps
KVM	SNMP trap receivers	SNMP over UDP	162	Management	Outgoing SNMPv2 traps
ASR Manager	ASR backend	HTTPS	443	Management	Telemetry messages sent to ASR backend
ASR Manager	ASR backend	HTTPS	8100 16161	Management	DiagPack uploads

Table 2-6 (Cont.) Open Ports for Firewall

Source	Target	Protocol	Port	Network	Application
Any	ILOMs for database servers, Exadata Storage Servers, and RDMA Network Fabric switches	HTTPS	443	Management	Web (user configurable)
Any	PDU	HTTPS	443	Management	Browser interface
Any	KVM	HTTPS	443	Management	Browser interface for MergePoint Utility switch and KVM sessions
Exadata Storage Servers	SMTSPS client	SMTSPS	465	Management	Simple Mail Transfer Protocol, Secure (if configured)
Database servers, Exadata Storage Servers, and ILOMs for database servers, Exadata Storage Servers, and RDMA Network Fabric switches	Syslog server	Syslog over UDP	514	Management	Outgoing Syslog
PDU	Syslog server	Syslog over UDP	514	Management	Outgoing Syslog
KVM	Syslog server	Syslog over UDP	514	Management	Outgoing Syslog

Table 2-6 (Cont.) Open Ports for Firewall

Source	Target	Protocol	Port	Network	Application
Any	ILOMs for database servers, Exadata Storage Servers, and RDMA Network Fabric switches (not applicable to X8M servers)	IPMI over UDP	623	Management	IPMI
Any	plathwsvcd	TCP	723	Management	
Any	evnd	TCP	791	Management	
Any	partitiond	TCP	867	Management	
Any	Database servers	TCP	1521	Client	Database listener
Any	KVM	TCP	2068	Management	KVM session data for keyboard and mouse transmission, or for video transmission
Any	tgtd	TCP	3260	Management	SCSI target daemon
Any		TCP	3872	Management	Java EM agent
Storage Servers	Fast node death detection (FNDD) on RDMA over Converged Ethernet (RoCE) systems	TCP	5053	Management	
Any	Cisco Smart Install	TCP	4786	Management	Cisco Smart Install
Any	ILOMs for database servers, and Exadata Storage Servers	TCP	5120	Management	ILOM remote console: CD

Table 2-6 (Cont.) Open Ports for Firewall

Source	Target	Protocol	Port	Network	Application
Any	ILOMs for database servers, and Exadata Storage Servers	TCP	5121	Management	ILOM remote console: keyboard and mouse
Any	ILOMs for database servers, and Exadata Storage Servers	TCP	5123	Management	ILOM remote console: diskette
Any	ILOMs for database servers, and Exadata Storage Servers	TCP	5555	Management	ILOM remote console: encryption
Any	ILOMs for database servers, and Exadata Storage Servers	TCP	5556	Management	ILOM remote console: authentication
Any	Oracle Notification Service	TCP	6200	Management	Oracle Notification Service
ASR Manager	ILOMs for database servers, and Exadata Storage Servers	HTTP	6481	Management	Service tag listener for asset activation
Any	ILOMs for database servers, and Exadata Storage Servers	TCP	6481	Management	ILOM remote console: Servicetag daemon
Any	ILOMs for database servers, and Exadata Storage Servers	TCP	7578	Management	ILOM remote console: video
Any	ILOMs for database servers, and Exadata Storage Servers	TCP	7579	Management	ILOM remote console: serial

Table 2-6 (Cont.) Open Ports for Firewall

Source	Target	Protocol	Port	Network	Application
Any	Database servers and Exadata Storage Servers	TCP	7777	Both	Oracle Enterprise Manager Grid Control HTTP console port
Any	Database servers and Exadata Storage Servers	TCP	7799	Both	Oracle Enterprise Manager Grid Control HTTPS console port

Table 2-6 (Cont.) Open Ports for Firewall

Source	Target	Protocol	Port	Network	Application
Any	Oracle WebLogic	TCP	7878 8888	Management	Management Server (MS)


 **N
o
t
e**
:
A
p
p
l
i
e
s
o
n
l
y
t
o
O
r
a
c
l
e
E
x
a
d
a
t
a
S
y
s
t
e
m
S
o
f
t
w
a

Table 2-6 (Cont.) Open Ports for Firewall

Source	Target	Protocol	Port	Network	Application
Diagpack uploads	Database servers and Exadata Storage Servers		8100 16161		
Any	rpc.statd	TCP	21408 40801 41460 47431	Management	rpc.statd
Any	gpnpd.bin	TCP	19193 (IPv6) 25194	Management	gpnpd.bin

 **See Also:**

Appendix E, "Managing Oracle Database Port Numbers", in the *Oracle Real Application Clusters Installation Guide for Linux and UNIX*.

3

Using Oracle Exadata Deployment Assistant

Oracle Exadata Deployment Assistant gathers your configuration details and creates the Oracle Exadata Rack configuration file. The configuration file drives the automated installation and configuration processes for Oracle Exadata Rack.

You can also use the OEDA command-line interface (OEDACLI) to perform Oracle Exadata Rack life-cycle management tasks.

You can download the latest version of OEDA from [Exadata Database Machine and Exadata Storage Server Supported Versions \(Doc ID 888828.1\)](#). OEDA is also available on Oracle Technology Network.

Note:

For ease of reading, Oracle Exadata Rack is used when information refers to both Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack.

- [Overview of Oracle Exadata Deployment Assistant](#)
OEDA is used to create configuration files and then implement the configuration on your engineered system.
- [OEDA Browser-based User Interface](#)
The Web-based interface for Oracle Exadata Deployment Assistant (OEDA) is available starting with the release October 2018 of OEDA. OEDA Web is available for Linux, OSX and Windows.

Related Topics

- [OEDA Command Line Interface](#)
When you run the `oedaccli` program, it starts the OEDA command line interface.

3.1 Overview of Oracle Exadata Deployment Assistant

OEDA is used to create configuration files and then implement the configuration on your engineered system.

OEDA is used for Oracle Exadata Database Machine, Oracle Zero Data Loss Recovery Appliance, and Oracle SuperCluster. Starting with Oracle Exadata System Software release 19.1.0, there is a new Web-based interface for OEDA. However, the overall process, described here, is the same regardless of whether you use the Web-based or Java-based user interface.

- Before your engineered system arrives, do the following:
 1. Work with your network and database administrators to evaluate the current network settings, such as current IP address use and network configuration.

2. Define the settings for the rack, such as network configuration and backup method.
3. Download the latest version of OEDA from Oracle Technology Network.
4. Unzip and run the configuration script on a supported platform, and select the type of system to configure.
5. Go through every page in OEDA and supply values for all required fields. You cannot advance to the next page if you do not supply all of the required values. You must provide naming details and DNS, NTP, and networking details.
6. At the end of the dialogue with OEDA, configuration files are generated on the client. The files are also listed at the bottom of the `InstallationTemplate.html` file that is generated by OEDA. Depending on your engineered system and configuration, OEDA generates all or some of the following files:

- `databasemachine.xml`
- `CustomerName-rackname.xml`
- `CustomerName-rackname-preconf_GUID.csv`
- `CustomerName-rackname-InstallationTemplate.html`
- `CustomerName-rackname-platinum.csv`
- `CustomerName-rackname-checkip.sh`
- `CustomerName-rackname.zip`
- `pkey_GUID.csv` and `pkey_racknamehostname_GUID.csv` — if you enabled InfiniBand partitioning for your virtual environments

The `CustomerName-hostname.zip` file contains all the generated files.

7. Review the `InstallationTemplate.html` file to check the entire configuration and verify all information was entered correctly.
- Shortly before your engineered system arrives, or is scheduled to be configured, validate the network configuration, as directed by Oracle. See [Verifying the Network Configuration Prior to Configuring the Rack](#).
 - After your engineered system arrives, the configuration files are copied to a database server, and the validation and installation is completed. See [Configuring Oracle Exadata Database Machine Using OEDA](#).

Accessibility Note: If you are using Oracle Exadata Deployment Assistant with assistive technology (for accessibility) on Windows, you need to enable Java Access Bridge.

- [Considerations and Requirement for Using OEDA](#)
Review this information before using Oracle Exadata Deployment Assistant (OEDA).

3.1.1 Considerations and Requirement for Using OEDA

Review this information before using Oracle Exadata Deployment Assistant (OEDA).

- Oracle Exadata Rack ships with the Oracle Linux operating system installed on the servers.

- The rack name is used to generate host names for all components. For example, if you specify `dbm0` for the rack name:
 - The database server host names are like `dbm0db01`
 - The storage server host names are like `dbm0ce101`
 - The InfiniBand Network Fabric switch names are like `dbm0sw-iba1`
 - The RoCE Network Fabric switch names are like `dbm0sw-rocea1`

If you have multiple racks, then, each Oracle Exadata Rack name should be a unique name that identifies the machine. Oracle recommends using `dm01` for the first Oracle Exadata Rack, `dm02` for the second, `dm03` for the third, and so on. Up to eight Oracle Exadata Racks can be cabled together without additional InfiniBand switches.

- The backup method and redundancy protection levels you use determine how to size the Oracle Automatic Storage Management (Oracle ASM) disk groups that are created during installation.
 - If you configure internal backups, then the database backups are created on disk in the Fast Recovery Area in the RECO disk group for the database. The Fast Recovery Area also contains Archived Redo Logs and Flashback Log Files. The division of disk space between the DATA disk group and the RECO disk group is 40% and 60%, respectively.
 - If you configure external backups, then database backups are created on disk or tape media that is external to currently deployed Oracle Exadata Database Machine. The Fast Recovery Area contains only objects such as Archived Redo Logs and Flashback Log Files.

Because you do not need to reserve additional space in the RECO disk group for backups, the division of disk space between the DATA disk group and the RECO disk group is now 80% and 20%, respectively.

- Oracle recommends the use of high redundancy disk groups for mission critical applications. Use the following guidelines to assist in sizing the disk groups based on the redundancy levels for the disk groups:

Table 3-1 Protection Levels and Disk Group Contents

Oracle ASM Redundancy Level for DATA Disk Group	Oracle ASM Redundancy Level for RECO Disk Group	DATA Group Contents	RECO Group Contents
High	High	Data files Temporary files Online redo logs Control file	Archive logs Flashback log files
High	Normal	Data files Online redo logs Control file.	Archive logs Temporary files Flashback log files
Normal	High	Data files Temporary files	Online redo logs Control file Archive logs Flashback log files

Table 3-1 (Cont.) Protection Levels and Disk Group Contents

Oracle ASM Redundancy Level for DATA Disk Group	Oracle ASM Redundancy Level for RECO Disk Group	DATA Group Contents	RECO Group Contents
Normal	Normal	Data files Temporary files Online redo logs Control file	Archive logs Flashback log files

- A valid time zone name is required when installing Oracle Exadata Database Machine. Time zone values provided for Oracle Exadata Database Machine and Oracle Linux comes from the Time Zone Database. You must specify a valid time zone name for each server using the `TZ` environment variable. The variable value is of the form *Area/Location*. For example, *America/New_York* is a valid `TZ` value; `EST`, `EDT`, and `UTC-4` are invalid values.
- OEDA configures all Exadata components, including the database servers (with or without virtualization), storage servers, and the RDMA Network Fabric. For successful configuration, OEDA requires access to:
 - The client and administration network interfaces on the database servers, including the VM server host and the VM guests where applicable.
 - The administration network interfaces on the storage servers.
 - The administration network interfaces on the RDMA Network Fabric switches.

Ensure that you run the OEDA deployment phase on a host with access to all of the required network interfaces.

Related Topics

- [iana Time Zone Database](#)
- Maximum Availability with Oracle ASM

3.2 OEDA Browser-based User Interface

The Web-based interface for Oracle Exadata Deployment Assistant (OEDA) is available starting with the release October 2018 of OEDA. OEDA Web is available for Linux, OSX and Windows.



Note:

The OEDA Web interface does not support Oracle SuperCluster. Instead, use the `configssc.sh` command.

You can use this interface to enter configuration details for new racks and also import configuration files from the old OEDA interface. Use the **Menu** button to import and save configurations.

To use the Web-based version of OEDA, perform the following steps.

1. Download the latest OEDA release from [My Oracle Support note 888828.1](#).

Extract the contents of the downloaded compressed file. When you extract the contents, it creates a directory based on the operating system, such as `linux-x64`, `macosx-x64`, or `windows-i586`, to store the extracted files. This is referred to as the `OEDA_HOME` directory.

2. Install the software.

Before you can use the Web-based interface, you must install and run the Web Application Server. In the created directory, locate and run the `installOedaServer` program. You do not have to be logged in as an administrator user to run this program. Use one of the following commands, where the `-p` option specifies the port number to use:

- On Linux, Apple, and UNIX:

```
./installOedaServer.sh -p 7072
```

- On Microsoft Windows:

```
installOedaServer.cmd -p 7072
```

Note:

The examples use port number 7072. However, you can specify any unused port number. It is not recommended to use port numbers less than 1024.

The `installOedaServer` program starts the latest version of the OEDA Web Server on the local system.

3. Access the local OEDA Web Server.

Once the OEDA Web Server is running, you can access the Web-based application by opening a browser and entering the following URL:

```
http://localhost:port_number/oeda
```

In the URL, `port_number` is the network port number that you specified when starting the `installOedaServer` program. For example, 7072.

The following topics describe the OEDA Web interface:

- [Using the Browser-based Version of Oracle Exadata Deployment Assistant](#)
- [Select Hardware Page](#)
- [Choose Operating System Page](#)
- [Rack Networks Page](#)
- [Users and Groups](#)
- [Define Clusters](#)
- [Diskgroups](#)
- [Database Home](#)
- [Database](#)

- [Cluster Networks](#)
- [Alerting](#)
- [Tape Library Connectivity](#)
On ZDLRA racks only, you can configure network connection information for tape backups.
- [Comments](#)

3.2.1 Using the Browser-based Version of Oracle Exadata Deployment Assistant

Use the deployment assistant for initial configuration, and when adding to an existing deployment. You can import an existing configuration when adding new racks or components.

Have the following available before running the deployment assistant:

- IP addresses
- Machine names
- DNS information
- NTP information for the rack

The following options are available from the OEDA **Menu** drop-down list:

- **About:** View information about the web interface software.
- **New:** Start a new configuration.
- **Import:** Import an existing configuration file. The file format is XML.
- **Save:** Save the current configuration to an XML file.
- **Diagnostics:** Creates a ZIP file that can be used to troubleshoot and diagnosis issues with the OEDA web interface.
- **Preferences:** Specify your preferences for how the web interface behaves.
- **Validate:** Check the current configuration for completeness and correctness.
- **CheckIp:** Checks the IP addresses.
- **Exit:** Exit OEDA

The following list describes the configuration information you specify in OEDA:

1. Select the hardware for the deployment on the **Select Hardware** page. Also provide customer information on this page.
2. Choose whether to configuration Oracle VM on the nodes.

 **Note:**

Creating more than 8 VMs on Exadata Database Servers is not supported.

3. Enter the management network and private network information on the **Rack Networks** page.

4. Specify the user and group configuration on the **Users and Groups** page. You can choose between Default, Role Separated, or a Custom configuration.
5. Define the clusters on the **Define Clusters** page for each cluster.
6. Define the disk group layout on the **Diskgroups** page.
7. Specify the Oracle Database home information for each Oracle home on the Create **Database Home** page.
8. Enter the client and backup network information on the **Cluster Networks** page.
9. Configure the email alerts, SNMP alerts, Auto Service Request, Oracle Config Manager, and Enterprise Manager alerting on the Alerting page.
10. Answer the configuration questionnaire and provide additional information about the deployment in the text field on the **Comments** page. Information entered in this field appears at the bottom of the `InstallationTemplate.html` file.
11. Click **Save and Download** to create the configuration files. You are prompted to choose a location in which to save a ZIP file to your local computer.
12. Copy the ZIP file to a directory on the first database server in Oracle Exadata Rack. Use a subdirectory of `/u01`, such as the `/u01/oeda` directory. These files are used by the field engineer to configure the Oracle Exadata Rack.

3.2.2 Select Hardware Page

Select the hardware type being configured and provide customer details.

- **Oracle Exadata Database Machine**

Supply the following additional information for this choice:

- **Rack Size:** The size of the rack you are configuring. Your choices are:
 - * **Elastic:** (X4-8 and later) A customized number of database servers and storage servers.
 - * **Full:** A standard configuration consisting of:
 - * 2-socket models (X8M-2, X8-2, X7-2, X6-2, and so on): Eight 2-socket database servers and 14 storage servers
 - * 8-socket models (X8M-8, X8-8, X7-8, X6-8, and so on): Two 8-socket database servers and 14 storage servers.
 - * Storage Expansion Racks (X2, X3, and X4): 18 storage servers
 - * **Half:** A standard configuration consisting of:
 - * 2-socket models (X8-2, X7-2, X6-2, and so on): Four 2-socket database servers and 7 storage servers
 - * Storage Expansion Racks (X2, X3, and X4): Nine storage servers
 - * **Quarter:** A standard configuration consisting of:
 - * 2-socket models (X8-2, X7-2, X6-2, and so on): Two 2-socket database servers and 3 storage servers
 - * Storage Expansion Racks (X2, X3, and X4): Four storage servers
 - * **Eighth:** A standard configuration consisting of 2 database servers and 3 storage servers with only half the cores enabled in each server. Available for 2-socket systems model X3-2 and later.

- * **Eighth Elastic:** A customizable configuration with 2 database servers and 3 storage servers. Available for 2-socket systems model X4-2 and later.
- **Compute Node, Storage Cell, and Disk Size:** The values you can select for each successive field are determined by the value you selected for the previous fields:
 - * **X8-2** Compute Node
 - * **X8** Storage Cell
 - * **HC 14TB** (High Capacity 14 TB drives)
 - * **EF 6.4 TB** (Extreme Flash 6.4 TB drives)
 - * **X8-8** Compute Node
 - * **X8** Storage Cell
 - * **HC 14TB** (High Capacity 14 TB drives)
 - * **EF 6.4 TB** (Extreme Flash 6.4 TB drives)
 - * **X7-2** Compute Node
 - * **X7** Storage Cell
 - * **HC 10TB** (High Capacity 10 TB drives)
 - * **EF 6.4 TB** (Extreme Flash 6.4 TB drives)
 - * **X7-8** Compute Node
 - * **X7** Storage Cell
 - * **HC 10TB** (High Capacity 10 TB drives)
 - * **EF 6.4 TB** (Extreme Flash 6.4 TB drives)
 - * **X6-2** Compute Node
 - * **X6** Storage Cell
 - * **EF 3.2TB** (Extreme Flash 3.2 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
 - * **X6-8** Compute Node
 - * **X6** Storage Cell
 - * **EF 3.2TB** (Extreme Flash 3.2 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
 - * **X5-8** Compute Node
 - * **X5** Storage Cell
 - * **EF 1.6TB** (Extreme Flash 1.6 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
 - * **X5-2** Compute Node
 - * **X5** Storage Cell
 - * **EF 1.6TB** (Extreme Flash 1.6 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
 - * **HC 4TB** (High Capacity 4 TB drives)

- * **X4-8** Compute Node
 - * **X5** Storage Cell
 - * **EF 1.6TB** (Extreme Flash 1.6 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
 - * **HC 4TB** (High Capacity 4 TB drives)
 - * **X4** Storage Cell
 - * **HP 1.2TB** (High Performance 1.2 TB drives)
 - * **HC 4TB** (High Capacity 4 TB drives)
- * **X4-2** Compute Node
 - * **X4** Storage Cell
 - * **HP 1.2TB** (High Performance 1.2 TB drives)
 - * **HC 4TB** (High Capacity 4 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
- * **X4 Storage Expansion** Compute Node
 - * **X4** Storage Cell
 - * **HP 1.2TB** (High Performance 1.2 TB drives)
 - * **HC 4TB** (High Capacity 4 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
- * **X3-2** Compute Node
 - * **X3** Storage Cell
 - * **HP** (High Performance drives)
 - * **HC 3TB** (High Capacity 3 TB drives)
 - * **HC 4TB** (High Capacity 4 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
- * **X3-8** Compute Node
 - * **X4** Storage Cell
 - * **HP 1.2TB** (High Performance 1.2 TB drives)
 - * **HC 4TB** (High Capacity 4 TB drives)
 - * **X3** Storage Cell
 - * **HP** (High Performance drives)
 - * **HC 3TB** (High Capacity 3 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
- * **X3 Storage Expansion** Compute Node
 - * **X3** Storage Cell
 - * **HP** (High Performance drives)
 - * **HC 3TB** (High Capacity 3 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)

- * **X2-2** Compute Node
 - * **X2** Storage Cell
 - * **HP** (High Performance drives)
 - * **HC 2TB** (High Capacity 2 TB drives)
 - * **HC 3TB** (High Capacity 3 TB drives)
 - * **HC 4TB** (High Capacity 4 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
 - * **X2-8** Compute Node
 - * **X2** Storage Cell
 - * **HP** (High Performance drives)
 - * **HC 2TB** (High Capacity 2 TB drives)
 - * **HC 3TB** (High Capacity 3 TB drives)
 - * **HC 8TB** (High Capacity 8 TB drives)
 - * **X2 Storage Expansion** Compute Node
 - * **X2** Storage Cell
 - * **HP** (High Performance drives)
 - * **HC 2TB** (High Capacity 2 TB drives)
 - * **HC 3TB** (High Capacity 3 TB drives)
 - **Add Rack (+)**: Click this button after you have selected the hardware to add the rack to the configuration and display the Rack Details options.
 - **Zero Data Loss Recovery Appliance**

Supply the following additional information for this choice:

 - **ZDLRA Server, Capacity, and Cell Count**: The model of the ZDLRA rack, the disk storage types and the cell count. The values you can select for each successive field are determined by the value you selected for the previous fields:
 - * Server **X8-2**
 - * Capacity **HC 14TB** (High Capacity)
 - * Server **X7-2**
 - * Capacity **HC 10TB** (High Capacity)
 - * Server **X6-2**
 - * Capacity **HC 8TB** (High Capacity)
 - * Server **X5-2**
 - * Capacity **HC 4TB** (High Capacity)
 - * Capacity **HC 8TB**
 - * Server **X4 Zdlra**
 - * Capacity: **HC 4TB** (High Capacity)

- * **Cell Count:** The number of cells in the rack. Choose the value from the drop-down list. Available values are from 3 to 13. You can also use the Search field in the drop-down list to find a value.
- **Add Rack (+):** Click this button after you have selected the hardware to add the rack to the configuration and display the Rack Details options.
- **Exadata Cloud**

Supply the following additional information for this choice:

 - **Exadata cloud options:** Select the Exadata deployment option from the list, for example, Cloud at Customer.
 - **Rack Size:** The size of the rack you are configuring. Your choices are:
 - * **Full:** A standard configuration consisting of eight 2-socket database servers and 12 storage servers.
 - * **Half:** A standard configuration consisting of four 2-socket database servers and 6 storage servers.
 - * **Quarter:** A standard configuration consisting of two 2-socket database servers and 3 storage servers.
 - * **Base Rack:** A configuration consisting of 2 database servers and 3 storage servers. You can specify the minimum and maximum number of compute cores (OCPUs) to enable on each database server.
 - **Compute Node, Storage Cell, and Disk Size:** The values you can select for each successive field are determined by the value you selected for the previous fields:
 - * **X8-2 Compute Node**
X8 Storage Cell
HC 14TB (High Capacity 14 TB drives)
 - * **X7-2 Compute Node**
X7 Storage Cell
HC 10TB (High Capacity 10 TB drives)
 - * **X6-2 Compute Node**
X6 Storage Cell
HC 8TB (High Capacity 8 TB drives)
 - **Add Rack (+):** Click this button after you have selected the hardware to add the rack to the configuration and display the Rack Details options.

Rack Details

Provide additional information about the rack and how it will be deployed.

- **Rack Name:** Value is supplied automatically based on the inputs in the previous section.
- **Rack Prefix:** Defines the prefix for network interfaces. The prefix can contain any alphanumeric character. The prefix *cannot* contain spaces. The maximum length is 20 characters. This field cannot be empty. The default value is dbm0.

For example, a value of dbm0 results in a database server host name of dbm0db01, and a storage server host name of dbm0ce101.

- **Spine Switch:** Click the check box if you want to include a spine switch as part of the rack configuration. Do this for each rack that will have a spine switch.

Customer Details

Provide information about the environment in which the rack will be deployed.

- **Customer Name:** Identifies the customer. This is a free-form field that can contain any alphanumeric characters, including spaces. This field cannot be empty.
- **Application Name:** Identifies the application used with the rack. This is a free-form field that can contain any alphanumeric characters, including spaces. This field cannot be empty.
- **Region:** Sets the region for the rack. Choose a value from the drop-down list. You can also use the Search field to locate a region. Your choices are:
 - Africa
 - America
 - Antarctica
 - Arctic
 - Asia
 - Atlantic
 - Australia
 - Europe
 - Indian
 - Other
 - Pacific
- **Timezone:** Sets the time zone where the rack will be installed. The time zones shown are based on the selected region. Choose a value from the drop-down list.
- **Compute Node #:** (only for X5-2 and X6-2 ZDLRA Elastic Rack) Specifies the number of compute nodes to configure
- **Storage Node #:** (only for X5-2 and X6-2 ZDLRA Elastic Rack) Specifies the number of storage nodes to configure
- **DNS Servers:** Sets the DNS IP Address. At least one IP address must be entered. For Exadata Cloud, these values are provided for you.
- **NTP Servers:** Sets the IP Address for the Network Time Protocol (NTP) server. At least one IP address must be entered. For Exadata Cloud, these values are provided for you.

WriteBack Flash Cache

Specify whether to enable WriteBack Flash Cache when the rack is deployed. WriteBack flash cache provides the ability to cache write I/Os directly to PCI flash in addition to read I/Os. This option is not available for Exadata Cloud.

Your choices are:

- **Automatic:** Default value. When selected, WriteBack Flash Cache will be enabled if the rack is running a software version that supports WriteBack Flash Cache, and if using WriteBack Flash Cache is a best practice for that software version.

- **Disable:** This options disables WriteBack Flash Cache, and only read IOs are stored in the Flash Cache.
- **Enable:** This option enables storing both read and write IOs in the Flash Cache.

Capacity on Demand

(Available for only Exadata racks X3 and later) Click this check box to reduce the initial number of active cores on the servers.

If you select this option, a slider bar appears which allows to select the number of active cores, with a range that is specific to the hardware model of the database server.

Note:

To qualify for Capacity on Demand, within 3 months after installation, this Exadata system must use either Platinum Services, Oracle Enterprise Manager Harvester (in disconnected or connected mode) or Oracle Configuration Manager (in connected mode).

Add Additional Racks

You can start at the top of the page and provide the information for an additional rack. Click the **Add Rack (+)** button to add the rack to the configuration. An additional tab appears in the Rack Details section.

Save

After you have entered all the information for the racks being configured, click the **Save** button in the lower right corner of the page.

3.2.3 Choose Operating System Page

Select the operating system to install on the servers in the rack.

This page is not available for Exadata Cloud configurations.

For Oracle Exadata Rack, your choices are:

- **All Linux VM:** All database servers running inside virtual machines (VMs) on Oracle Linux.

Starting with Oracle Exadata Database Machine X8M-2, Oracle Linux KVM is the virtualization technology for systems that use RoCE Network Fabric. Otherwise, Oracle VM is a Xen-based virtualization technology, which is used across Oracle Exadata Database Machine systems that use InfiniBand Network Fabric.

- **All Linux Physical:** All database servers are installed with Oracle Linux and no VM server configuration.
- **Custom:** Specify which database servers are installed with a VM server, and which are not.

Choosing this option displays a list of available database servers. You must select one of the following options for each database server:

- **Linux Physical:** No VM server.

- **Linux Dom0:** VM server.

Starting with Oracle Exadata Database Machine X8M-2, Oracle Linux KVM is the virtualization technology for systems that use RoCE Network Fabric. Otherwise, Oracle VM is a Xen-based virtualization technology, which is used across Oracle Exadata Database Machine systems that use InfiniBand Network Fabric.

You can also select the option just under the heading Linux Physical or Linux Dom0 to select that option for all the database servers in the list.

3.2.4 Rack Networks Page

Provide the information necessary to configure the administration and private networks on the rack.

Each compute node and storage server has two network interfaces for administration. One network provides management access to the operating system through the Ethernet interface, and the other network provides access to the Integrated Lights Out Manager (ILOM) through the ILOM Ethernet interface. Default administration network host names are usually changed to meet corporate naming standards or compatibility with existing systems.

In the upper right corner of the Admin Network section is a line that displays the Total IP Addresses that will be configured for the rack. This value is based on the hardware you selected on the Hardware page.

You provide the following information on this page:

Admin Network

Supply values for the following fields. *For Exadata Cloud, these values are provided for you.*

- **Default gateway for database servers:** A check box that indicates that the gateway IP address is the default gateway for the database servers.
- **Gateway:** Sets the IP address for the gateway. This field cannot be empty.
- **Start IP Address:** Assigns the starting IP address for the administration network of the servers and switches. This field cannot be empty.

Oracle Exadata Deployment Assistant (OEDA) generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified during the configuration process.

Note:

These addresses are also assigned to the ILOM ports on the database and storage servers

- **Subnet Mask:** Sets the IP address mask for the subnet. You can choose a subnet mask from the drop-down list which contains supported values from 255.0.0.0 to 255.255.255.252
- **Domain Name:** Sets the network domain name. Use lowercase letters for the name. This field cannot be empty.

The domain name specified on this page is used as the domain for all networks. If different domains are used for the client and backup networks, then specify the administration network on this page. The client and backup networks can be changed on the Cluster Networks page.

- **Separate ILOM Network:** *(Not available with Exadata Cloud)* A check box that indicates that ILOM uses a network that is different from the administration network. If you check this option, you must specify the following additional information for the ILOM network:
 - **Gateway**
 - **Start IP Address**
 - **Subnet Mask**
 - **Domain Name**

Advanced

(Not available with Exadata Cloud) If you click the **Advanced** button in the upper right corner of this page, you can select the following options:

- **Remove PDU's** — Select this option if you do not want to add administration IP addresses to the PDUs. If you select this option, then OEDA does not allocate IP addresses for each of the two PDUs in the rack, and the administration names do not appear in the Admin network Review and Edit section of this page.
- **Enable Vlan** — Allows you to create VLANs on database servers and storage servers for the admin network, ILOM, client and the backup access network. If you select this option, then the following additional fields are visible on the Admin Network page:
 - **Vlan** — Specifies the VLAN identifier. The value must be a positive integer.

Note:

- * If your system will use more than 10 VIP addresses in the cluster and you have VLAN configured for the Oracle Clusterware client network, then you must use 3 digit VLAN ids.
 - * Do not use 4 digit VLAN ids because the VLAN name can exceed the 15 character operating system interface name limit.
 - * If VLAN tagging is to be used, then it must be enabled by the customer after the initial deployment.
- See the My Oracle Support notes for details.

Modify Masks

(Not available with Exadata Cloud) If you click the button **Modify Masks**, it opens a pop-up window where you can enter the following information:

- **Compute Name Mask** — There is a preview field at the bottom of this window that shows how the generated administration names will appear based on the values you specify for the following fields:

 **Note:**

In the following fields, the name prefix can contain alphanumeric characters. The percent signs (%) in the names are replaced by the numbers entered into the Start ID field. At least one percent sign is required in each mask field. If you use more than one percent sign in a mask field, the symbols must be contiguous. For example, you can use `test%-ilom`, but not `test%new%-ilom`.

- **Compute Name Mask:** The string used to generate administration names for each database server in this rack. The default string is `{Rack Prefix}adm%%`, which results in names such as `dbm0adm01`, `dbm0adm02`, and so on. This field cannot be empty.
- **Compute Name Start Id:** The starting value to use when generating the database server administration names. The default value is 1. This field cannot be empty.
- **Compute ILOM Mask** — There is a preview field at the bottom of this section that shows how the generated administration names will appear based on the values you specify for the following fields:
 - **Compute ILOM Mask:** The string used to generate administration names for each database server ILOM in this rack. The default string is `{Compute Name Mask}%-ilom`, which results in names such as `dbm0adm01-ilom`, `dbm0adm02-ilom`, and so on. This field cannot be empty.
 - **ILOM Name Start Id:** The starting value to use when generating the ILOM administration names. The default value is the same value as Compute Name Start Id. This field cannot be empty.

 **Note:**

Oracle recommends using the same starting IDs for the servers and server ILOMs.

- **Storage Name Mask** — There is a preview field at the bottom of this section that shows how the generated administration names will appear based on the values you specify for the following fields:
 - **Storage Name Mask:** The string used to generate administration names for each storage server. The default string is `{Rack Prefix}celadm%%`, which results in names such as `dbm0celadm01`, `dbm0celadm02`, and so on. This field cannot be empty.
 - **Storage Name Start Id:** The starting value to use when generating the storage server network names. The default value is 1. This field cannot be empty.
- **Storage ILOM Mask** — There is a preview field at the bottom of this section that shows how the generated server names will appear based on the values you specify for the following fields:
 - **Storage ILOM Mask:** The string used to generate administration names for each storage server ILOM in this rack. The default string is `{Storage`

Name Mask}%-ilom, which results in names such as dbm0celadm01-ilom, dbm0celadm02-ilom, and so on. This field cannot be empty.

- **Storage Name Start Id:** The starting value to use when generating the ILOM network names. The default value is the same value as Storage Name Start Id. This field cannot be empty.

 **Note:**

Oracle recommends using the same starting IDs for the servers and server ILOMs.

- Switches and PDUs
 - **Cisco Switch** — The string used to generate the administration name for the Cisco Management Switch in this rack. The default value is `{Rack Prefix}sw-adm%`, which results in names such as dbm0sw-adm1.
 - **RDMA Network Fabric Spine Switch** — The string used to generate the administration names for the spine switch, if your rack is configured with one. The default value is `{Rack Prefix}sw-ibs%`, which results in names such as dbm0sw-ibs1.
 - **RDMA Network Fabric Leaf Switch** — The string used to generate the administration names for the lower RDMA Network Fabric Leaf switch in this rack. The default value is `{Rack Prefix}sw-typea%`, which results in names such as dbm0sw-iba1 or dbm0sw-rocea0.
 - **RDMA Network Fabric Leaf Switch** — The string used to generate the administration names for the upper RDMA Network Fabric Leaf switch in this rack. The default value is `{Rack Prefix}sw-typeb%`, which results in names such as dbm0sw-ibb1 or dbm0sw-roceb0.
 - **PDU-A** — The string used to generate the administration names for Power Distribution Unit (PDU) A in this rack. The default value is `{Rack Prefix}sw-pdua%`, which results in names such as dbm0sw-pdua1.
 - **PDU-B** — The string used to generate the administration names for Power Distribution Unit (PDU) B in this rack. The default value is `{Rack Prefix}sw-pdub%`, which results in names such as dbm0sw-pdub1.
 - **KVM Switch** — (Exadata X2-2 only) The string used to generate the administration name for the keyboard, video, and mouse (KVM) switch. The default value is `{Rack Prefix}sw-kvm%`, which results in names such as dbm0sw-kvm1.

After you have entered the custom network name masks, click the **Save** button at the bottom of the page. To undo your changes, you can either close the window or click the **Cancel** button at the bottom of the page.

Admin network Review and Edit

After you save the input provided for the admin network, the Admin network Review and Edit section appears above the Private Network section. Here you can view all the IP addresses and administration names that will be configured for your rack, and modify them if needed. If you make any modifications to the administration name or their IP addresses, click the **Update** button when you have finished to save your changes.

If you are connected to the same network on which the rack will be deployed, you can check the validity of the network configuration data. Click the **Lookup IP** button at the bottom of this section. OEDA will display any invalid IP addresses. *(Not available with Exadata Cloud)*

Private Network

Provide the information necessary to configure the private network on the rack.

In the upper right corner of the Private Network section is a line that displays the Total IP Addresses that will be configured for the rack. This value is based on the hardware you selected on the Hardware page. *For Exadata Cloud, these values are provided for you.*

- **Enable active bonding** — *(Exadata and ZDLRA racks only, X4-2 or higher)*
Specifies whether active bonding should be enabled for only the private network on compute nodes. By default this option is selected.

You should only use a non-default option if you are adding a compute node to a cluster that contains other compute nodes that are X3 or earlier models.

- **Start IP Address:** Assigns the starting IP address for the private network of the servers. This field cannot be empty.

OEDA generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified during the configuration process.

- **Subnet Mask:** Sets the IP address mask for the subnet. You can choose a subnet mask from the drop-down list. The list contains supported values from 255.255.240.0 to 255.255.255.252
- **Domain Name:** Sets the network domain name. The default value is the domain name value entered for the Admin Network. Use lowercase letters for the name. This field cannot be empty.

(Not available with Exadata Cloud) If you click the **Modify Masks** button, it opens a popup window where you can modify the generated names for the database and storage server private names. There are preview fields at the bottom of this window that shows how the generated private names will appear based on the values in the name mask fields. The default strings used to generate the private names are:

- Compute private name mask: `{Rack prefix}adm%-priv`, which results in names such as `dbm0adm01-priv`, `dbm0adm02-priv`, and so on
- Cell private name mask: `{Rack Prefix}celadm%-priv`, which results in names such as `dbm0celadm01-priv`, `dbm0celadm02-priv`, and so on.

After you have entered the custom private network name masks, click the **Save** button at the bottom of the page. To undo your changes, you can either close the window or click the **Cancel** button at the bottom of the page.

After you have entered the private network configuration information, click the **Save** button at the bottom of the page.

Private network Review and Edit

After you save the input provided for the private network, the Private network Review and Edit section appears above the User(s) and Groups section. Here you can view all the IP addresses and private host names that will be configured for your rack, and modify them if needed. If you make any modifications to the private host names or

their IP addresses, click the **Update** button at the bottom of this section when you have finished to save your changes.

3.2.5 Users and Groups

Provide information about the operating system users to use when installing Oracle software. The user information you provide is saved in a collection, which can be referenced later in the configuration process.

For Exadata Cloud, the Role Separated option is selected by default and all the values are populated automatically. The only values you can change are the group ID numbers.

The first selection is the type of users you want to define. Your options are:

- **Default:** *(Not available for Exadata Cloud)* Use the default user configuration. This option creates a single operating system user to own the Oracle Grid Infrastructure and Oracle Database software.
- **Role Separated:** *(Not available for ZDLRA)* Create users for specific roles for the software installation. The new users are granted only the privileges necessary for that role.
- **Custom:** *(Exadata racks only)* Create software installation users that have combined roles or special roles and privileges.

If you plan to have multiple installations of Oracle software on your rack that require different operating system users, then, after you have entered the user information, you can click the **Add group (+)** button. The information you entered is saved with the label `UserCollection-1_id`. A new section appears on the page where you can enter information for an additional group of software installation users with distinct names and permissions. The new user information collection is automatically assigned the label `UserCollection-2_id`.

If you want to remove a user information collection, click the **Delete Group (-)** button next to the ID of the group you want to remove.

After you have entered all the user information, click the **Save** button at the lower right side of the page to continue.

Default User Information

If you selected to create Default users, then you must supply the following information:

- **User Name:** The operating system user name for the database software owner. This field cannot be empty. The default value is `oracle`.
- **User Id:** The operating system identification number for this user. This field cannot be empty. The default value is `1001`.
- **Home Directory:** The user home directory. This field cannot be empty. The default value is `/home/oracle`.
- **User Type:** This field is not available if you select the Default option. The named user will be created as the Grid and Database software owner.
- **Group Name:** The name of the operating system group that will be the DBA group. All the operating system privileges necessary for managing the Oracle Database software are granted to this group. This field cannot be empty. The default value is `dba`.

- **Group Id:** The operating system identification number for this group. This field cannot be empty. The default value is 1001.
- **Group Type:** This field is not available if you select the Default option. The named group will be created as a `DBA GROUP` group.
- **Group Name:** The name of the operating system group that will be the installation owner group. All the operating system privileges necessary for installing the Oracle Grid Infrastructure and Oracle Database software are granted to this group. This field cannot be empty. The default value is `oinstall`.
- **Group Id:** The operating system identification number for this group. This field cannot be empty. The default value is 1002.
- **Group Type:** This field is not available if you select the Default option. The named group will be created as an `OINSTALL` group.

Role Separated User Information

If you selected to create Role Separated users, then there are two types of users for which you must provide user information: `Grid` and `database`. The page is split into two sections, one for each user type. The first section defines the `Grid` user and the second section defines the `database` user.

- User Type: `Grid`

In this section you specify the user to own the Oracle Grid Infrastructure software installation. You also specify the software installation group, and 3 operating system groups used to manage Oracle ASM. All of the following settings are mandatory and must contain valid values:

- Oracle Grid Infrastructure software installation user information:
 - * **User Name:** The operating system user name for the Oracle Grid Infrastructure software owner. The default value is `grid`.
 - * **User ID:** The operating system identification number for this user. The default value is 1000.
 - * **Home Dir:** The user home directory. The default value is `/home/grid`.
- Oracle Grid Infrastructure software installation group information:
 - * **Group Name:** The name of the operating system group that owns the Oracle Inventory. The default value is `oinstall`.
 - * **Group Id:** The operating system identification number for the Oracle Inventory group. The default value is 1001.
 - * **Group Type:** The group type. The required value is `OINSTALL`, .
- Oracle ASM DBA group information:
 - * **Group Name:** The name of the operating system group that will be the Oracle ASM DBA group. The default value is `asmdba`.
Membership in this group enables access to the files managed by Oracle ASM.
 - * **Group Id:** The operating system identification number for the Oracle ASM DBA group. The default value is 1004.
 - * **Group Type:** The group type. The required value is `OSDBA`.
- Oracle ASM Operator group information:

- * **Group Name:** The name of the operating system group that will be the Oracle ASM Operator group. The default value is `asmoper`.
Membership in this group enables a limited set of Oracle ASM instance administration privileges that includes starting and stopping the Oracle ASM instance.
- * **Group Id:** The operating system identification number for the Oracle ASM Operator group. The default value is `1005`.
- * **Group Type:** The group type. The required value is `OSOPER`.
- Oracle ASM Administrator group information:
 - * **Group Name:** The name of the operating system group that will be the Oracle ASM Administrator group. The default value is `asmadmin`.
Membership in this group enables users to connect to an Oracle ASM instance as `SYSASM` using operating system authentication. The `SYSASM` privileges permit mounting and dismounting of disk groups, and other storage administration tasks. `SYSASM` privileges provide no access privileges on an Oracle Database instance.
 - * **Group Id:** The operating system identification number for the Oracle ASM Administrator group. The default value is `1006`.
 - * **Group Type:** The group type. The required value is `OSASM`.
- User Type: database

In this section you specify the user to own the Oracle Database software installation. You also specify the software installation group, and 2 operating system groups used to manage Oracle Database. You must also provide details for the Oracle ASM DBA group, which enables Oracle Database to access data files stored in Oracle ASM. All of the following settings are mandatory and must contain valid values:

- Oracle Database software installation user information:
 - * **User Name:** The operating system user name for the Oracle Database software owner. The default value is `oracle`.
 - * **User ID:** The operating system identification number for this user. The default value is `1001`.
 - * **Home Dir:** The user home directory. The default value is `/home/oracle`.
- Oracle Database software installation group information:
 - * **Group Name:** The name of the operating system group that owns the Oracle Inventory. The default value is `oinstall`.
 - * **Group Id:** The operating system identification number for the Oracle Inventory group. The default value is `1001`.
 - * **Group Type:** The group type. The required value is `OINSTALL`.
- Oracle Database DBA group information:
 - * **Group Name:** The name of the operating system group that will be the Oracle Database DBA group. The default value is `dba`.
 - * **Group Id:** The operating system identification number for the Oracle Database DBA group. The default value is `1002`.

- * **Group Type:** The group type. The required value is `DBA GROUP`.
- Oracle Database Operator group information:
 - * **Group Name:** The name of the operating system group that will be the Oracle Database Operator group. The default value is `racoper`.
Membership in this group provides users with a limited set of Oracle Database instance administration privileges that includes starting and stopping the Oracle RAC database instances.
 - * **Group Id:** The operating system identification number for the Oracle RAC Operator group, or other group type that you select. The default value is `1003`.
 - * **Group Type:** The group type. The required value is `OPER GROUP`.
- Oracle ASM DBA group information:

Ensure that you match the information that you supplied in the Grid user section.

 - * **Group Name:** The name of the operating system group that will be the Oracle ASM DBA group. The default value is `asmdba`.
Membership in this group enables access to the files managed by Oracle ASM.
 - * **Group Id:** The operating system identification number for the Oracle ASM DBA group. The default value is `1004`.
 - * **Group Type:** The group type. The required value is `OSDBA`.

Custom User and Group Information

If you selected the Custom option, then you must provide user and group information for managing the Oracle software installations. You can specify multiple users. You provide information for the software administration operating system users and groups.

For each user you must provide the following information:

- **User Name:** The operating system user name that will own the Oracle software. This field cannot be empty. The default value is `oracle`.
- **User Id:** The operating system identification number for this user. This field cannot be empty.
- **Home Directory:** The user home directory. This field cannot be empty. The default value is `/home/username`.
- **User Type:** The type of software user. You can choose from the following values:
 - Grid and Database: The user will have permissions for both the Oracle Database and the Oracle Grid Infrastructure software installations.
 - Grid: The user will have permissions for the Oracle Grid Infrastructure software installation.
 - Database: The user will have permissions for the Oracle Database software installation.
 - custom: The user type does not fit a pre-defined model.
- **Group Name:** The name of the operating system group to which the user will be assigned. This field cannot be empty.

- **Group Id:** The operating system identification number for the operating system group. This field cannot be empty.
- **Group Type:** The group type. You can choose from the following list of pre-defined group types:
 - OINSTALL
 - OSDBA
 - OSOPER
 - OSASM
 - ORADBA
 - DBA GROUP
 - OPER GROUP
 - ASMDBA
 - ASMOPER
 - ASMAADMIN

After you have entered the user and group information, you can click the **Add User (+)** button to add a new section. You can also click the **Add Group (+)** button to add an additional operating system group for the specified user.

Save

After you have entered all the information for the racks being configured, click the **Apply** button. Be sure to go to the **Menu** and select **Save**.

3.2.6 Define Clusters

More than one cluster can be defined for a hardware racks. Larger environments support multiple clusters. Each cluster requires a minimum of two physical or virtual compute nodes, and three storage cells. You cannot mix physical and virtual compute nodes in the same cluster.

If you selected **All Linux Ovm** on the Operating System page and you have specified multiple clusters, then the compute nodes can appear in more than one cluster and OEDA will create a guest for each cluster on the node. Each cluster is independent and can have different software versions.

For Exadata Cloud configurations, there is one cluster created. You can specify the name of the cluster and the usable space for the cluster, but all other details are not modifiable.

Provide the information for the first cluster and assign nodes to the cluster using the **Add Node (>)** button to move selected nodes or the **Add All Nodes (>>)** button to assign all available nodes to the cluster. If you want to add a second cluster, then click the **Add Cluster (+)** button next to the cluster name tab.

Provide the following information for each cluster in the rack:

- **Cluster Name:** Is the cluster name, which must be unique across your enterprise. The cluster name is case-insensitive, must be at least one character long and no more than 15 characters in length, must be alphanumeric, cannot begin with a

numeral, and may contain hyphens (-). Underscore characters (_) are not allowed. This field cannot be empty.

- **Inventory Location:** Is the Oracle inventory (`oraInventory`) directory path. The default value is `/u01/app/oraInventory`.
- **Grid Home Owner:** Use the drop-down list to select the collection of user information previously created on the Users page. By default, the first user collection is selected.
- **Grid Infrastructure Home Location:** Is the Oracle Grid Infrastructure directory path, which typically includes the software release number. The default value is `/u01/app/release_number/grid`, for example, `/u01/app/18.2.0.0/grid`. This value is automatically set depending on the value selected for Cluster Version. You can overwrite the default value.
- **Cluster Version:** After you select the nodes to assign to this cluster, you will be able to select the cluster version from the drop-down list. The list contains all supported versions of the Oracle Grid Infrastructure. Choose any of the supported values.
- **Base Directory:** Is the base directory for the Oracle Grid Infrastructure software. The default value is `/u01/app/grid`.

After you assign nodes to the cluster, you can modify the default values for each cluster, if needed. The following customer information fields appear below the Available Machines list:

- **Customer Name:** Is the customer name. The default value is the value you provided earlier on the Hardware page in the Customer Details section.
- **Application:** Is the application name or description. The default value is the value you provided earlier on the Hardware page in the Customer Details section.
- **Region:** Is the geographical location of the rack. The default value is the value you provided earlier on the Hardware page in the Customer Details section.
- **Time zone:** Is the time zone in which the rack is located. The default value is the value you provided earlier on the Hardware page in the Customer Details section.
- **DNS Servers:** Are the IP addresses for the DNS servers. The default values are the values you provided earlier on the Hardware page in the Customer Details section.
- **NTP Servers:** Are the IP addresses for the NTP servers. The default values are the values you provided earlier on the Hardware page in the Customer Details section.

The following additional information fields help allocate resources across multiple clusters:

- **Cores:** Is the number of cores to use for each guest VM for this cluster. The minimum value is 2 cores per guest. This field cannot be empty.
- **vCPU:** Is the number of virtual CPUs for each guest VM in this cluster. The minimum value is 4 vCPUs per guest. This is not modifiable, and is automatically set to twice the core count.
- **Memory:** Is the amount of memory in GB allocated to the guest VMs, or user domains (`domU`), for this cluster. This field cannot be empty.

If you assign Oracle VMs to your cluster, then you need to provide the following information:

- **Guest Image Version:** Is the Oracle Exadata System Software image version for guest VMs. Use the drop-down list to select the value. Selecting the `default` option sets the guest image version to match the image version of `dom0`.

 **Note:**

The guest image version must be the same as or an earlier version than the image version of the management domain (`dom0`).

3.2.7 Diskgroups

Define the Oracle ASM disk groups and their sizes for each cluster in the rack. *This page is not available for Exadata Cloud.*

If you specified more than one cluster, use the tabs at the top of this page to select each cluster and provide the disk group information.

ZDLRA Racks

For ZDLRA racks, the following information is used:

- **Diskgroup name:** Automatically set to CATALOG and DELTA
- **Diskgroup Type:** Automatically set to CATALOG and DELTA
- **Redundancy:** The CATALOG disk group is configured to use NORMAL redundancy. For the DELTA disk group, you can choose either HIGH or NORMAL (default).
- **Size:** The CATALOG disk group has a size of Default, and the DELTA disk group has a size of 95%.
- **Size Type:** The disk group sizes are specified as percent.

After you have entered the values for the disk groups and clicked Save, the following fields are then populated with the estimated values:

- **CATALOG disk group** — The usable space and raw size are set to default.
- **DELTA disk group** — The values depends on the hardware selected, for example, with NORMAL redundancy selected, there is 312007GB of usable space and 624013GB raw disk size for HC 10TB disks.

Exadata Racks

For Exadata racks, you must provide the following information:

- **Diskgroup layout:** Are the ratios to use when creating the disk groups. The options are:
 - **Legacy 80% 20%** — Is useful when using external backups. The DATA disk group is assigned 80% of the available disk space and the RECO disk group is assigned 20% of the available disk space.
 - **Legacy 40% 60%** — Is useful when using internal backups. The DATA disk group is assigned 40% of the available disk space and the RECO disk group is assigned 60% of the available disk space.

- **Custom** — Select this option when you want a non-standard layout for the disk groups. For example, if you want to create additional disk groups, or use percentages that are not 20%, 40%, 60%, or 80%.
- **Diskgroup name:** Is the name for the disk group. You must specify a name for each disk group you specify. Default names are provided, such as DBFSC1, DATAC1 and RECOC1.

 **Note:**

If you selected X8 or X7 hardware, then you cannot specify any information for the DBFS disk group.

- **Type:** Is the type of data to store in the disk group. Possible values are DATA, RECO, DBFS (automatically supplied).
- **Redundancy:** Is the redundancy level for the disk group. Possible values are HIGH, NORMAL, or UNUSED.

 **Note:**

Oracle recommends the use of HIGH redundancy disk groups for mission critical applications.

- **Size:** Is the amount of available disk space to use when creating this disk group. You can enter either a percentage, such as 40%, or the actual size specified in GB or TB, such as 250GB. For the DBFS disk group, the size is set to Default.
- **Size Type:** Specifies whether the value was a percentage or actual disk size. The drop-down list contains the values percent, usable size, or slice size. You must use the same Size Type for all disk groups.

You can configure additional disk groups by clicking the **Add new diskgroup (+)** button. Enter the same information for the new diskgroup as you did for the DATA and RECO disk groups. For the Type, you must specify a different type, such as DBFS_DG. If you try to create more than one DATA or one RECO disk group, you will get an error when you click Save. If you want to remove a disk group, you can click the **remove last diskgroup (-)** button. You must have at least one DATA and one RECO disk group.

After you have entered the values for the disk groups, click **Save**. The following fields are then populated with the estimated values:


- Usable Space — The amount of usable disk space available in the disk group, accounting for redundancy, for example 204355GB.
- Raw Size — The total disk space used by the disk group, for example 613066GB.

Advanced Options

In the upper right corner of this page is a check box labeled Advanced. If you click this box, then the following options are available:

- **Enable sparse diskgroup** — Enables the creation of sparse disk groups. Sparse grid disks can be used to create a sparse disk group to store database files that will use a small portion of their allocated space. Sparse grid disks allocate space

as new data is written to the disk, and therefore have a virtual size that can be much larger than the actual physical size.

 **Note:**

To create a sparse disk group, the version of Oracle Grid Infrastructure software needs to be release 12.1.0.2 BP5 or later.

If you choose this option, a new disk group appears in the list, with a default name of SPARSE1. Also, the Type field for this disk group is set to `SPARSE`.

- **Configure Acfs** — is a button that opens a popup window where you can configure Oracle ASM Cluster File System (Oracle ACFS). Oracle ACFS provides a general purpose file system. The following information is required to configure Oracle ACFS:
 - **Enable ACFS for disk group** — you must check this option.
 - **Diskgroup** — Is the disk group to use as the Oracle ACFS disk group. The drop-down list contains the names of the disk groups specified on the Diskgroups page.

 **Note:**

You should use a different group than DATA or RECO for Oracle ACFS. Use the **Add new diskgroup (+)** button on the Diskgroups page to add another disk group.

- **ACFS Volume Name** — Is the volume name to associate with this disk group. This field cannot be empty.
- **ACFS Volume Size** — Is the size of the Oracle ACFS volume. This field cannot be empty.
- **ACFS Mount Path** — Is the mount path for the Oracle ACFS volume. This field cannot be empty.

When you have finished configuring Oracle ACFS, click the **Save** button in the ACFS Configuration window.

3.2.8 Database Home

Provide information for configuring the Oracle Database software installation.

This page is not available for Exadata Cloud configurations.

Specify how to configure the Oracle Database software installation by providing the following information:

- **DB Home Type:** The type of Oracle Database deployment — either RAC Home (for Oracle Real Application Clusters (Oracle RAC) or Oracle RAC One Node) or single instance.
- **Select Cluster:** The cluster on which to deploy the Oracle Database software. This is a drop-down list that is populated with the cluster names specified on the Clusters page. By default, the first cluster is selected.

- **Database Home Name:** A name given to the Oracle Database software installation. This name identifies the program group associated with a particular home and the Oracle services installed on this home. The Oracle home name must be between 1 to 127 characters long, and can include only alphanumeric characters and underscores. This field cannot be empty. The default value is DbHome_1.
- **Database Home Version:** The version of the Oracle Database software to install in the cluster. Choose the value from the drop-down list. The most recent version is selected by default.
- **Database Home Location:** A directory path where the Oracle Database software will be installed. This value is set automatically based on the value specified for the Database Home Version. For example, if you select version 12.1.0.2 DBBP13, then the Database Home Location is set to `/u01/app/oracle/product/12.1.0.2/dbhome_1`. After specifying the Database Home Version, you can type over the default value to use a different location, if needed. This field cannot be empty.
- **Base Directory:** The Oracle Base directory location. The Oracle Database binary home is located underneath the Oracle base path. You can use the same Oracle base directory for multiple installations of Oracle software. If different operating system users install Oracle software on the same system, then you must create a separate Oracle base directory for each installation owner. This field cannot be empty. The default value is `/u01/app/oracle`.
- **Database Home Owner:** Use the drop-down list to select the collection of user information previously created on the Users page. By default, the first user collection is selected.
- **Available Computes and Selected Computes** — If you chose Single Instance Home for the DB Home Type, then you select the database servers on which you want to install the Oracle Database software. If you selected RAC Home for the DB Home Type, then the Oracle Database software is automatically installed on every database server in the specified cluster. Use the **Add Node (>)** or **Add All Nodes (>>)** buttons to move servers from the Available Computes list to the Selected Computes list. Use the **Delete Node (>)** or **Delete All Nodes (>>)** buttons to remove servers from the Selected Computes list.

To specify multiple Oracle homes, use the **Add Database Home (+)** button at the top of this page, next to the tab for the existing Oracle home. To remove an Oracle Home, click the **X** on the tab, next to the name of the Oracle home.

When you have finished configuring the Oracle home information, click **Apply** and then **Save** from the **Menu**.

3.2.9 Database

Provide information about the databases to create in the specified database home directories.

For Exadata Cloud configurations, you can specify the name of the database. All other options are pre-configured and cannot be modified.

- **Database Name:** Is the name of the database you want to create in the database home. The database name (`DB_UNIQUE_NAME`) portion is a string of no more than 30 characters that can contain alphanumeric, underscore (`_`), dollar (`$`), and pound (`#`) characters, but must begin with an alphabetic character.

No other special characters are permitted in a database name. The `DB_NAME` parameter for a database is set to the first 8 characters of the database name. The domain portion of the global database name (`DB_DOMAIN`) uses the value you provided for the domain when configuring the networks. The values for `DB_UNIQUE_NAME.DB_DOMAIN` in its entirety must be unique within the enterprise. This field cannot be empty.

- **Database Home:** Is the name of the Database Home in which you want to create the database. This value is automatically populated with the first home specified on the Database Home page.

Use the drop down list to select the database home name if you specified multiple database homes on the Database Home page.

- **Database Type:** (Exadata racks only) Is the type of database to create in the database home selected. If the database software was installed with the type Oracle Real Application Clusters (Oracle RAC) (as specified on the Database Home page), then you can choose to create either a Oracle RAC database or an Oracle RAC One Node database. If you selected a DB Home Type of single instance for the database home, then the Single Instance DB option is selected automatically and cannot be changed.
- **Database Owner:** Use the drop-down list to select the collection of user information previously created on the Users page. By default, the first user collection is selected.
- **Character Sets:** (Exadata racks only) Specify the character set to use with the database. Use the drop-down list to select one of the available values. Character sets are the encoding schemes used to display characters on your computer screen. The character set determines what languages can be represented in the database. The recommended value is Unicode AL32UTF8.
- **Database Template:** (Exadata racks only) Is the template to use when configuring the database. Preconfigured databases are based on templates that Oracle provides or that you create. Each Oracle-provided template is optimized for a particular workload type. The two options are online transaction processing (OLTP) or data warehousing (DW).
- **Block Size:** (Exadata racks only) Is the database block size. The default value is 8192. You can use the drop-down list to select a different value. Supported values range from 8KB to 32KB.
- **DATA Diskgroup:** (Exadata racks only) Is the name of the Oracle ASM disk group to use when creating data files for the database. You can use the drop-down list to select one of the disk groups you configured on the Diskgroups page. The disk group selected must have disk group type set to DATA.
- **RECO Diskgroup:** (Exadata racks only) Is the name of the Oracle ASM disk group to use when creating the recovery area for the database. You can use the drop-down list to select one of the disk groups you configured on the Diskgroups page. The disk group selected must have disk group type set to RECO.
- **Enable CDB:** (Exadata racks only) Select this check box to create the database as a Oracle Multitenant container database (CDB) with a single pluggable database (PDB). See "Deciding Between Multitenant Container Databases and Non-CDBs in Oracle RAC" in [Oracle Real Application Clusters Installation Guide for Linux and UNIX](#) for information about whether to select this option. If you select this option, then you must provide the following additional information:
 - **PDB Name:** Specify the name of the PDB to create in the container database.

You can use the **Add Pdb Name (+)** button to add additional PDBs to the configuration. If you want to remove a PDB, you can use the **Remove Pdb Name (-)** button that appears to the right of the PDB name you want to remove.

- **Available Computes** and **Selected Computes** — (Exadata racks only) If you chose RAC One Node DB or Single Instance DB for the Database Type, then you select the database servers on which you want to create the database instance. Selecting one node deploys Oracle RAC One Node on a single node. Oracle recommends that you select all nodes in the cluster to which you want Oracle RAC One Node to be able to fail over.

Use the **Add Node (>)** or **Add All Nodes (>>)** buttons to move servers from the Available Computes list to the Selected Computes list. Use the **Delete Node (>)** or **Delete All Nodes (>>)** buttons to remove servers from the Selected Computes list.

When you have finished configuring the database configuration information, click **Save** from the **Menu**.

3.2.10 Cluster Networks

For Exadata racks and Exadata Cloud configurations, configure the client and backup networks for the rack.

For ZDLRA racks, configure the Ingest and Replication networks.

Applications access the database servers or compute nodes through the client Ethernet network using Single Client Access Name (SCAN) and Oracle RAC Virtual IP (VIP) addresses. Oracle databases running in Oracle VM guests on Oracle Exadata Database Machine are accessed through the client Ethernet network as defined in this section. Client network configuration in both the management domain (dom0) and user domains (domUs) is done automatically when the OEDA installation tool creates the first user domain during initial deployment.

In the upper right corner of the Cluster Networks page is a line that displays the Total IP Addresses that will be configured for the rack. This value is based on the hardware you selected on the Hardware page. At the top of the section are tabs, each designating one of the clusters you specified on the Clusters page.

Advanced

At the very top of the Cluster Networks section, before any network configuration details, there is a button for **Advanced** configuration. If you click this button, the Advanced options window appears. Depending on the system configuration entries that you have made so far, you may see the following options:

- **Enable Vlan:** Select this option if there is a need for virtual deployments on the rack to access additional virtual local area networks (VLANs) on the network, such as enabling network isolation across user domains.

If you select this option, then in the client or ingest network configuration details, an additional field appears:

- **Vlan:** Specify the VLAN identifier. The VLAN ID is a numeric value between 1 and 4095. You also need to select the `Bonded` option to enable VLAN. If the Backup or Replication network uses VLAN, then the Client or Ingest network must also be configured to use VLAN. If the VLAN ID field is empty, then VLAN is not enabled.

 **Note:**

For all networks, you are responsible for configuring VLAN on the switch.

- **Enable Secure Fabric:** Select this option to configure Exadata Secure RDMA Fabric Isolation. Selecting this option implements network separation between different clusters on Oracle Exadata Database Machine systems that use RoCE Network Fabric.
- **Enable InfiniBand Partitioning:** Enable InfiniBand partitioning for Exadata systems that use InfiniBand Network Fabric. InfiniBand partitions can be defined on a per-cluster basis. If storage servers are shared among multiple clusters, then all clusters will use the same storage partition key. This option is automatically selected for Exadata Cloud configurations.

Admin Network

For some Exadata rack configurations, at the top of the Cluster Networks section, on the right side, is a check box **Enable Admin network**. You can use this to optionally configure an Admin network inside the virtual compute servers. If you select this option, you must specify the network configuration by using the following:

- **Default gateway for database servers:** Select this option if the gateway IP address is the default gateway for all database servers
- **Defines Domain hostname for database servers:** Select this option if you want the name used for the Admin network on a virtual compute server to be used as the host name for the virtual compute server. If you do not select this option, then the virtual compute server uses the network name assigned to the client network as the host name.
- **Gateway:** Sets the IP address for the gateway for the client network. This field cannot be empty.
- **Start IP Address:** Specify the first IP address that will be used for client access to the database servers. This value is filled in automatically, based on the value supplied for the gateway. OEDA generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified using the **Modify Masks** button. This field cannot be empty.
- **Subnet Mask:** Specify the IP address mask for the subnet. You can choose a subnet mask from the drop-down list which contains supported values from 255.0.0.0 to 255.255.255.252.
- **Domain Name:** Specify the network domain name for client connections. Use lowercase letters for the name. This field cannot be empty. For Exadata Cloud configurations, this value is provided for you.

After you have specified all the configuration details for the Admin network, click the **Apply** button at the bottom of the section, or click **Modify Masks** to customize the network configuration.

Modify Masks

If you click the **Modify Masks** button, it opens a pop-up window where you can modify the generated names for the client or ingest network. There are preview fields at the bottom of each section that show how the generated network names will appear

based on the values in the name mask fields. The default strings used to generate the network names are:

- Client Access Details:
 - **Name:**
 - * Exadata racks: Enter `{Rack prefix}db%%`, which results in names such as `dbm0db01`, `dbm0db02`, and so on
 - * ZDLRA racks: Enter `{Rack prefix}ingest%%`, which results in names such as `{Rack prefix}ingest01`, `{Rack prefix}ingest02`, and so on
 - **Start id:** Specify the number to start with when generating sequential names.
- Vip details:
 - **Name:**
 - * Exadata racks: Enter `{Rack prefix}db%%-vip`, which results in names such as `dbm0db01`, `dbm0db02`, and so on
 - * ZDLRA racks: Enter `{Rack prefix}ingest%%-vip`, which results in names such as `{Rack prefix}ingest01-vip`, `{Rack prefix}ingest02-vip`, and so on
 - **Start id:** Specify the number to start with when generating sequential names.
- Scan details: The configuration details for the single client access network (SCAN) listener
 - **Name:** Enter the name of the listener. The default value is `{Rack prefix}-scan1`.
 - **Port:** Specify the network port which the SCAN listener monitors for incoming connection requests. The default value is 1521.

After you have completed the modifications to the client or ingest network configuration, click the **Apply** button at the bottom of the Client Network Masks window. To undo your changes, you can either close the window or click the **Cancel** button at the bottom of the Client Network Masks window.

Client Network or Ingest Network

Provide the following information to configure either the Client networks for Exadata or the Ingest networks for ZDLRA racks.

For Exadata racks and ZDLRA racks (*but some not for Exadata Cloud*):

- **Default gateway for database servers:** Select this checkbox if the gateway IP address is the default gateway for all database or compute servers.
- **Default hostname for database servers:** Select this checkbox if the network names define the host names for the database or compute servers.
- **Network Media:** Select the network cable type, either Optical or Copper.
- If you select one of the SFP28 network types, then you must additionally specify whether to use 10 Gbit or 25 Gbit connections.

Client (Ingest) network format: (For X7 or later racks for Exadata and ZDLRA) Specifies the Ethernet cards, based on the hardware selected on the first page and the type of cable being used.

Options for X7-2 and Later Systems:

On X7 or later racks for Exadata and ZDLRA, there are 3 available formats for client networks or ingest networks:

- RJ45/SFP28 Combined on Motherboard
 - * RJ45 10 Gbit
 - * SFP28 10 Gbit
 - * SFP28 25 Gbit
- SFP28 PCI Dual Port Card
 - * SFP28 10 Gbit
 - * SFP28 25 Gbit
- SFP28 2nd PCI Dual Port Card (extra cost option on X8 systems, not available my default)
- RJ45 PCI Quad Port Card (extra cost option, not available my default)

Options for Exadata X7-8 or Later Systems:

On Exadata X7-8 or later systems, there are 2 available formats for client and backup networks:

- RJ45 on Motherboard
- SFP28 PCI Dual Port Cards
 - * SFP28 10 Gbit
 - * SFP28 25 Gbit

- **Bonded:** For X7 or later racks and Exadata Cloud, only bonded networks are supported.

For X6 and earlier racks of Exadata and ZDLRA, or Exadata Cloud, select this option if using a bonded network.

- **LACP:** Check this option to enable the Link Aggregation Control Protocol (LACP) on the network.
- **Gateway:** Sets the IP address for the gateway for the network. This field cannot be empty.
- **Start IP Address:** Set the first IP address that will be used for access to the database servers on this network. This value is filled in automatically, based on the value supplied for the gateway. OEDA generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified using the **Modify Masks** button. This field cannot be empty.
- **Subnet Mask:** Set the IP address mask for the subnet. You can choose a subnet mask from the drop-down list which contains supported values from 255.0.0.0 to 255.255.255.252.
- **Domain Name:** Set the network domain name for client or ingest connections. Use lowercase letters for the name. This field cannot be empty. For Exadata Cloud configurations, this value is provided for you.

After you have specified all the configuration details for the client or ingest network configuration, click the **Apply** button at the bottom of the section.

Client network or Ingest network Review and Edit

After you save the input provided for the client or ingest network, the Client network Review and Edit or Ingest network Review and Edit section appears. Here you can view all the host names and IP addresses that will be configured for client access to the database servers in the rack. You can modify any of the settings if needed. If you make any modifications to the administration name or their IP addresses, click the **Apply** button when you have finished to save your changes.

If you are connected to the same network on which the rack will be deployed, you can check the validity of the network configuration data. Click the **Lookup IP Address** button at the bottom of this section. OEDA will display any invalid IP addresses.

Private Network

In the upper right corner of the Private Network section is a line that displays the Total IP Addresses that will be configured for the rack for this section. At the top of the section are tabs, each designating one of the clusters you specified on the Clusters page.

If you have not selected any options in the Advanced options window, then you provide the following information to configure the private network:

- **Start IP Address:** Set the starting IP address for the cluster's network. Oracle recommends using the default IP address.
- **Name Mask:** (Not available for Exadata Cloud) Define the mask used to generate network names for this cluster. The default value is `{Rack Prefix}adm%-priv`, which results in names such as `dbm0adm01-priv`, `dbm0adm02-priv`, and so on.
- **Subnet Mask:** (Not available for Exadata Cloud) Set the subnet mask for the private network. The drop-down list contains values from `255.255.240.0` to `255.255.255.252`.
- **Domain Name:** (Not available for Exadata Cloud) Set the domain name for the private network, if it is different.
- **Start Id:** (Not available for Exadata Cloud) Set the starting identifier for this cluster. The default value is 1.

If you selected the option to **Enable Secure Fabric** in the Advanced options window, then you provide the following information to configure the private network:



Note:

Ensure that there is no overlap between the cluster VLAN subnets, and that the storage VLAN subnet does not overlap with any of the cluster VLAN subnets.

- Cluster vlan:
 - **Vlan:** Set the VLAN ID for the cluster VLAN.
Each cluster must have a unique VLAN ID, which is different from the storage VLAN ID. The VLAN ID must be in the range: 2744-3967.
 - **Domain Name:** Set the domain name for the cluster VLAN, if it is different.

- **Start IP Address:** Set the starting IP address for the cluster VLAN. Oracle recommends using the default IP address.
- **Subnet Mask:** Set the subnet mask for the cluster VLAN. The drop-down list contains values from 255.255.240.0 to 255.255.255.252.
- **Name Mask:** Define the mask used to generate network names in the cluster VLAN. The default value is `{Rack Prefix}%vmclu%-priv`, which results in names such as `dbm001vmclu1-priv`, `dbm002vmclu1-priv`, and so on.
- **Start Id:** Set the starting identifier to generate network names in the cluster VLAN. The default value is 1.
- Storage vlan:
 - **Vlan:** Set the VLAN ID for the storage VLAN.
The storage VLAN ID must be different from all of the cluster VLAN IDs. The VLAN ID must be in the range: 2744-3967.
 - **Domain Name:** Set the domain name for the storage VLAN, if it is different.
 - **Start IP Address:** Set the starting IP address for the storage VLAN. Oracle recommends using the default IP address.
 - **Subnet Mask:** Set the subnet mask for the storage VLAN. The drop-down list contains values from 255.255.240.0 to 255.255.255.252.
 - **Compute Name Mask:** Define the mask used to generate compute node network names in the storage VLAN. The default value is `{Rack Prefix}db%vm%clu1-priv`, which results in names such as `dbm0db01vm1clu1-priv`, `dbm0db02vm1clu1-priv`, and so on.
 - **Start Id:** Set the starting identifier to generate compute node network names in the storage VLAN. The default value is 1.
 - **Storage Name Mask:** Define the mask used to generate storage server network names in the storage VLAN. The default value is `{Rack Prefix}cell%vmstr1-priv`, which results in names such as `dbm0cell101vmstr1-priv`, `dbm0cell102vmstr1-priv`, and so on.
 - **Start Id:** Set the starting identifier to generate storage server network names in the storage VLAN. The default value is 1.

If you selected the option to **Enable InfiniBand Partitioning** in the Advanced options window, then you provide the following information to configure the private network:

- Cluster pkey:
 - **Pkey:** (Not available for Exadata Cloud) Set the address for the database servers or compute nodes. This address is used by all database servers or compute nodes in the cluster. The first cluster address is set to `0xa010`, the second cluster address is set to `0xa011`, and so on.

 **Note:**

Each cluster must have a unique address.

By default, OEDA deployment with pkeys will set the default membership in the InfiniBand fabric to limited. However, you can choose to have both full and limited membership, or shared membership.

- **Domain Name:** (Not available for Exadata Cloud) Set the domain name for the InfiniBand network, if it is different.
- **Start IP Address:** Set the starting IP address for the cluster's network. Oracle recommends using the default IP address.
- **Subnet Mask:** (Not available for Exadata Cloud) Set the subnet mask for the InfiniBand network. The drop-down list contains values from 255.255.240.0 to 255.255.255.252.
- **Name Mask:** (Not available for Exadata Cloud) Define the mask used to generate network names for this cluster. The default value is `{Rack Prefix}adm%-priv`, which results in names such as `dbm0adm01-priv`, `dbm0adm02-priv`, and so on.
- **Start Id:** (Not available for Exadata Cloud) Set the starting identifier for this cluster. The default value is 1.
- **Storage pkey** (Not available for Exadata Cloud):
 - **Pkey:** Set the address for the storage servers or cells. This address is used by all storage servers or cells in the rack. The default value is `0xaa10`.
 - **Domain Name:** Set the domain name for the InfiniBand network for the storage servers or cells, if it is different.

Private network Review and Edit

Here you can view all the host names and IP addresses that will be configured for the private network. You can modify any of the settings if needed. If you make any modifications, then when you are finished click the **Update** button to save your changes.

Backup Network or Replication Network

In the upper right corner of the Backup Network section for Exadata is the option **Enable Backup network**. For ZDLRA racks, this option is labeled **Enable replication network**. If you select this option, it enables the configuration of a network dedicated to backups or replication.

The **Enable Backup network** option is automatically selected for Exadata Cloud.

In the upper right corner of the Backup Network or Replication Network section is a line that displays the Total IP Addresses that will be configured for each cluster. At the top of the section are tabs, each designating one of the clusters you specified on the Clusters page.

Provide the following information to configure the backup network or replication network for each cluster:

- **Shared client network port:** (Not available for Exadata Cloud) Select this option if the client and backup networks use the same physical network cables. This option requires that both networks have VLAN enabled and share the same network interface cards (NIC). If this option is not selected, then you must use one set of cables for the client network and a separate set of cables for the backup network.

- **Network Media:** (For Exadata X6 racks and earlier and ZDLRA racks) Select the network cable type, either Optical or Copper.
- **Backup network format:** (For X7 or later racks for Exadata and ZDLRA only) (Not available for Exadata Cloud) Specify the Ethernet cards, based on the hardware selected on the first page and the type of cable being used.

Options for X7 or later racks for Exadata and ZDLRA:

There are 3 available formats for client networks.

- RJ45/SFP28 Combined on Motherboard
 - * RJ45 10 Gbit
 - * SFP28 10 Gbit
 - * SFP28 25 Gbit
- SFP28 PCI Dual Port Card
 - * SFP28 10 Gbit
 - * SFP28 25 Gbit
- RJ45 PCI Quad Port Card (extra cost option)

Options for Exadata X7-8 or Later Racks:

There are 2 available formats for client and backup networks:

- RJ45 on Motherboard
- SFP28 PCI Dual Port Cards
 - * SFP28 10 Gbit
 - * SFP28 25 Gbit

If you select one of the SFP28 network types on X7 or later racks, then you must additionally specify whether to use 10 Gbit or 25 Gbit connections.

- **Bonded:** (Not available on X7 or later racks or Exadata Cloud) Select this option to use network bonding. For X7 or later racks for Exadata and ZDLRA, and Exadata Cloud, only bonded networks are supported.
- **LACP:** Check this option to enable the Link Aggregation Control Protocol (LACP) on the network.
- **Gateway:** Set the IP address for the gateway for the network. This field cannot be empty.
- **Start IP Address:** Set the first IP address that will be used for the network. This value is filled in automatically, based on the value supplied for the gateway. OEDA generates a pool of consecutive IP addresses based on the specified Start IP Address. If consecutive IP addresses are not available, then specific IP addresses can be modified using the **Modify Masks** button. This field cannot be empty.
- **Subnet Mask:** Set the IP address mask for the subnet. You can choose a subnet mask from the drop-down list which contains supported values from 255.0.0.0 to 255.255.255.252.
- **Domain Name:** (Not available for Exadata Cloud) Set the network domain name for network connections. Use lowercase letters for the name. This field cannot be empty.

After you have specified all the configuration details for the network configuration, click the **Apply** button at the bottom of the section.

(Not available for Exadata Cloud) If you click the **Modify Masks** button, it opens a pop-up window where you can modify the generated names for the network. There is a preview field at the bottom of the window that shows how the generated network names will appear based on the values in the name mask fields. The default strings used to generate the network names are:

- **Client Access Details:**
 - **Name:**
 - * Exadata racks: Enter `{Rack prefix}db%%`, which results in names such as `dbm0db01`, `dbm0db02`, and so on
 - * ZDLRA racks: Enter `{Rack prefix}ingest%%`, which results in names such as `{Rack prefix}ingest01`, `{Rack prefix}ingest02`, and so on
 - **Start id:** Set the number to start with when generating sequential names.
- **Vip details:**
 - **Name:**
 - * Exadata racks: Enter `{Rack prefix}db%%-vip`, which results in names such as `dbm0db01-vip`, `dbm0db02-vip`, and so on
 - * ZDLRA racks: Enter `{Rack prefix}ingest%%-vip`, which results in names such as `{Rack prefix}ingest01-vip`, `{Rack prefix}ingest02-vip`, and so on
 - **Start id:** Set the number to start with when generating sequential names.
- **Scan details:** Are the configuration details for the single client access network (SCAN) listener.
 - **Name:** Specify the name of the listener. The default value is `{Rack prefix}-scan1`.
 - **Port:** Specify the network port which the SCAN listener monitors for incoming connection requests. The default value is 1521.

(Not available for Exadata Cloud) After you have completed the modifications to the backup or replication network configuration, click the **Apply** button at the bottom of the Backup Network Masks or Replication Network Masks page. To undo your changes, you can either close the window or click the **Cancel** button at the bottom of the window. The buttons might be out of view and require scrolling down in the window to see.

After you have specified all the configuration details for the backup or replication network configuration, click the **Apply** button at the bottom of the Backup Network or Replication Network section.

Backup or Replication network Review and Edit

If you enabled the Backup network or the Replication network option, then this section appears after you save the input provided for the network. Here you can view all the host names and IP addresses that will be configured for the network. You can modify any of the settings if needed. If you make any modifications, then click the **Update** button when you have finished to save your changes.

(Not available for Exadata Cloud) If you are connected to the same network on which the rack will be deployed, you can check the validity of the network configuration data. Click the **Lookup IP Address** button at the bottom of this section. OEDA will display any invalid IP addresses.

Related Topics

- [Enabling 802.1Q VLAN Tagging in Exadata Database Machine over client networks \(My Oracle Support Doc ID 1423676.1\)](#)
- [Implementing Tagged VLAN Interfaces in Oracle VM Environments on Exadata \(My Oracle Support Doc ID 2018550.1\)](#)

3.2.11 Alerting

In this section you configure alerting (*Not available for Exadata Cloud*).

You use the following methods for receiving or handling alerts:

- Email with Simple Mail Transfer Protocol (SMTP)
- Simple Network Management Protocol (SNMP)
- Oracle Auto Service Request (ASR)
- Oracle Configuration Manager
- Oracle Enterprise Manager

The alerting options include:

- **Enable Email Alerting:** Enables alerts to be delivered automatically through email.
 - **Name:** Sets the SMTP email user name that appears in the alert notifications, for example, `Oracle Exadata Database Machine`. This field cannot be empty.
 - **Email Address:** Sets the SMTP email address that sends alert notifications, for example, `dm0db01@example.com`. This field cannot be empty.
 - **SMTP Server:** Sets the SMTP email server used to send alert notifications, for example, `mail.example.com`. This field cannot be empty.
 - **Port:** Sets the SMTP email server port used to send alert notifications, for example, 25 or 465. This field cannot be empty.
 - **Send Email Alerts to:** Sets the email address for the recipients of the alerts. You can click the **Add Email (+)** button to add additional recipients. You can use the **Remove Email (-)** button to remove an email address that you added.
- **Enable SNMP Alerting:** Enables using SNMP to deliver alerts.
 - **SNMP Server:** Sets the SNMP server host name, for example, `snmp.example.com`. This field cannot be empty.

 **Note:**

Additional SNMP targets may be defined after installation of the Exadata rack. Refer to the `ALTER CELL CellCLI` command or the `ALTER DBSERVER DBMCLI` command.

- **Port:** Sets the SNMP email server port used to send alert notifications. The default value is 162. This field cannot be empty.
- **Community:** Sets the SNMP server community. The default value is `public`. This field cannot be empty.

- **Enable ASR Alerting:** Enables Oracle Auto Service Request (ASR). Oracle ASR is designed to automatically open service requests when specific rack hardware faults occur.

Oracle recommends that you install Oracle ASR Manager on an external, standalone server. The server must be installed have connectivity to the Exadata servers or Recovery Appliance nodes and an outbound Internet connection using HTTPS or an HTTPS proxy. This server must run an Oracle Solaris or Linux operating system.

- **ASR Manager Request:** The host name of the Oracle ASR Manager server. This field cannot be empty.
 - **HTTP Proxy Host:** If you want to enable an HTTP proxy to upload to Oracle ASR, then specify the proxy server host name.
 - **Contact Name:** The name of the technical contact person for the rack. This field cannot be empty.
 - **Contact Email:** The e-mail address of the technical contact for the rack. This field cannot be empty.
 - **HTTP Proxy User:** If you want to enable an HTTP proxy to upload to Oracle ASR, then specify the proxy server user name.
 - **Port:** If you want to enable an HTTP proxy to upload to Oracle ASR, then specify the HTTP proxy server port number.
 - **My Oracle Support Account Name:** The email address of the My Oracle Support account. This field cannot be empty.
- **Enable OCM Alerting:** Enables Oracle Configuration Manager to collect configuration information.

Oracle Configuration Manager is used to collect configuration information and upload it to an Oracle repository, for access when needed by Oracle Support Services.

- **MOS Email Address:** The My Oracle Support account (email address) to receive updates from My Oracle Support. This field cannot be empty.
- **HTTP Proxy Host:** If you use an HTTP proxy to upload configuration information to the Oracle repository, then specify the proxy server host name. The proxy server may be between the following:
 - * Database (compute) servers and Oracle (applicable when you do not have Oracle Support Hub)
 - * Database (compute) servers and Oracle Support Hub (applicable when you have only Oracle Support Hub)
 - * Oracle Support Hub and Oracle

If passwords are needed, then provide them during installation.

- **Support Hub Hostname:** The Support Hub server host name. Oracle Support Hub allows Oracle Configuration Manager instances to connect to a single internal port (the Oracle Support Hub), and upload configuration data. This eliminates the need for each individual Oracle Configuration Manager instance in the database servers to access the Internet. This field cannot be empty.
- **HTTP Proxy User:** If you want to enable an HTTP proxy to upload configuration information to the Oracle repository, then specify the HTTP proxy server user name.

- **Hub User Name:** The Oracle Support Hub operating system user name. This field cannot be empty.
- **Port:** If you want to enable an HTTP proxy to upload configuration information to the Oracle repository, then specify the HTTP proxy server port number. The default value is 80.
- **Enable EM Alerting:** Enables alerts in Oracle Enterprise Manager.
Oracle Enterprise Manager Cloud Control or Oracle Enterprise Manager Grid Control is used to monitor the rack. When Oracle Enterprise Manager agents are installed on the rack, they can send information to the Oracle Management Server. If you do not have an existing Oracle Management Server, then install a new Oracle Management Server on a separate host.
 - **EM Home Base Location:** The base directory for the Oracle Enterprise Manager agents installed on the database or compute servers. The default value is `/u01/app/oracle/product/EMbase`. This field cannot be empty.
 - **OMS HTTPS Upload Port:** The port the Oracle Enterprise Manager agent should use to upload information to the Oracle Management Server. The default value is 4900. This field cannot be empty.
 - **OMS Host Name:** The name of the server or host where the Oracle Management Server repository is installed. This field cannot be empty.

3.2.12 Tape Library Connectivity

On ZDLRA racks only, you can configure network connection information for tape backups.

You can configure the tape library for ZDLRA only if you use Oracle Secure Backup as the media manager for tape backups. You can manually configure other media managers as clients only.

- **Use Tape:** Select this option if you plan to use Oracle Secure Backup as the media manager for tape backups. Then you can enter the configuration details.
- **Vip IP Address:** The IP address of the VIP.
- **Vip Name:** The alphanumeric name of the VIP. The default value is `osbadminvip.{domain_name}`.

If you are connected to the same network on which the rack will be deployed, you can check the validity of the tape library address information. Click the **Lookup IP** button at the bottom of this section. OEDA will display any invalid IP addresses.

- **Number of tape libraries:** The number of tape libraries available to the ZDLRA rack. The default is 1. All tape libraries must be hosted in the same domain. If you select a value greater than 1, a new tab appears for each tape library. You can enter the configuration details for each tape library on the separate tabs.
- **ACLS ACS API:** Use the Automated Cartridge System Library Software (ACSL) API for the tape library. Fill out the following additional fields:
 - **ACSL Server:** The name of the server running ACSL
 - **ACS ID:** The ACS identification number
- **ACSL FC-SCSI:** Use the Automated Cartridge System Library Software (ACSL) FC-SCSI for the tape library. Fill out the following additional fields:
 - **ACSL Server:** The name of the server running ACSL

- **Logical Library ID:** The logical library identification number
- **SCSI:** Use SCSI to communicate with the tape library. The Simple SAN tape library has no configuration fields. This is the default option.

3.2.13 Comments

Enter any additional information about the deployment in the text field.

For Oracle Exadata Rack, this section contains a list of questions about the configuration to help ensure a smooth installation and configuration process for your rack. You can type in the answers in the text field.

For all deployments, you can provide additional comments about your environment or the deployment in this field.

When you have finished providing the information, click **Save and Download**. You are prompted to save a ZIP file containing the deployment and configuration files. The main configuration files are:



Note:

Do not edit or modify any of the generated files.

- *Customer_name-rack_prefix.xml*: This is the main configuration file used by the installation process.
- *Customer_name-rack_prefix-checkip.sh*: A script you run on the installed rack to verify all the IP addresses you specified in OEDA are available and ready for deployment.
- *Customer_name-rack_prefix-InstallationTemplate.html*: This file is an HTML report of all the data captured in the web user interface. The questions and answers from the Comments section of OEDA appear at the bottom of this file.
- *Customer_name-rack_prefix-platinum.csv*: This file is used when first imaging of the nodes in the rack by Oracle Advanced Customer Support.
- *Customer_name-rack_prefix-preconf_pkey.csv*: If you configured InfiniBand partitioning, this file is created. This file is used when PKEY security is enabled in the interconnect.

4

OEDA Command Line Interface

When you run the `oedacli` program, it starts the OEDA command line interface.

- [About the OEDA Command Line Interface](#)
You can use the OEDA command-line interface (OEDACLI) to perform Exadata life cycle management tasks.
- [Starting the OEDACLI Utility](#)
- [OEDACLI Command Reference](#)
These OEDACLI commands are arranged by the object they act upon.

4.1 About the OEDA Command Line Interface

You can use the OEDA command-line interface (OEDACLI) to perform Exadata life cycle management tasks.

Supported tasks include:

- Add nodes to and remove nodes from a VM cluster
- Add or remove Oracle Database home directories
- Add or remove storage cells
- Resize Oracle ASM disk groups
- Add or remove additional databases

OEDACLI supports life cycle management tasks for Oracle Exadata systems only. OEDACLI does not support life cycle management operations on Oracle SuperCluster and Oracle Zero Data Loss Recovery Appliance (ZDLRA).

OEDACLI allows you to update an existing Exadata configuration file (`es.xml`). The updates are called Actions. An **Action** is a single atomic task.

An Action can have many sub commands, however most actions are single commands. Examples of multi-command steps are - [CLONE GUEST](#) and [CLONE CELL](#)

OEDACLI has two operating modes:

- Edit mode:
 - Used against an `es.xml` file *before* you run the `install.sh` command steps
 - Used to edit an XML file, for example to add a second Database Home, then save the XML file
 - Can be used as a logical XML editor for the `es` model prior to deployment
- Deploy mode:
 - Used to edit an `es.xml` file *after* you have run the `install.sh` command steps

- Can be used to Clone Oracle VM Guests or grow or shrink Oracle ASM Disk groups
- Has a smaller list of supported changes than Edit mode
- Can be used as a logical XML editor and deployment tool after the initial deployment

Edit Mode Workflow

1. When the command line interface starts, it has no target `es.xml` file to work on, so you first load an `es.xml` file.

```
LOAD FILE name=es.xml
```

2. Modify an in-memory version of the `es.xml`. Specify the additional actions to perform during deployment. For example, to clone an Oracle VM guest to a new guest:

```
clone guest srcname=asd
set privnet
name1='test08adm01-priv1',ip1='192.168.216.102',name2='testadm01-
priv2',ip2='192.168.216.103'
set parent name='test08adm01a'
set adminnet name='test08adm01vm02',ip='10.128.218.70'
set clientnet name='test08client01vm02',ip='10.128.196.63'
set vipnet name='test08client01vm02-vip',ip='10.128.196.71'
```

3. After you have provided all the commands for an action, save the action to a list of actions to be processed within the command line interface. At this point, the data in the `es.xml` file is not yet updated.

```
SAVE ACTION
```

4. Merge the actions into the main body of the in-memory copy of the `es.xml` using the `FORCE` keyword.

```
MERGE ACTIONS FORCE
```

5. Save the memory copy of the `es.xml`, overwriting the input file or writing to a new file that you specify.

```
SAVE FILE [name=new_es.xml]
```

Deploy Mode Workflow

1. When the command line interface starts, it has no target `es.xml` file to work on, so you first load an `es.xml` file.

```
LOAD FILE name=es.xml
```

2. Modify an in-memory version of the `es.xml`. Specify the actions to perform.

3. After you have provided all the commands for an action, save the action to a list of actions to be processed within the command line interface. At this point, the data in the `es.xml` file is not yet updated.

```
SAVE ACTION
```

4. Merge the actions into the main body of the in-memory copy of the `es.xml`.

```
MERGE ACTIONS
```

5. Deploy the actions into the main body of the `es.xml` file. This step will perform the specified actions such as creating the new Oracle VM guest.

```
DEPLOY ACTIONS
```

6. Save the memory copy of the `es.xml`, overwriting the input file or writing to a new file that you specify.

```
SAVE FILE [name=new_es.xml]
```

4.2 Starting the OEDACLI Utility

Starting with August 2017 Release of Oracle Exadata Deployment Assistant (OEDA), a command-line utility — `oedacli` — can be used for various routine tasks on Exadata.

Prerequisites

To use `oedacli`, you must meet the following requirements:

- Download the August 2017 release or later of OEDA. The August 2017 release of OEDA is only supported on Oracle Linux.
- Have the original OEDA configuration XML file used for deployment
- Use the default password on the nodes in the cluster (for the `root`, `oracle` and `grid` users)

Note:

If your environment does not use the default passwords for the specified users, then refer to My Oracle Support Doc ID 2163751.1.

File Path

The directory that contains the unzipped files from the OEDA download. For example `/u01/joc/linux-x64`.

Syntax

```
./oedacli
```

Usage Notes

To display the help output, use one of the following:

- `HELP`
- `HELP command name`
- [Displaying Help Output for OEDACLI](#)
You can display syntax help for the OEDACLI utility using the `HELP` command.
- [Displaying Help Output for an OEDACLI Command](#)
This example shows the help output listing the OEDACLI `ALTER DATABASE` command.

Related Topics

- [OEDA Deployment - Generating new or non-default hash password for root and non-root users \(My Oracle Support Doc ID 2163751.1\)](#)

4.2.1 Displaying Help Output for OEDACLI

You can display syntax help for the OEDACLI utility using the `HELP` command.

Example 4-1 Help Output for OEDACLI

This example shows the help output listing the available objects and command for the OEDACLI 18c (18.1.0) release:

```
oedacli> HELP
HELP [topic]
  Available Objects:
ACFSVOLUME
    LIST ACFSVOLUMES
CELL
    LIST CELLS
CLUSTER
    LIST CLUSTERS
DISKGROUP
    LIST DISKGROUPS
DATABASE
    LIST DATABASES
DATABASEHOME
    LIST DATABASEHOMES
DOM0
    LIST DOM0S
GUEST
    CLONE GUEST
    DELETE GUEST
    LIST GUESTS
ILOM
    LIST ILOMS
MACHINE
    LIST MACHINES
NETWORK
    LIST NETWORKS
SCAN
```

```

        LIST SCANS
SWITCH
        LIST SWITCHES
VIP
        LIST VIPS
CLI COMMANDS
        LOAD FILE
        SAVE FILE
        SAVE FILES
        LIST XMLACTIONS
        SAVE ACTION
        MERGE ACTIONS
        DEPLOY ACTIONS

```

4.2.2 Displaying Help Output for an OEDACLI Command

This example shows the help output listing the OEDACLI ALTER DATABASE command.

Example 4-2 Help Output for the OEDACLI ALTER DATABASE command

This example shows how to display the syntax for the ALTER DATABASE command of OEDACLI. The help output is shown for the 18c (18.1.0) release.

```

oedacli> HELP ALTER DATABASE
Usage:
ALTER DATABASE
  BLOCKSIZE = <blocksize> |
  CHARSET = <characterset> |
  DATADG = <datadg> |
  DBLANG = <dblang> |
  DBNAME = <dbname> |
  DBTEMPLATE = <dbtemplate> |
  DBTYPE = <dbtype> |
  HOSTNAMES = <hostnames> |
  RECODG = <recodg>
WHERE
  ID = <databaseid> |
  CLUSTERNUMBER = <clusternumber> DATABASENAME = <databasename> |
  CLUSTERNAME = <clustername> DATABASENAME = <databasename> |
  CLUSTERID = <clusterid> DATABASENAME = <databasename>

Purpose:
  Alters an existing database
Arguments:
  <blocksize>      : Default: 8192, not required for PDB
  <characterset>   : Default: AL32UTF8, not required for PDB
  <datadg>         : *The DATA diskgroup, not required for PDB
  <dblang>         : Default: all_langs, not required for PDB
  <dbname>         : *Database name
  <dbtemplate>    : Default: ADMIN, otherwise ADMIN or DW, not
required for PDB
  <dbtype>         : Default: normal DB, otherwise CDB or PDB
  <hostnames>     : Default: list of nodes from db home
  <recodg>        : *The RECO diskgroup, not required for PDB
Where:

```


<databaseid> : The es xml id of the database
<clusternumber> : The cluster number in the es xml, starting at 1
<clustername> : The name of the cluster
<clusterid> : The es xml id of the cluster
<databasename> : The name of the database in the cluster

4.3 OEDACLI Command Reference

These OEDACLI commands are arranged by the object they act upon.

- **OEDACLI Commands**
These commands are used to interact with OEDACLI, to load and save files, and to merge, save, and deploy actions.
- **ACFSVOLUME**
You can add, alter, delete, or list Oracle ACFS volumes.
- **CELL**
You can clone, delete, or list Oracle Exadata storage cells.
- **CLUSTER**
You can alter or list the Oracle RAC Clusters on the Oracle Exadata Database Machine.
- **COMPUTE**
Use these commands to add or delete a physical compute node as a cluster member.
- **DATABASE**
You can add, alter, delete, or list the Oracle databases.
- **DATABASE HOME**
You can add, alter, delete, or list the Oracle database homes.
- **DISKGROUP**
You can add, alter, delete, or list the Oracle ASM disk groups.
- **ES**
You can discover an engineered system or alter a limited number of attributes of an engineered system.
- **GUEST**
You can clone or delete an Oracle VM guest.
- **ILOM**
You can alter or list the ILOMS in the loaded es.xml file.
- **MACHINE**
You can alter or list the machines in the es.xml file.
- **NETWORK**
You can add, alter, delete, or list the networks defined in the XML configuration file.
- **SCAN**
You can add, alter, delete, or list the SCANS for a cluster.
- **SSHKEYS**
You can generate and deploy SSH keys for root user authentication on Oracle Exadata Rack.

- **SWITCH**
You can alter or list the switches.
- **VIP**
You can add, alter, delete, or list the VIPs for the clusters.

4.3.1 OEDACLI Commands

These commands are used to interact with OEDACLI, to load and save files, and to merge, save, and deploy actions.

- **LOAD FILE**
- **SAVE ACTION**
- **MERGE ACTIONS**
- **DEPLOY ACTIONS**
- **SAVE FILE**
- **SAVE FILES**
- **SET SECURITY**
This command performs security-related actions, like setting the non-default root password for the OEDACLI session or controlling password-based authentication.
- **SET SU**
This command enables and disables using `su` from the `root` user to run remote `non-root` commands from OEDACLI.
- **LIST STEPS**
This command lists the steps for a given cluster.
- **LIST XMLACTIONS**

4.3.1.1 LOAD FILE

Use this command to load an engineered system XML file (`es.xml`) into OEDACLI memory for editing.

Syntax

```
LOAD FILE NAME=filename
```

Arguments

- **NAME** : Specifies the name of an existing engineered system XML file

Example 4-3 Loading a File into OEDACLI for Editing

```
oedacli> LOAD FILE NAME=Oracle-test12.xml
SUCCESS - file loaded OK
Customer : testadm01 - Oracle
```

4.3.1.2 SAVE ACTION

Use this command to save the current action to engineered system XML memory copy.

Syntax

```
SAVE ACTION [FORCE]
```

Arguments

Specify `FORCE` if the action should be auto-merged into the engineered system model.



Note:

The engineered system configuration should be saved as a file using a subsequent `SAVE FILE` command.

4.3.1.3 MERGE ACTIONS

Use this command to merge the current actions into the other actions in the Engineered System model.

Syntax

```
MERGE ACTIONS [FORCE]
```

Arguments

Specify `FORCE` if you do not want this action to be processed by a `DEPLOY ACTIONS` command.

4.3.1.4 DEPLOY ACTIONS

Use this command to execute all actions from the Actions section of the engineered systems XML file.

Syntax

```
DEPLOY ACTIONS
```

Example 4-4 Deploying the Actions to Add a Database

```
oedacli> DEPLOY ACTIONS
Deploying Action ID : 4 add database dbname='testdb' DATADG='DATA1'
RECODG='
RECO1' where DBHOMEID='c0_databaseHome'
Deploying ADD DATABASE
Running datapatch on database 'testdb'
Done...
Done
```

4.3.1.5 SAVE FILE

Use this command to save a current set of actions to a new engineered system XML file.

Syntax

```
SAVE FILE [ NAME=filename ]
```

Arguments

The optional argument, *NAME*, is used to specify the name of a new engineered system XML file. If you do not specify the *NAME* argument, then this command overwrites the source file.

4.3.1.6 SAVE FILES

Use this command to generate a new set of configuration files from OEDA.

Syntax

```
SAVE FILES LOCATION = dirname
```

Arguments

The *LOCATION* argument specifies the directory in which the new file set will be created.

Usage Notes

If the specified directory does not exist, it will be created.

4.3.1.7 SET SECURITY

This command performs security-related actions, like setting the non-default root password for the OEDACLI session or controlling password-based authentication.

Prerequisites

Before you can use this command with the `ENABLEPASSWORDLOGIN` option, you must configure SSH key-based authentication on the target system. See [DEPLOY SSHKEYS](#).

Syntax

```
SET SECURITY PASSWORD=root_password
```

```
SET SECURITY ENABLEPASSWORDLOGIN={true|false}  
[ HOSTNAMES="host1[,host2[,host3 ...]]" ]
```

Arguments

- `PASSWORD` specifies the non-default root password for the OEDACLI session.

- `ENABLEPASSWORDLOGIN` specifies whether to enable or disable password-based authentication for the `root` user.
- `HOSTNAMES` optionally specifies the list of host names on which to act.

4.3.1.8 SET SU

This command enables and disables using `su` from the `root` user to run remote non-`root` commands from OEDACLI.

Syntax

```
SET SU ENABLE={true|false}
```

Arguments

- `ENABLE` specifies whether or not OEDACLI uses `su` from the `root` user to run remote non-`root` commands.

4.3.1.9 LIST STEPS

This command lists the steps for a given cluster.

Each step is named, for example `Create Virtual Machine`. Each step also has a one word identifier, for example, `CREATEVM`. The short name identifier is used when specifying `STEPS` or `SKIPSTEPS` for the `CREATE CLUSTER` command.

The list of valid steps for a cluster depends on the features selected for the cluster. The steps listed are valid for the identified cluster. If the XML file contains only 1 cluster, then there is no need to identify the cluster explicitly.

The step numbers listed may change depending on features implemented for the cluster, but the short name identifier is always the same for a given step.

Syntax

```
oedacli> LIST STEPS WHERE clustername = Cluster-cl
{
  "step" : [ {
    "stepName" : "Validate Configuration File",
    "shortName" : "VALIDATE",
    "id" : "1"
  }, {
    "stepName" : "Create Virtual Machine",
    "shortName" : "CREATEVM",
    "id" : "2"
  }, {
    "stepName" : "Create Users",
    "shortName" : "USERS",
    "id" : "3"
  }, {
    "stepName" : "Setup Cell Connectivity",
    "shortName" : "CELLCONN",
    "id" : "4"
  }, {
```

```
    "stepName" : "Calibrate Cells",
    "shortName" : "CALIBRATE",
    "id" : "5"
  }, {
    "stepName" : "Create Cell Disks",
    "shortName" : "CELLDISKS",
    "id" : "6"
  }, {
    "stepName" : "Create Grid Disks",
    "shortName" : "GRIDDISKS",
    "id" : "7"
  }, {
    "stepName" : "Install Cluster Software",
    "shortName" : "INSTALLGI",
    "id" : "8"
  }, {
    "stepName" : "Initialize Cluster Software",
    "shortName" : "INITGI",
    "id" : "9"
  }, {
    "stepName" : "Install Database Software",
    "shortName" : "INSTALLDB",
    "id" : "10"
  }, {
    "stepName" : "Relink Database with RDS",
    "shortName" : "RELINK",
    "id" : "11"
  }, {
    "stepName" : "Create ASM Diskgroups",
    "shortName" : "ASMDG",
    "id" : "12"
  }, {
    "stepName" : "Create Databases",
    "shortName" : "DATABASE",
    "id" : "13"
  }, {
    "stepName" : "Apply Security Fixes",
    "shortName" : "APPLYFIXES",
    "id" : "14"
  }, {
    "stepName" : "Install Exachk",
    "shortName" : "EXACHK",
    "id" : "15"
  }, {
    "stepName" : "Create Installation Summary",
    "shortName" : "INSTALLSUMMARY",
    "id" : "16"
  }, {
    "stepName" : "Resecure Machine",
    "shortName" : "RESECURE",
    "id" : "17"
  }
} ]
}
```

4.3.1.10 LIST XMLACTIONS

This command lists the current actions created by OEDACLI.

Syntax

```
LIST XMLACTIONS
```

4.3.2 ACFSVOLUME

You can add, alter, delete, or list Oracle ACFS volumes.

- [ADD ACFSVOLUME](#)
- [ALTER ACFSVOLUME](#)
- [DELETE ACFSVOLUME](#)
- [LIST ACFSVOLUME](#)

4.3.2.1 ADD ACFSVOLUME

This command adds Oracle ASM Cluster File System (Oracle ACFS) volumes.

Syntax

```
ADD ACFSVOLUME
  VOLUMENAME = volume_name |
  VOLUMESIZE = volume_size |
  MOUNTPATH = mountpath
WHERE
  DISKGROUPID = diskgroup_id |
  CLUSTERNUMBER = cluster_number DISKGROUPNAME = diskgroup_name |
  CLUSTERNAME = cluster_name DISKGROUPNAME diskgroup_name |
  CLUSTERID = cluster_id DISKGROUPNAME = diskgroup_name
```

Arguments

You can specify the following attributes for the new Oracle ACFS volume:

- **VOLUMENAME** : Specifies a unique name for the new volume
- **VOLUMESIZE** : A number that specifies the size of the new volume in GB
- **MOUNTPATH** : Specifies the file system mount path for the new volume

The following arguments are available in the **WHERE** clause:

- **DISKGROUPID** : Specifies the `es.xml` ID for the container disk group
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster
- **DISKGROUPNAME** : Specifies the name of the container disk group

4.3.2.2 ALTER ACFSVOLUME

This command alters an existing Oracle ASM Cluster File System (Oracle ACFS) volume.

Syntax

```
ALTER ACFSVOLUME
  VOLUMENAME = volume_name |
  VOLUMESIZE = volume_size |
  MOUNTPATH = mountpath
WHERE
  ID = acfs_volume_id |
  DISKGROUPID = diskgroup_id VOLUMENAME = volume_name |
  CLUSTERNUMBER = cluster_number DISKGROUPNAME = diskgroup_name
  VOLUMENAME = volume_name |
  CLUSTERNAME = cluster_name DISKGROUPNAME diskgroup_name VOLUMENAME =
  volume_name |
  CLUSTERID = cluster_id DISKGROUPNAME = diskgroup_name VOLUMENAME =
  volume_name
```

Arguments

You can modify the following attributes of the Oracle ACFS volume:

- **VOLUMENAME** : Specifies a new, unique volume name
- **VOLUMESIZE** : A number that specifies the new size of the volume in GB
- **MOUNTPATH** : Specifies the new file system mount path for the volume

The following arguments are available in the `WHERE` clause:

- **ID** : Specifies the `es.xml` ID for an existing Oracle ACFS volume
- **DISKGROUPID** : Specifies the `es.xml` ID for the container disk group
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster
- **DISKGROUPNAME** : Specifies the name of the container disk group
- **VOLUMENAME** : Specifies the name of the existing Oracle ACFS volume

4.3.2.3 DELETE ACFSVOLUME

This command deletes an existing Oracle ASM Cluster File System (Oracle ACFS) volume.

Syntax

```
DELETE ACFSVOLUME
WHERE
  ID = acfs_volume_id |
  DISKGROUPID = diskgroup_id VOLUMENAME = volume_name |
```



```

    CLUSTERNUMBER = cluster_number DISKGROUPNAME = diskgroup_name
    VOLUMENAME = volume_name |
    CLUSTERNAME = cluster_name DISKGROUPNAME diskgroup_name VOLUMENAME =
    volume_name |
    CLUSTERID = cluster_id DISKGROUPNAME = diskgroup_name VOLUMENAME =
    volume_name

```

Arguments

The following arguments are available in the WHERE clause:

- ID : Specifies the `es.xml` ID for the Oracle ACFS volume
- DISKGROUPID : Specifies the `es.xml` ID for the container disk group
- CLUSTERNUMBER : Specifies the cluster number in the `es.xml`, starting at 1
- CLUSTERNAME : Specifies the name of the cluster
- CLUSTERID : Specifies the `es.xml` ID of the cluster
- DISKGROUPNAME : Specifies the name of the container disk group
- VOLUMENAME : Specifies the name of the existing Oracle ACFS volume

4.3.2.4 LIST ACFSVOLUME

This command lists the Oracle ASM Cluster File System (Oracle ACFS) volumes.

Syntax

```

LIST ACFSVOLUMES
WHERE
  ID = id |
  DISKGROUPID = diskgroup_id |
  CLUSTERNUMBER = cluster_number [DISKGROUPNAME = diskgroup_name] |
  CLUSTERNAME = cluster_name [DISKGROUPNAME diskgroup_name] |
  CLUSTERID = cluster_id [DISKGROUPNAME = diskgroup_name]

```

Arguments

The following arguments are available in the WHERE clause:

- ID : Specifies the `es.xml` ID for the Oracle ACFS volume
- DISKGROUPID : Specifies the `es.xml` ID for the container disk group
- CLUSTERNUMBER : Specifies the cluster number in the `es.xml`, starting at 1
- CLUSTERNAME : Specifies the name of the cluster
- CLUSTERID : Specifies the `es.xml` ID of the cluster
- DISKGROUPNAME : Specifies the name of the container disk group

Usage Notes

There is no WHERE clause that will list all the Oracle ACFS volumes in the `es.xml` file.

4.3.3 CELL

You can clone, delete, or list Oracle Exadata storage cells.

- [CLONE CELL](#)
- [CLONE NEWCELL](#)
- [DELETE CELL](#)
- [LIST CELLS](#)

4.3.3.1 CLONE CELL

This command clones an existing storage cell.

 **Note:**

This command is deprecated and replaced by `CLONE NEWCELL`.

Syntax

```
CLONE CELL SRCNAME = source_cell TGTNAME=new_cell_name WHERE  
STEPNAME=step
```

```
SET ADMINNET NAME=admin_name IP=admin_ip [ NETMASK=netmask ]  
[ GATEWAY=gatewayip ]  
SET PRIVNET NAME1=priv_name_1 IP1=priv_ip1 NAME2=priv_name_2  
IP2=priv_ip2  
SET ILOMNET NAME=ilom_name IP=ilom_ip [ NETMASK=netmask ]  
[ GATEWAY=gatewayip ]  
SET RACK NUM=rack_number ULOC=ulocation
```

Arguments

- SRCNAME: Specifies the name of an existing cell from which the new cell is derived.
- TGTNAME: Specifies the name of the new cell being added.
- STEPNAME: Specifies the deployment step to perform, which is one of the following:
 - CONFIG_CELL
 - CREATE_GRIDDISKS
 - ADD_DISKS_TO_ASM
 - REBALANCE

You can set the following options for a cloned cell:

- ADMINNET : Specifies the DNS name and IP address to use for the administration network. The DNS name may optionally include the domain name. You can also optionally specify the netmask and gateway for the administration network.

- **PRIVNET** : Specifies the names and IP addresses to use for the first and second private networks. The names may optionally include the domain name.
- **ILOMNET** : Specifies the DNS name and IP address to use for the ILOM network. The DNS name may optionally include the domain name. You can also optionally specify the netmask and gateway for the ILOM network.
- **RACK** : Specifies the rack number in a multi-rack configuration, where 1 is the first rack, and the physical location (slot number) of the cell in the rack.

The following example shows the command sequence for adding cell `dbm01celadm14` by cloning `dbm01celadm09`.

```
# Clone
CLONE CELL SRCNAME=dbm01celadm09 TGTNAME=dbm01celadm14 where
STEPNAME=CONFIG_CELL
SET ADMINNET NAME=dbm01celadm14 IP=203.0.161.160
SET PRIVNET NAME1=dbm01celadm14-priv1 IP1=192.168.160.19
NAME2=dbm01celadm14-priv2 IP2=192.168.160.20
SET ILOMNET NAME=dbm01celadm14-ilom IP=203.0.161.182
SET RACK NUM=1, ULOC=12

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

# Create Griddisks
CLONE CELL SRCNAME=dbm01celadm09 TGTNAME=dbm01celadm14 where
STEPNAME=CREATE_GRIDDISKS

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

# Add Griddisks to ASM Diskgroups
CLONE CELL SRCNAME=dbm01celadm09 TGTNAME=dbm01celadm14 where
STEPNAME=ADD_DISKS_TO_ASM

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

# Rebalance ASM Diskgroups
CLONE CELL SRCNAME=dbm01celadm09 TGTNAME=dbm01celadm14 where
STEPNAME=REBALANCE

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS
```

4.3.3.2 CLONE NEWCELL

This command clones an existing storage cell.

This command replaces the deprecated `CLONE CELL` command. Instead of using the `STEPNAME` argument, `CLONE NEWCELL` works in conjunction with the `ALTER CLUSTER ADDCELLS` command to add the new storage into an existing cluster.

Syntax

```
CLONE NEWCELL SRCNAME = source_cell TGTNAME=new_cell_name
[ TYPE=cell_type ]

SET ADMINNET NAME=admin_name IP=admin_ip [ NETMASK=netmask ]
[ GATEWAY=gatewayip ]
SET PRIVNET NAME1=priv_name_1 IP1=priv_ip1 NAME2=priv_name_2
IP2=priv_ip2
SET ILOMNET NAME=ilom_name IP=ilom_ip [ NETMASK=netmask ]
[ GATEWAY=gatewayip ]
SET RACK NUM=rack_number ULOC=ulocation
```

Arguments

- `SRCNAME`: Specifies the name of an existing cell from which the new cell is derived.
- `TGTNAME`: Specifies the name of the new cell being added.
- `TYPE`: Optionally specifies the cell type. Use this option only for situations where the clone source is a high capacity (HC) storage server, and the target is an extended (XT) storage server. The following values are allowed:
 - `X8_CELL_XT`: Specifies that the target is an XT storage server with InfiniBand Network Fabric.
 - `X8_ROCE_CELL_XT`: Specifies that the target is an XT storage server with RoCE Network Fabric.

You can set the following options for a cloned cell:

- `ADMINNET` : Specifies the DNS name and IP address to use for the administration network. The DNS name may optionally include the domain name. You can also optionally specify the netmask and gateway for the administration network.
- `PRIVNET` : Specifies the names and IP addresses to use for the first and second private networks. The names may optionally include the domain name.
- `ILOMNET` : Specifies the DNS name and IP address to use for the ILOM network. The DNS name may optionally include the domain name. You can also optionally specify the netmask and gateway for the ILOM network.
- `RACK` : Specifies the rack number in a multi-rack configuration, where 1 is the first rack, and the physical location (slot number) of the cell in the rack.

The following example shows the command sequence for adding cell `dbm02celadm06` by cloning `dbm02celadm01`.

```
# Clone
CLONE NEWCELL SRCNAME=dbm02celadm01 TGTNAME=dbm02celadm06
SET ADMINNET NAME=dbm02celadm06 IP=203.0.96.17
SET PRIVNET NAME1=dbm02cell06vm1str-priv1 IP1=192.168.3.47
NAME2=dbm02cell06vm1str-priv2 IP2=192.168.3.48
SET ILOMNET NAME=dbm02celadm06-ilom IP=203.0.10.103
```

```
SET RACK NUM=1 ULOC=12

SAVE ACTION

# Add Cell to Cluster
ALTER CLUSTER ADDCELLS='dbm02celadm06' POWER=4 WHERE CLUSTERNUMBER=1

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS
```

4.3.3.3 DELETE CELL

This command deletes an existing storage cell.

Syntax

```
DELETE CELL WHERE SRCNAME = cell_host_name
```

Arguments

SRCNAME is used to specify the host name of an existing cell that you want to remove.

4.3.3.4 LIST CELLS

This command lists the storage cells.

Syntax

```
LIST CELLS
```

4.3.4 CLUSTER

You can alter or list the Oracle RAC Clusters on the Oracle Exadata Database Machine.

- [ALTER CLUSTER](#)
- [CREATE CLUSTER](#)
This command is used to run all the steps required to create the virtual guests for a cluster and the following steps to configure and deploy an Oracle RAC database.
- [DESTROY CLUSTER](#)
This command is used to cleanly destroy a cluster.
- [LIST CLUSTERS](#)
This command lists the existing Oracle RAC clusters.
- [UPGRADE CLUSTER](#)
This command is used to upgrade Oracle Clusterware running on the virtual hosts.

4.3.4.1 ALTER CLUSTER

This command alters attributes for an existing Oracle RAC Cluster.

Syntax

```
ALTER CLUSTER {  
    CLUSTERNAME = cluster_name |  
    GIVERSION = grid_version |  
    GIHOMELOC = grid_home_location |  
    INVLOC = inventory_location |  
    BASEDIR = oracle_base_directory |  
    PATCHLIST = patch_list |  
    INSTALLTYPE = CLONE | ZIP }  
WHERE  
    CLUSTERNUMBER = cluster_number |  
    CLUSTERNAME = cluster_name |  
    CLUSTERID = cluster_id
```

Arguments

You can modify the following attributes of the Oracle RAC Cluster:

- **CLUSTERNAME** : Specifies a new, unique name for the cluster
- **GIVERSION** : Specifies the clusterware version, including the bundle patch, for example 12.1.0.2.170418
- **GIHOMELOC** : Specifies a new path for the Grid home
- **INVLOC** : Specifies a new path for the inventory
- **BASEDIR** : Specifies a new path for the Oracle base directory
- **PATCHLIST** : Specifies a full list of patches separated by commas and enclosed in single quotes. By default, this is derived from the bundle patch specified by GIVERSION.
- **INSTALLTYPE**: Specifies the software installation method to use. The options are:
 - **CLONE**: Uses image files to clone a software installation. The image files include the base software release, and may also include additional patches and updates. By including updates in the image files, this method reduces deployment time and improves deployment consistency.

The **CLONE** software installation method is the only supported method for a cluster in a virtual machine (VM) deployment. This option is also available for clusters in a physical (bare metal) deployment.
 - **ZIP**: Performs a fresh software installation by using compressed (zip) files that contain the software release. This method does not include additional patches or updates.

The **ZIP** software installation method is the default method for clusters in a physical (bare metal) deployment. This option is not available for clusters in a virtual machine (VM) deployment.

The following arguments are available in the **WHERE** clause:

- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster

4.3.4.2 CREATE CLUSTER

This command is used to run all the steps required to create the virtual guests for a cluster and the following steps to configure and deploy an Oracle RAC database.

Syntax

The simplest form of this command does not require any arguments. However, if the XML file being used by OEDACLI contains more than one cluster, then the `WHERE` clause is mandatory.

```
CREATE CLUSTER [WHERE CLUSTERNAME=cluster_name]
```

There are additional parameters and `WHERE` clause options for the command that you can use.

```
CREATE CLUSTER [STEPS='short_names'] [SKIPSTEPS='short_names']  
[DELETEDB={true|false}] [WHERE {CLUSTERNAME=cluster_name |  
CLUSTERNUMBER=cluster_number | CLUSTERID=cluster_id }]
```

Arguments

- **STEPS:** This parameter allows you to specify which steps to run. Each step in the list is separated by a space and the list is enclosed in single quotes. When you run the `CREATE CLUSTER` command with this parameter, it performs only the steps listed in the `STEPS` clause, excluding any steps listed in the `SKIPSTEPS` clause.
- **SKIPSTEPS:** When you run the `CREATE CLUSTER` command, it performs all the steps listed in `LIST STEPS` excluding those listed in the `SKIPSTEPS` clause. Each step in the list is separated by a space and the list is enclosed in single quotes.
- **DELETEDB:** If you do not require the database created by OEDA, then the database can be removed upon complete of the `CREATE CLUSTER` command by specifying the parameter `DELETEDB=true`. The default value for this parameter is `false`, which leaves the database in-place.

The following arguments are available in the `WHERE` clause:

- **CLUSTERNUMBER:** Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME:** Specifies the name of the cluster
- **CLUSTERID:** Specifies the `es.xml` ID of the cluster

Example 4-5 Using the STEPS Parameter with CREATE CLUSTER

If you only want to create the guest domains for the Cluster-c1 cluster, then you would use the following command:

```
CREATE CLUSTER STEPS='CREATEVM' WHERE CLUSTERNAME=Cluster-c1
```

Example 4-6 Using the SKIPSTEPS and DELETEDB Parameters with CREATE CLUSTER

If there are 17 steps for the `CREATE CLUSTER` command, and `SKIPSTEPS` is not specified, then OEDACLI will complete all 17 steps. If you use the following command, then only 13 steps are completed, after which the database is deleted.

```
CREATE CLUSTER SKIPSTEPS='VALIDATE CALIBRATE EXACHECK RESECURE'
          DELETEDB=true WHERE CLUSTERNAME=Cluster-cl
```

4.3.4.3 DESTROY CLUSTER

This command is used to cleanly destroy a cluster.

If the XML file contains only 1 cluster there is no need to identify the cluster explicitly. The `MERGE` step will indicate which machines will be destroyed and which cells will be affected.

Syntax

```
DESTROY CLUSTER WHERE {CLUSTERNUMBER = cluster_number |
                       CLUSTERNAME = cluster_name | CLUSTERID = cluster_id }
```

Arguments

The following arguments are available in the `WHERE` clause:

- `CLUSTERNUMBER`: Specifies the cluster number in the `es.xml`, starting at 1
- `CLUSTERNAME`: Specifies the name of the cluster
- `CLUSTERID`: Specifies the `es.xml` ID of the cluster

Example 4-7 Using the DESTROY CLUSTER Command

The following example shows the simplest case, where there is only a single cluster in the XML file.

```
oedacli> DESTROY CLUSTER
oedacli> SAVE ACTION
oedacli> MERGE ACTIONS
oedacli> DEPLOY ACTIONS
```

4.3.4.4 LIST CLUSTERS

This command lists the existing Oracle RAC clusters.

Syntax

```
LIST CLUSTERS
[ WHERE
  CLUSTERNUMBER = cluster_number |
  CLUSTERNAME = cluster_name |
  CLUSTERID = cluster_id ]
```


Arguments

The following arguments are available in the `WHERE` clause:

- `CLUSTERNUMBER`: Specifies the cluster number in the `es.xml`, starting at 1
- `CLUSTERNAME`: Specifies the name of the cluster
- `CLUSTERID`: Specifies the `es.xml` ID of the cluster

Usage Notes

The `WHERE` clause is optional.

Example 4-8 Using the LIST CLUSTERS Command

This example shows how to identify the Oracle RAC clusters on the host.

```
OEDACLI> LIST CLUSTERS
  version : "CloneInstall"
  clusterName : "DBMclu1"
  clusterOwner : "grid"
  clusterVersion : "19.3.0.0.190416"
  clusterHome : "/u01/app/19.0.0.0/grid"
  inventoryLocation : "/u01/app/oraInventory"
  asmScopedSecurity : "false"
```

4.3.4.5 UPGRADE CLUSTER

This command is used to upgrade Oracle Clusterware running on the virtual hosts.

Syntax

```
UPGRADE CLUSTER {GIVERSION = Grid_relnum | GIHOMELOC = Grid_home_dir }
[WHERE {CLUSTERNAME=cluster_name | CLUSTERNUMBER=cluster_number |
        CLUSTERID=cluster_id } [STEPNAME = step_name] ]
```

Arguments

- `GIVERSION`: The version of the Oracle Grid Infrastructure software you are upgrading to
- `GIHOMELOC`: The target location of the Oracle Grid Infrastructure home directory

The following arguments are available in the `WHERE` clause:

- `CLUSTERNUMBER`: Specifies the cluster number in the `es.xml`, starting at 1
- `CLUSTERNAME`: Specifies the name of the cluster
- `CLUSTERID`: Specifies the `es.xml` ID of the cluster
- `STEPNAME`: Specifies the name of the individual upgrade step to be run. If this parameter is not specified, then `ADD_HOME` and `CONFIG_HOME` operations are performed.

Usage Notes

- This option is only available for virtual environments.
- The `RUN_ROOTSCRIPT` step runs the `rootupgrade.sh` script in a rolling mode on all nodes. This step also restarts the database instances on each of the nodes.

Example 4-9 Using the UPGRADE CLUSTER Command

This example shows how to upgrade the Oracle Grid Infrastructure software to release 18.6 for both nodes in the first cluster.

```
OEDACLI> UPGRADE CLUSTER GIVERSION=18.6.0.0.190416 GIHOMELOC=/u01/app/
18.0.0.0/grid
WHERE CLUSTERNUMBER=1
deploy actions
Deploying Action ID : 3 UPGRADE CLUSTER GIVERSION=18.6.0.0.190416
GIHOMELOC=/u01/app/18.0.0.0/grid
  WHERE CLUSTERNUMBER=1 STEPNAME=RUN_ROOTSCRIPT
Deploying UPGRADE CLUSTER
Upgrading Cluster
Running Cluster Verification Utility for upgrade readiness..
Pre-upgrade Validation completed successfully.
Running rootupgrade.sh on node dbm01adm01vm01.example.com
Checking file
root_dbm01adm01vm01.example.com_2019-05-02_05-56-08-814110215.log on
node
  dbm01adm01vm01.example.com
Checking status of cluster on dbm01adm01vm01.example.com
Running rootupgrade.sh on node dbm01adm02vm01.example.com
Checking file
root_dbm01adm02vm01.example.com_2019-05-02_06-14-40-524921970.log on
node
  dbm01adm02vm01.example.com
Checking status of cluster on dbm01adm02vm01.example.com
Running Configuration Assistants on dbm01adm01vm01.example.com
Active cluster version after upgrade [18.0.0.0]
Done...
Done
```

4.3.5 COMPUTE

Use these commands to add or delete a physical compute node as a cluster member.

- **CLONE COMPUTE**
Use this command to clone an existing physical compute cluster node.
- **DELETE COMPUTE**
This command deletes an existing compute node, removes the instance, home and clusterware from the cluster.

4.3.5.1 CLONE COMPUTE

Use this command to clone an existing physical compute cluster node.

Syntax

```
CLONE COMPUTE SRCNAME=source_compute_node TGTNAME=target_compute_node  
[ WHERE STEPNAME = stepname ]
```

```
SET ADMINNET NAME=admin_name IP=admin_ip  
SET PRIVNET NAME1=priv_name_1 IP1=priv_ip1 NAME2=priv_name_2  
IP2=priv_ip2  
[ SET INTERCONNECT NAME1=priv_name_1 IP1=priv_ip1 [ NAME2=priv_name_2  
IP2=priv_ip2 ] ]  
[ SET CLIENTNET NAME=client_name IP=client_ip ]  
[ SET VIPNET NAME=vip_name IP=vip_ip ]  
[ SET BACKUP NAME=backup_name IP=backup_ip ]  
SET ILOMNET NAME=ilom_name IP=ilom_ip  
SET RACK NUM=rack_number ULOC=ulocation
```

Arguments

- SRCNAME: Specifies the name of an existing compute server from which the compute server is derived.
- TGTNAME: Specifies the name of the new compute server being added.
- STEPNAME: Optionally specifies the deployment step to perform, which is one of the following:
 - CREATE_USERS
 - CELL_CONNECTIVITY
 - ADD_NODE
 - EXTEND_DBHOME
 - ADD_INSTANCE

If STEPNAME is not specified, then all steps are performed.

You can set the following options for a cloned compute node:

- ADMINNET : Specifies the DNS name and IP address to use for the administration network. The DNS name may optionally include the domain name.
- PRIVNET : Specifies the names and IP addresses to use for the first and second private networks. The names may optionally include the domain name.
- INTERCONNECT: Only required if the compute nodes are configured with Exadata Secure RDMA Fabric Isolation or InfiniBand security (PKEYS). Specifies the names and IP addresses to use for the first and second private networks. The names may optionally include the domain name.
- CLIENTNET : Specifies the DNS name and IP address to use for the client network. The DNS name may optionally include the domain name. Not required on virtual machine (VM) hosts (Oracle Linux KVM or Oracle VM).
- VIPNET : Specifies the DNS name and IP address to use for the virtual IP (VIP) network. The DNS name may optionally include the domain name. Not required on virtual machine (VM) hosts (Oracle Linux KVM or Oracle VM).

- **BACKUP** : Specifies the DNS name and IP address to use for the backup network. The DNS name may optionally include the domain name. Not required on servers without a backup network.
- **ILOMNET**: Specifies the DNS name and IP address to use for the ILOM network. The DNS name may optionally include the domain name.
- **RACK** : Specifies the rack number in a multi-rack configuration, where 1 is the first rack, and the physical location (slot number) of the server in the rack.

Usage Notes

- For configurations with Exadata Secure RDMA Fabric Isolation or InfiniBand security (PKEYS), **INTERCONNECT** defines the private network that connects the database nodes and **PRIVNET** defines the private storage network.

Example 4-10 Cloning a Compute Node

This example clones the dbadm03 compute server to a new server, dbadm04.

```
CLONE COMPUTE srcname=dbadm03 tgtname=dbadm04
SET ADMINNET name=dbadm04 ip=10.128.161.192
SET PRIVNET name1=dbadm04-priv1 ip1=192.168.161.195 name2=dbadm04-priv2
ip2=192.168.161.196
SET CLIENTNET name=dbadm04-client ip=198.51.100.112
SET VIPNET name=dbadm04-vip ip=10.128.177.116
SET ILOMNET name=dbadm04-ilom ip=10.128.161.203
SET RACK num=1 uloc=17

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS
```

Example 4-11 Cloning a Compute Node Using Steps

This example clones the dbadm03 compute server to a new server, dbadm04, running individual steps. The example shows only three steps and does not show all the steps involved in cloning the compute node.

```
CLONE COMPUTE srcname=dbadm03 tgtname=dbadm04 WHERE
STEPNAME=CREATE_USERS
SET ADMINNET name=dbadm04 ip=10.128.161.192
SET PRIVNET name1=dbadm04-priv1 ip1=192.168.161.195 name2=dbadm04-priv2
ip2=192.168.161.196
SET CLIENTNET name=dbadm04-client ip=198.51.100.112
SET VIPNET name=dbadm04-vip ip=10.128.177.116
SET ILOMNET name=dbadm04-ilom ip=10.128.161.203
SET RACK num=1 uloc=17

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

CLONE COMPUTE compute srcname=dbadm03 tgtname=dbadm04 WHERE
STEPNAME=CELL_CONNECTIVITY
```

```
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS
```

```
CLONE COMPUTE srcname=dbadm03 tgtname=dbadm04 WHERE STEPNAME=ADD_NODE
```

```
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS
```

...

4.3.5.2 DELETE COMPUTE

This command deletes an existing compute node, removes the instance, home and clusterware from the cluster.

Syntax

```
DELETE COMPUTE WHERE SRCNAME = cell_host_name
  [ STEPNAME = stepname ]
```

Arguments

SRCNAME is used to specify the host name of the guest that you want to remove.

STEPNAME is used to clean up steps from a failed CREATE COMPUTE command. Specify the name of the individual creation step to be reverted. The values for *stepname* are:

- ADD_INSTANCE
- EXTEND_DBHOME
- ADD_NODE
- CELL_CONNECTIVITY
- CREATE_USERS

Example 4-12 Deleting a Compute Node - Simple Case

This example removes the dbadm04 compute node from the cluster.

```
LOAD FILE NAME=4-4compute.xml
DELETE COMPUTE where srcname=dbadm04
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS
SAVE FILE name=3-4compute.xml
```

Example 4-13 Deleting a Compute Node with Individual Steps

This example removes the `dbadm04` compute node from the cluster, running each step individually. This example does not show all the steps involved in removing the compute node.

```
LOAD FILE NAME=4-4compute.xml
DELETE COMPUTE where srcname=dbadm04 stepname=ADD_INSTANCE
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

DELETE COMPUTE where srcname=dbadm04 stepname=EXTEND_DBHOME
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

DELETE COMPUTE where srcname=dbadm04 stepname=ADD_NODE
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS
...
```

4.3.6 DATABASE

You can add, alter, delete, or list the Oracle databases.

- [ADD DATABASE](#)
- [ALTER DATABASE](#)
- [DELETE DATABASE](#)
- [LIST DATABASES](#)

4.3.6.1 ADD DATABASE

This command adds an Oracle database, or a CDB or PDB to an existing CDB.

Syntax

```
ADD DATABASE
  [ BLOCKSIZE = blocksize ]
  [ CHARSET = database_character_set ]
  DATADG = data_diskgroup
  [ DBLANG = database_language ]
  DBNAME = database_name
  [ DBTEMPLATE = database_template ]
  [ DBTYPE = database_type ]
  [ HOSTNAMES = 'host_names' ]
  RECODG = reco_diskgroup
WHERE
  { DBHOMEID = database_home_id |
  CLUSTERNUMBER = cluster_number DBHOMELOC = database_home_location |
  CLUSTERNAME = cluster_name DBHOMELOC = database_home_location |
  CLUSTERID = cluster_id DBHOMELOC = database_home_location }
```

```
CLUSTERNUMBER = cluster_number CDBNAME = container_database_name |  
CLUSTERNAME = cluster_name CDBNAME = container_database_name |  
CLUSTERID = cluster_id CDBNAME = container_database_name }
```

Arguments

You can specify the following attributes for the new Oracle database:

- **BLOCKSIZE** : Specifies the block size for the new database. The default value is 8192. This argument is not required for pluggable databases.
- **CHARSET** : Specifies the character set to use for the new database. The default value is AL32UTF8. This argument is not required for pluggable databases.
- **DATADG** : (Mandatory) Specifies the name of the DATA disk group for the new database. This argument is not required for pluggable databases.
- **DBLANG** : Specifies the language to use for the new database. The default value is all_langs. This argument is not required for pluggable databases.
- **DBNAME** : (Mandatory) Specifies the name of the new database.
- **DBTEMPLATE** : Specifies the template to use when creating the new database, ADMIN or DW. The default value is ADMIN. This argument is not required for pluggable databases.
- **DBTYPE** : Specifies the type of database to create. The accepted values are normal DB, CDB, or PDB. The default value is normal DB.
- **HOSTNAMES** : Specifies a comma-separated list of host names on which the database should run. The default value is the list of nodes registered with the database home.
- **RECODG** : (Mandatory) Specifies the name of the RECO disk group for the new database. This argument is not required for pluggable databases.

The following arguments are available in the WHERE clause:

- **DBHOMEID** : Specifies the es.xml ID of the database home.
- **CLUSTERNUMBER** : Specifies the cluster number in the es.xml, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the es.xml ID of the cluster
- **DBHOMELLOCATION** : Specifies the path for the target database home.
- **CDBNAME** : Specifies the database name for the container database. This argument is required only when DBTYPE is set to PDB.

Example 4-14 Adding a New Oracle Database Using OEDACLI

This example shows how to add a new database to the configuration.

1. Load the OEDA XML configuration file.

```
oedacli> LOAD FILE NAME=Oracle-test.xml  
SUCCESS - file loaded OK  
Customer : ora07adm01 - Oracle
```

2. Use the command `LIST DATABASEHOMES` to identify the `databasehome_id` associated with the target database home. This database home will be used to create the new database.

```
oedacli> LIST DATABASEHOMES
version : "CloneInstall"
cluster :
id : "c0_clusterHome"
databaseHomeName : "c0_DbHome_0"
databaseSwOwner : "c0_oracle"
databaseVersion : "12.2.0.1.170718"
databaseHomeLoc : "/u01/app/oracle/product/12.2.0.1/dbhome_1"
inventoryLocation : "/u01/app/oraInventory"
language : "all_langs"
machines :
machine : ...
...
patches :
patch :
patchNumber : "26133434"
basedir : "/u01/app/oracle"
useZfs : "false"
id : "c0_databaseHome1"
```

3. Create an action for creating the new database. You must provide a name for the new database, and the names of the DATA and RECO disk groups that the new database should use. Also, use the ID retrieved in the previous step to specify the Oracle Home in which to create this database.

```
oedacli> ADD DATABASE DBNAME='testdb' DATADG='DATA1'
RECODG='RECO1'
WHERE DBHOMEID='c0_databaseHome'
```

4. Save the action.

```
oedacli> SAVE ACTION
```

5. Merge all actions.

```
oedacli> MERGE ACTIONS
processMerge
processMergeActions
Merging Action : add database dbname='testdb' DATADG='DATA1'
RECODG='RECO1'
where DBHOMEID='c0_databaseHome'
Merging ADD DATABASE
Action Validated and Merged OK
```

6. Save the action to a new Engineered Systems XML configuration file.

```
oedacli> SAVE FILE NAME ='cli-test-2databases.xml'
File : cli-test-2databases.xml saved OK
```


7. Deploy the actions.

```
oedacli> DEPLOY ACTIONS
Deploying Action ID : 4 add database dbname='testdb'
DATADG='DATA1' RECODG='
RECO1' where DBHOMEID='c0_databaseHome'
Deploying ADD DATABASE
Running datapatch on database 'testdb'
Done...
Done
```

4.3.6.2 ALTER DATABASE

This command alters an existing Oracle database.

Syntax

```
ALTER DATABASE
  [ BLOCKSIZE = blocksize ]
  [ CHARSET = database_character_set ]
  DATADG = data_diskgroup
  [ DBLANG = database_language ]
  DBNAME = database_name
  [ DBTEMPLATE = database_template ]
  [ DBTYPE = database_type ]
  [ HOSTNAMES = 'host_names' ]
  RECODG = reco_diskgroup
WHERE
  ID = database_id |
  CLUSTERNUMBER = cluster_number DATABASENAME = database_name |
  CLUSTERNAME = cluster_name DATABASENAME = database_name |
  CLUSTERID = cluster_id DATABASENAME = database_name
```

Arguments

You can modify the following attributes for the Oracle database:

- **BLOCKSIZE** : Specifies the new block size for the database. The default value is 8192. This argument is not required for pluggable databases.
- **CHARSET** : Specifies the new character set to use for the database. The default value is AL32UTF8. This argument is not required for pluggable databases.
- **DATADG** : (Mandatory) Specifies the name of the DATA disk group for the database. This argument is not required for pluggable databases.
- **DBLANG** : Specifies the new language to use for the database. The default value is all_langs. This argument is not required for pluggable databases.
- **DBNAME** : (Mandatory) Specifies the name of the database.
- **DBTEMPLATE** : Specifies the template to use when creating the database, ADMIN or DW. The default value is ADMIN. This argument is not required for pluggable databases.
- **DBTYPE** : Specifies the type of database to create. The accepted values are normal DB, CDB, or PDB. The default value is normal DB.

- **HOSTNAMES** : Specifies a comma-separated list of host names on which the database should be created. The default value is the list of nodes registered with the database home.
- **RECODG** : (Mandatory) Specifies the name of the RECO disk group for the database. This argument is not required for pluggable databases.

The following arguments are available in the `WHERE` clause:

- **ID** : Specifies the `es.xml` ID of the database
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster
- **DATABASENAME** : Specifies the name of the database in the cluster

4.3.6.3 DELETE DATABASE

This command deletes an existing Oracle database.

Syntax

```
DELETE DATABASE
WHERE
  ID = database_id |
  CLUSTERNUMBER = cluster_number DATABASENAME = database_name |
  CLUSTERNAME = cluster_name DATABASENAME = database_name |
  CLUSTERID = cluster_id DATABASENAME = database_name
```

Arguments

The following arguments are available in the `WHERE` clause:

- **ID** : Specifies the `es.xml` ID of the database
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster
- **DATABASENAME** : Specifies the name of the database in the cluster

Example 4-15 Deleting an Oracle Database Using OEDACLI

This example shows how to delete an existing database from the configuration.

1. Load the OEDA XML configuration file.

```
oedacli> LOAD FILE NAME=Oracle-test.xml
SUCCESS - file loaded OK
Customer : ora07adm01 - Oracle
```

2. List the databases contained in the XML configuration file, so that you can identify the database ID associated with the database you want to delete.

```
oedacli> LIST DATABASES
version : "2"
```

```
databaseOwner : "c0_oracle"
databaseSid : "dbm01"
databaseBlockSize : "8192"
characterSet : "AL32UTF8"
databaseHome :
id : "c0_databaseHome"
machines :
machine :
.....
.....
diskGroups :
diskGroup :
id : "c0_datadg"
id : "c0_recodg"
id : "c0_dbm01"
.....
.....
databaseHome :
id : "c0_databaseHome3"
machines :
machine :
....
....
databaseTemplate : "OLTP"
databaseStyle : "ADMIN"
language : "all_langs"
diskGroups :
diskGroup :
id : "c0_datadg"
id : "c0_recodg"
id : "c0_testdb01"
```

3. Create an action to delete the target database. Use the database id identified in the previous step

```
oedacli> delete database where id='c0_testdb01'
```

4. Save the action.

```
oedacli> SAVE ACTION
```

5. Merge all actions.

```
oedacli> MERGE ACTIONS
processMerge
processMergeActions
Merging Action : delete database where id='c0_testdb01'
Merging DELETE DATABASE
Action Validated and Merged OK
```

6. Deploy the actions into the current configuration.

```
oedacli> DEPLOY ACTIONS
Deploying Action ID : 1 delete database where id='c0_testdb01'
```

```

Deploying DELETE DATABASE
Deleting database testdb
Done...
Done

```

4.3.6.4 LIST DATABASES

This command lists the existing Oracle databases.

Syntax

```

LIST DATABASES
[ WHERE
  ID = database_id |
  CLUSTERNUMBER = cluster_number |
  CLUSTERNAME = cluster_name |
  CLUSTERID = cluster_id ]

```

Arguments

The following arguments are available in the WHERE clause:

- **ID** : Specifies the `es.xml` ID of the database
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster

4.3.7 DATABASE HOME

You can add, alter, delete, or list the Oracle database homes.

- [ADD DATABASEHOME](#)
- [ALTER DATABASEHOME](#)
- [DELETE DATABASEHOME](#)
- [LIST DATABASEHOMES](#)

4.3.7.1 ADD DATABASEHOME

This command adds an Oracle database home to an existing cluster.

Syntax

```

ADD DATABASEHOME
  OWNER = owner DBVERSION = version
  DBHOMELOC = database_home_path [DBHOMENAME = database_home_name]
  [INVLOC = inventory_location] [DBLANG = database_language]
  [MACHINELIST = 'machine_list'] [PATCHLIST = 'patch_list']
  [BASEDIR = base_directory] [VOLUMESIZE = volume_size]
  [INSTALLTYPE = CLONE | ZIP]
WHERE

```

```
{ CLUSTERNUMBER = cluster_number | CLUSTERNAME = cluster_name |  
  CLUSTERID = cluster_id }
```

Arguments

You can specify the following attributes for the new Oracle database home:

- **OWNER** : (Mandatory) Specifies the operating system user that owns the Oracle home.
- **DBVERSION** : (Mandatory) Specifies the version of the software in the database home, including the bundle patch, for example 12.1.0.2.170418.
- **DBHOMELOC** : (Mandatory) Specifies the file system path for the new Oracle database home directory.
- **DBHOMENAME** : Specifies the name of the new database home. The default value is the standard template name, for example, `dbhome_1`.
- **INVLOC** : Specifies the location of the Oracle inventory directory. The default value is derived from the cluster home.
- **DBLANG** : Specifies the language to use with the Oracle home. The default value is `all_langs`.
- **MACHINELIST** : Specifies a list of nodes to create the Oracle home on. The default value is list of nodes registered with the Grid home.
- **PATCHLIST** : Specifies a comma-separated list of patches for the database home. The default value is derived from `DBVERSION`.
- **BASEDIR** : Specifies the location of the Oracle Base directory. The default value is derived from `OWNER`.
- **VOLUMESIZE** : Specifies the size of the home directory volume. The default value is 50 GB. This argument is applicable only for database homes for Oracle VM clusters.
- **INSTALLTYPE**: Specifies the software installation method to use. The options are:
 - **CLONE**: Uses image files to clone a software installation. The image files include the base software release, and may also include additional patches and updates. By including updates in the image files, this method reduces deployment time and improves deployment consistency.

The `CLONE` software installation method is the only supported method for database homes in a virtual machine (VM) deployment. This option is also available for database homes in a physical (bare metal) deployment.
 - **ZIP**: Performs a fresh software installation by using compressed (zip) files that contain the software release. This method does not include additional patches or updates.

This option is not available for database homes in a virtual machine (VM) deployment.

If you do not specify the `INSTALLTYPE` setting, then the default setting matches the `INSTALLTYPE` setting in the containing cluster.

The following arguments are available in the `WHERE` clause:

- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1

- CLUSTERNAME : Specifies the name of the cluster
- CLUSTERID : Specifies the es.xml ID of the cluster

Example 4-16 Adding a Database Home to a Deployment using OEDACLI

This example shows how to add a new Oracle home. `dbversion` should specify both the base version and the target Bundle Patch version. In this example, we are setting up a new Oracle database home with 12.1.0.2 DBBP 170418. The software will be installed under `/u01/app/oracle/product/12.1.0.2/dbhome_2`.

```
./oedacli

oedacli> LOAD FILE NAME=Oracle-test12.xml
SUCCESS - file loaded OK
Customer : ora07adm01 - Oracle

oedacli> ADD DATABASEHOME OWNER='oracle' DBVERSION='12.1.0.2.170418'
DBHOMELOC='/u0
1/app/oracle/product/12.1.0.2/dbhome_2' WHERE CLUSTERNUMBER=1

oedacli> SAVE ACTION

oedacli> MERGE ACTIONS
processMerge
processMergeActions
Merging Action : add databasehome owner='oracle'
dbversion='12.1.0.2.170418'
dbhomeloc='/u01/app/oracle/product/12.1.0.2/dbhome_2' where
clusternumber=1
Merging ADD DATABASEHOME
Action Validated and Merged OK

oedacli> DEPLOY ACTIONS
Deploying Action ID : 1 add databasehome owner='oracle'
dbversion='12.1.0.2.170418'
dbhomeloc='/u01/app/oracle/product/12.1.0.2/dbhome_2' where
clusternumber=1
Deploying ADD DATABASEHOME
Checking local files..
Required files are:
SUCCESS: Required file /scratch/cpabba/Oeda/main/linux-x64/WorkDir/
p26022196_121020
_Linux-x86-64.zip exists...
SUCCESS: Required file /scratch/cpabba/Oeda/main/linux-x64/WorkDir/
p26027162_121020
_Linux-x86-64.zip exists...
SUCCESS: Required file /scratch/cpabba/Oeda/main/linux-x64/WorkDir/
p6880880_122010_
Linux-x86-64.zip exists...
SUCCESS: Required file /scratch/cpabba/Oeda/main/linux-x64/WorkDir/
linuxamd64_12102
_database_2of2.zip exists...
SUCCESS: Required file /scratch/cpabba/Oeda/main/linux-x64/WorkDir/
linuxamd64_12102
_database_lof2.zip exists...
```

```
Copying required files...
Checking status of remote files...
Checking status of existing files on remote nodes...
Getting status of local files...
Copying file: linuxamd64_12102_database_1of2.zip to node
nodeladm01.example.com
Copying file: linuxamd64_12102_database_2of2.zip to node
nodeladm01.example.com
Copying file: p26022196_121020_Linux-x86-64.zip to node
nodeladm01.example.com
Copying file: p26027162_121020_Linux-x86-64.zip to node
nodeladm01.example.com
Copying file: p26022196_121020_Linux-x86-64.zip to node
nodeladm02.example.com
Copying file: p26027162_121020_Linux-x86-64.zip to node
nodeladm02.example.com
Completed copying files...
Extracting files...
Completed extracting files...
Validating nodes for database readiness...
Installing database software ...
Running database installer on node nodeladm01.example.com ... Please
wait...
After running database installer...
Patching...
Done...
Done
```

4.3.7.2 ALTER DATABASEHOME

This command alters an existing Oracle database home.

Syntax

```
ALTER DATABASEHOME {
  OWNER = owner |
  DBVERSION = version |
  DBHOMELOC = database_home_path |
  DBHOMENAME = database_home_name |
  INVLOC = inventory_location |
  DBLANG = language |
  MACHINELIST = 'machine_list' |
  PATCHLIST = 'patch_list' |
  BASEDIR = base_directory |
  INSTALLTYPE = CLONE | ZIP }
WHERE
  ID = database_home_id |
  CLUSTERNUMBER = cluster_number DBHOMELOC = database_home_path |
  CLUSTERNAME = cluster_name DBHOMELOC = database_home_path |
  CLUSTERID = cluster_id DBHOMELOC = database_home_path
```

Arguments

You can modify the following attributes for the Oracle database home:

- **OWNER** : Specifies the operating system user that owns the Oracle home.
- **DBVERSION** : Specifies the version of the software in the database home, including the bundle patch, for example 12.1.0.2.170418.
- **DBHOMELOC** : Specifies the file system path for the new Oracle database home directory.
- **DBHOMENAME** : Specifies the new name of the database home. The database home name must be unique within a cluster.
- **INVLOC** : Specifies the location of the Oracle inventory directory.
- **DBLANG** : Specifies the supported languages for the Oracle home.
- **MACHINELIST** : Specifies a list of nodes. The nodes must exist within the cluster.
- **PATCHLIST** : Specifies a comma-separated, custom list of patches for the database home.
- **BASEDIR** : Specifies the location of the Oracle Base directory.
- **INSTALLTYPE**: Specifies the software installation method to use. The options are:
 - **CLONE**: Uses image files to clone a software installation. The image files include the base software release, and may also include additional patches and updates. By including updates in the image files, this method reduces deployment time and improves deployment consistency.

The **CLONE** software installation method is the only supported method for database homes in a virtual machine (VM) deployment. This option is also available for database homes in a physical (bare metal) deployment.
 - **ZIP**: Performs a fresh software installation by using compressed (zip) files that contain the software release. This method does not include additional patches or updates.

This option is not available for database homes in a virtual machine (VM) deployment.

The following arguments are available in the **WHERE** clause:

- **ID** : Specifies the `es.xml` ID for the database home
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster
- **DBHOMELOC** : Specifies the path for the database home directory.

4.3.7.3 DELETE DATABASEHOME

This command removes an existing Oracle database home.

Syntax

```
DELETE DATABASEHOME
WHERE
  ID = database_home_id |
  CLUSTERNUMBER = cluster_number DBHOMELOC = database_home_path |
```



```
CLUSTERNAME = cluster_name DBHOMELOC = database_home_path |  
CLUSTERID = cluster_id DBHOMELOC = database_home_path
```

Arguments

The following arguments are available in the `WHERE` clause:

- `ID` : Specifies the `es.xml` ID for the database home
- `CLUSTERNUMBER` : Specifies the cluster number in the `es.xml`, starting at 1
- `CLUSTERNAME` : Specifies the name of the cluster
- `CLUSTERID` : Specifies the `es.xml` ID of the cluster
- `DBHOMELOC` : Specifies the path for the database home directory.

Usage Notes

The database home you are deleting must not contain any dependent databases.

Example 4-17 Removing an Existing Oracle Database Home

This example shows how to delete an existing Oracle Database home directory.

```
oedacli> LOAD FILE NAME=Oracle-test12.xml  
SUCCESS - file loaded OK  
Customer : ora07adm01 - Oracle  
  
oedacli> DELETE DATABASEHOME WHERE id='c0_databaseHome1'  
  
oedacli> SAVE ACTION  
  
oedacli> MERGE ACTIONS  
processMerge  
processMergeActions  
Merging Action : delete databasehome where id='c0_databaseHome1'  
Merging DELETE DATABASEHOME  
Action Validated and Merged OK  
  
oedacli> DEPLOY ACTIONS  
Deploying Action ID : 3 delete databasehome where id='c0_databaseHome1'  
Deploying DELETE DATABASEHOME  
Validating Oracle home..  
Deinstalling database home c0_DbHome_1  
Unmounting file system..  
Updating /etc/fstab entries...  
Done...  
Done
```

4.3.7.4 LIST DATABASEHOMES

This command lists the existing Oracle database homes.

Syntax

```
LIST DATABASEHOMES
[ WHERE
  ID = database_home_id |
  CLUSTERNUMBER = cluster_number NAME = database_home_name |
  CLUSTERNAME = cluster_name NAME = database_home_name |
  CLUSTERID = cluster_id NAME = database_home_name ]
```

Arguments

The following arguments are available in the `WHERE` clause:

- `ID` : Specifies the `es.xml` ID for the target database home
- `CLUSTERNUMBER` : Specifies the cluster number in the `es.xml`, starting at 1
- `CLUSTERNAME` : Specifies the name of the cluster
- `CLUSTERID` : Specifies the `es.xml` ID of the cluster
- `NAME` : Specifies the name of the target database home

4.3.8 DISKGROUP

You can add, alter, delete, or list the Oracle ASM disk groups.

- [ADD DISKGROUP](#)
- [ALTER DISKGROUP](#)
- [ALTER DISKGROUPS](#)
- [DELETE DISKGROUP](#)
- [LIST DISKGROUPS](#)

4.3.8.1 ADD DISKGROUP

This command adds an Oracle ASM disk group.

Syntax

```
ADD DISKGROUP
[ ACFSNAME=acfs_name ]
[ ACFSPTH=acfs_path ]
[ ACFSIZE=acfs_size ]
[ CELLLIST='cell_list' ]
DISKGROUPNAME=diskgroup_name
DISKGROUPSIZE=diskgroup_size
[ OCRVOTE=ocr_voting_disks_included ]
[ QUORUMDISK=quorum_disks_included ]
REDUNDANCY=redundancy_level
SLICESIZE=slice_size
[ SPARSE=sparse ]
[ SPARSEVIRTUALSIZE=sparse_virtual_size ]
[ TYPE = diskgroup_type [ DATABASENAME = database_name ]]
```

```
WHERE
  CLUSTERNUMBER = cluster_number |
  CLUSTERNAME = cluster_name |
  CLUSTERID = cluster_id
```

Arguments

You can specify the following attributes for the new Oracle ASM disk group:

- **ACFSNAME** : Specifies the name for the Oracle ACFS volume on this diskgroup
- **ACFSPATH** : Specifies the path for the Oracle ACFS file system
- **ACFSSIZE** : Specifies the size for the Oracle ACFS volume on this diskgroup
- **CELLLIST** : A comma separated list of cells for this disk group, enclosed in single quotes
- **DISKGROUPNAME** : (Mandatory) Specifies the name of the disk group
- **DISKGROUPSIZE** : (Mandatory) Specifies the size of the disk group
- **OCRVOTE** : A boolean variable, where **TRUE** indicates the disk group should store the OCR and voting disks for the cluster.
- **QUORUMDISK** : A boolean variable, where **TRUE** indicates the disk group requires a quorum
- **REDUNDANCY** : (Mandatory) Specifies the redundancy level for the disk group. The accepted values are **NORMAL** or **HIGH**.
- **SLICESIZE** : (Mandatory) Specifies the slice size on disk for each grid disk for this disk group.
- **SPARSE** : A boolean variable, where **TRUE** indicates this should be sparse disk group.
- **SPARSEVIRTUALSIZE** : Specifies the sparse virtual size.
- **TYPE** : Specifies the diskgroup type, which can be **DATA**, **RECO**, **DBFS** or **OTHER**. The default value is **OTHER**. If you want the database to use this disk group, then the type must be set to **DATA** or **RECO**.
- **DATABASENAME** : Specifies the name of the target database for disk groups of type **DATA** or **RECO**.

The following arguments are available in the **WHERE** clause:

- **CLUSTERNUMBER** : Specifies the cluster number in the *es.xml*, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the *es.xml* ID of the cluster

4.3.8.2 ALTER DISKGROUP

This command alters attributes for an existing Oracle ASM disk group.

Syntax

```
ALTER DISKGROUP {
  ACFSNAME=acfs_name |
```

```

ACFSPATH=acfs_path |
ACFSSIZE=acfs_size |
CELLLIST='cell_list' |
DISKGROUPNAME=diskgroup_name |
DISKGROUPSIZE=diskgroup_size |
OCRVOTE=ocr_voting_disks_included |
QUORUMDISK=quorum_disks_included |
REDUNDANCY=redundancy_level |
SLICESIZE=slice_size |
SPARSE=sparse |
SPARSEVIRTUALSIZE=sparse_virtual_size }
WHERE {
  ID = diskgroup_id|
  CLUSTERNUMBER = cluster_number DISKGROUPNAME = diskgroup_name|
  CLUSTERNAME = cluster_name DISKGROUPNAME = diskgroup_name|
  CLUSTERID = cluster_id DISKGROUPNAME = diskgroup_name }

```

Arguments

You can modify the following attributes of an Oracle ASM disk group:

- **ACFSNAME** : Specifies a new name for the Oracle ACFS volume on this diskgroup
- **ACFSPATH** : Specifies the new path for the Oracle ACFS file system
- **ACFSSIZE** : Specifies a new size for the Oracle ACFS volume on this diskgroup
- **CELLLIST** : A comma separated list of cells for this disk group, enclosed in single quotes
- **DISKGROUPNAME** : Specifies a new name for the disk group
- **DISKGROUPSIZE** : Specifies a new size for the disk group
- **OCRVOTE** : A boolean variable, where **TRUE** indicates the disk group should store the OCR and voting disks for the cluster.
- **QUORUMDISK** : A boolean variable, where **TRUE** indicates the disk group requires a quorum
- **REDUNDANCY** : Specifies the redundancy level for the disk group. The accepted values are **NORMAL** or **HIGH**.
- **SLICESIZE** : Specifies the slice size on disk for each grid disk for this disk group.
- **SPARSE** : A boolean variable, where **TRUE** indicates this should be sparse disk group.
- **SPARSEVIRTUALSIZE** : Specifies the sparse virtual size.

The following arguments are available in the **WHERE** clause:

- **ID** : Specifies the *es.xml* ID for the disk group
- **CLUSTERNUMBER** : Specifies the cluster number in the *es.xml*, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the *es.xml* ID of the cluster
- **DISKGROUPNAME** : Specifies the name of the disk group

Usage Notes

If you modify the redundancy, the volume size will be maintained, however the slice size will be recalculated.

Example 4-18 Resizing an Oracle ASM Diskgroup Using OEDACLI

This example shows how to resize an existing Oracle ASM disk group using OEDACLI.

1. Load the OEDA XML configuration file.

```
oedacli> LOAD FILE NAME=Oracle-test.xml
SUCCESS - file loaded OK
Customer : ora07adm01 - Oracle
```

2. Use the command `LIST DISKGROUPS` to view the disk group details.

```
oedacli> LIST DISKGROUPS
version : "PhysicalLayoutLegacy"
diskGroupName : "DBFS_DG"
redundancy : "NORMAL"
sliceSize : "default"
machines :
machine :
...
...
cellDisks : "ALL"
diskGroupSize : "436G"
ocrVote : "true"
quorumDisk : "false"
id : "c0_dbfsdg"
version : "PhysicalLayoutLegacy"
diskGroupName : "DATA1"
redundancy : "NORMAL"
sliceSize : "2953G"
machines :
..
..
cellDisks : "ALL"
diskGroupSize : "53154G"
ocrVote : "false"
quorumDisk : "false"
id : "c0_datadg"
version : "PhysicalLayoutLegacy"
diskGroupName : "RECO1"
redundancy : "NORMAL"
sliceSize : "ALL"
machines :
..
..
cellDisks : "ALL"
diskGroupSize : "13284G"
ocrVote : "false"
```

```
quorumDisk : "false"  
id : "c0_recodg"
```

3. Create an action for resizing the disk group.

```
oedacli> ALTER DISKGROUP DISKGROUPSIZE='2500G' WHERE  
CLUSTERNUMBER=1 DISKGROUPN  
AME='DATA1'
```

4. Save the action.

```
oedacli> SAVE ACTION
```

5. Merge all actions.

```
oedacli> MERGE ACTIONS  
processMerge  
processMergeActions  
Merging Action : alter diskgroup diskgroupsize='2500G' where  
diskgroupname='DAT  
A1' clusternumber=1  
Merging ALTER DISKGROUP  
Action Validated and Merged OK
```

6. Deploy the actions.

```
oedacli> DEPLOY ACTIONS  
Deploying Action ID : 3 alter diskgroup diskgroupsize='2500G' where  
diskgroupna  
me='DATA1'  
Deploying ALTER DISKGROUP  
Deploying Action ID : 4 alter diskgroup diskgroupsize='2500G' where  
diskgroupna  
me='DATA1' clusternumber=1  
Deploying ALTER DISKGROUP  
Validating ASM Disk status..  
Disks in ASM Disk Group DATA1 have valid status  
Shrinking Disk Group size to 2500 GB  
Resizing ASM Disk Group..  
Checking Reblance operations in disk group DATA1  
Rebalance operations completed in disk group DATA1  
Original Grid Disk Size 2953 GB, new size 138 GB on  
node1celadm02.example.com  
Original Grid Disk Size 2953 GB, new size 138 GB on  
node1celadm03.example.com  
Original Grid Disk Size 2953 GB, new size 138 GB on  
node1celadm01.example.com  
Updated/New Disk Group size 2484 GB, original size 53154 GB.  
Done...  
Done
```

4.3.8.3 ALTER DISKGROUPS

This command alters disk group sizing attributes for existing Oracle ASM disk groups.

Syntax

```
ALTER DISKGROUPS
  DATASPLIT=data_split_percent
  RECOSPLIT=reco_split_percent
  DATAREDUNDANCY=data_redundancy_level
  COREDUNDANCY=reco_redundancy_level
  SPLITDISK=percent_of_cell_disk
WHERE
  DATADG = data_diskgroup
  RECODG = reco_diskgroup
  CLUSTERNAME = cluster_name | CLUSTERNUMBER = cluster_number |
  CLUSTERID = cluster_id
```

Arguments

You can modify the following sizing attributes of the DATA and RECO Oracle ASM disk groups:

- **DATASPLIT** : Specifies a percentage of the Oracle ASM volume allocated to the DATA disk group
- **RECOSPLIT** : Specifies a percentage of the Oracle ASM volume allocated to the RECO disk group
- **DATAREDUNDANCY** : Specifies the redundancy level for the DATA disk group. The accepted values are NORMAL or HIGH.
- **COREDUNDANCY** : Specifies the redundancy level for the RECO disk group. The accepted values are NORMAL or HIGH.
- **SPLITDISK** : Specifies the percentage of the cell disk to allocate to the DATA and RECO disk groups. The specified value must be an integer between 1 and 100.

The following arguments are available in the WHERE clause:

- **DATADG** : Specifies the name of the DATA disk group in the cluster
- **RECODG** : Specifies the name of the RECO disk group in the cluster
- **CLUSTERNUMBER** : Specifies the cluster number in the *es.xml*, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the *es.xml* ID of the cluster

4.3.8.4 DELETE DISKGROUP

This command deletes an existing Oracle ASM disk group.

Syntax

```
DELETE DISKGROUP
WHERE
  ID = diskgroup_id |
  CLUSTERNUMBER = cluster_number DISKGROUPNAME = diskgroup_name |
  CLUSTERNAME = cluster_name DISKGROUPNAME = diskgroup_name |
  CLUSTERID = cluster_id DISKGROUPNAME = diskgroup_name
```

Arguments

The following arguments are available in the `WHERE` clause:

- `ID` : Specifies the `es.xml` ID for the disk group
- `CLUSTERNUMBER` : Specifies the cluster number in the `es.xml`, starting at 1
- `CLUSTERNAME` : Specifies the name of the cluster
- `CLUSTERID` : Specifies the `es.xml` ID of the cluster
- `DISKGROUPNAME` : Specifies the name of the disk group

4.3.8.5 LIST DISKGROUPS

This command lists the existing Oracle ASM disk groups.

Syntax

```
LIST DISKGROUPS
[ WHERE
  ID = diskgroup_id |
  CLUSTERNUMBER = cluster_number |
  CLUSTERNAME = cluster_name |
  CLUSTERID = cluster_id ]
```

Arguments

The following arguments are available in the `WHERE` clause:

- `ID` : Specifies the `es.xml` ID for the disk group
- `CLUSTERNUMBER` : Specifies the cluster number in the `es.xml`, starting at 1
- `CLUSTERNAME` : Specifies the name of the cluster
- `CLUSTERID` : Specifies the `es.xml` ID of the cluster
- `DISKGROUPNAME` : Specifies the name of the disk group

4.3.9 ES

You can discover an engineered system or alter a limited number of attributes of an engineered system.

- [ALTER ES](#)
This command alters the attributes of an engineered system.
- [DISCOVER ES](#)
This command discovers an existing configuration of an engineered system.

4.3.9.1 ALTER ES

This command alters the attributes of an engineered system.

Syntax

```
ALTER ES
  PAAS = 'pass_mode' |
  CUSTOMERNAME = customer_name |
  CUSTOMERDEPT = customer_dept
```

Arguments

You can modify the following attributes for the engineered system:

- **PAAS** : A boolean value, `TRUE` or `FALSE`, which specifies whether to enable PaaS mode.
- **CUSTOMERNAME** : Specifies the customer name.
- **CUSTOMERDEPT** : Specifies the customer department.

4.3.9.2 DISCOVER ES

This command discovers an existing configuration of an engineered system.

Syntax

```
DISCOVER ES
  HOSTNAMES = host_names
  LOCATION = directory_name
```

Arguments

You can use the following arguments when discovering an engineered system configuration:

- **HOSTNAMES**: Specifies the list of nodes to be discovered, with each host name separated by a ',' or a space. The list must be enclosed in quotation marks, for example, 'node1,node2' or 'node1 node2'. For virtual environments with user domains, you need to provide a list of only management domain (dom0) host names and storage servers.
- **LOCATION**: Specifies the target directory for file creation.

Usage Notes

The following are requirements and limitations of the `DISCOVER ES` command:

- Supported only on Exadata. ZDLRA is not supported.
- All hosts must have the same root password.
- Discover will not work with custom Operating System profiles and customized SQL*Plus prompts.
- Only RAC databases are discovered. Single instance databases and SIHA databases are not discovered.
- Oracle Clusterware must be running.
- Only online databases are discovered.

- If `DISCOVER ES` cannot log into databases as `SYSDBA` using Operating System authentication databases are skipped.
- Only database resources registered with Oracle Clusterware are discovered.
- Virtual machines that are not running are not discovered.

4.3.10 GUEST

You can clone or delete an Oracle VM guest.

- **CLONE GUEST**
This command clones an existing VM guest and creates a new fully functional guest with storage, users, and Oracle RAC instances.
- **DELETE GUEST**
This command deletes an existing Oracle VM guest, and removes all of its configuration: instances, software installations, storage, and users from the cluster.

4.3.10.1 CLONE GUEST

This command clones an existing VM guest and creates a new fully functional guest with storage, users, and Oracle RAC instances.

Prerequisites

Before you can clone an virtual machine (VM) guest, you must download the appropriate software images from My Oracle Support and extract the files. Place the extracted "klone.zip" files for Oracle Grid Infrastructure and Oracle Database in the `/EXAVMIMAGES/onecommand/<platform>/WorkDir` directory in the management domain (dom0) where the new guest clone will be created.

Syntax

```
CLONE GUEST SRCNAME = source_guest [ TGTNAME = target_guest WHERE
STEPNAME = stepname ]

SET PARENT NAME=parent_name
[ SET ADMINNET NAME=admin_name IP=admin_ip ]
[ SET BACKUP NAME=backup_name IP=backup_ip ]
SET PRIVNET NAME1=priv_name_1 IP1=priv_ip1 NAME2=priv_name_2
IP2=priv_ip2
[ SET INTERCONNECT NAME1=priv_name_1 IP1=priv_ip1 NAME2=priv_name_2
IP2=priv_ip2 ]
SET CLIENTNET NAME=client_name IP=client_ip [ NATIP=nat_ip
NATHOSTNAME=nat_name NATDOMAINNAME=nat_domain NATNETMASK=nat_mask ]
SET VIPNET NAME=vip_name IP=vip_ip
```

Arguments

- **SRCNAME**: Specifies the name of an existing guest VM from which the new guest VM is derived.
- **TGTNAME**: Specifies the host name of the new guest VM that you are creating with the clone operation. This argument is used with the `WHERE STEPNAME` clause.

- **STEPNAME:** Specifies the deployment step to perform, which is one of the following:
 - CREATE_GUEST
 - CREATE_USERS
 - CELL_CONNECTIVITY
 - ADD_NODE
 - EXTEND_DBHOME
 - ADD_INSTANCE

You can specify the following additional settings when specifying `CLONE GUEST` with the `WHERE STEPNAME = CREATE_GUEST` clause or when specifying `CLONE GUEST` without using any `WHERE STEPNAME` clause:

- **PARENT:** Specifies the name of the virtual machine (VM) host (Oracle Linux KVM or Oracle VM) that hosts the new guest.
- **ADMINNET :** Specifies the DNS name and IP address to use for the administration network.
- **BACKUP :** Specifies the DNS name and IP address to use for the backup network. The DNS name may optionally include the domain name. Not required on guests without a backup network.
- **PRIVNET :** Specifies the names and IP addresses to use for the first and second private networks.
- **INTERCONNECT:** Only required if the compute nodes are configured with Exadata Secure RDMA Fabric Isolation or InfiniBand security (PKEYS). Specifies the names and IP addresses to use for the first and second private networks.
- **CLIENTNET :** Specifies the DNS name and IP address to use for the client network. May optionally include host name, IP address, domain name, and netmask details to support Network Address Translation (NAT).
- **VIPNET :** Specifies the DNS name and IP address to use for the virtual IP (VIP) network.

Usage Notes

- For configurations with Exadata Secure RDMA Fabric Isolation or InfiniBand security (PKEYS), `INTERCONNECT` defines the private network that connects the database nodes and `PRIVNET` defines the private storage network.
- When you load the XML configuration file to use when cloning a guest domain, the XML file should have the management domain (dom0) node of the source VM defined but not allocated to the cluster that is being extended into the guest domain.
- You should save the modified XML file after the actions are merged so that you have a new XML file that reflects the addition of the new guest domain.

Example 4-19 Cloning a Guest VM

This example shows how to clone a new guest from an existing node in the cluster. In this example, the source node name is `exa01adm01vm01.example.com` and the new node is `exa01adm03vm01.example.com`. A series of `SET` commands are used to specify the configuration for the new guest VM.

The XML file loaded at the beginning of this example has the dom0 node (host01adm01.example.com) defined but not allocated to the cluster that is being extended into the guest.

```
oedacli> LOAD FILE NAME=exa01-pre-cloning-node3.xml
oedacli> CLONE GUEST SRCNAME='exa01adm01vm01.example.com'
oedacli> SET PARENT NAME='exa01adm03.example.com'
oedacli> SET ADMINNET NAME='exa01adm03vm01.example.com' IP='10.xxx.xx.x'
oedacli> SET CLIENTNET NAME='exa01client03vm01.example.com'
IP='10.xxx.xx.x'
oedacli> SET PRIVNET NAME1='exa01adm03vm01-priv' IP1='192.168.16.8'
NAME2='exa01adm01vm03-priv2' IP2='192.168.16.9'
oedacli> SET VIPNET NAME='exa01client03m01-vip.example.com'
IP='10.xxx.xx.x'
oedacli> SAVE ACTION
oedacli> MERGE ACTIONS
oedacli> SAVE FILE NAME='exa01-cloned-node3-rac.xml'
oedacli> DEPLOY ACTIONS
```

Related Topics

- Expanding an Oracle VM RAC Cluster on Exadata Using OEDACLI
- Expanding an Oracle RAC Cluster in Oracle Linux KVM Using OEDACLI

4.3.10.2 DELETE GUEST

This command deletes an existing Oracle VM guest, and removes all of its configuration: instances, software installations, storage, and users from the cluster.

Syntax

```
DELETE GUEST WHERE SRCNAME = node_name
    [STEPNAME = stepname ]
```

Arguments

- SRCNAME is used to specify the host name of an existing DomU that you want to delete.
- The WHERE STEPNAME clause allows you to undo individual steps of the clone operation.

The value of *stepname* can be one of the following:

- CREATE_GUEST
- CREATE_USERS
- CELL_CONNECTIVITY
- ADD_NODE
- EXTEND_DBHOME
- ADD_INSTANCE

When used with DELETE GUEST, the STEPNAME clauses should be used in the reverse order of CLONE GUEST steps, and you should not skip any steps. For

example, if you want to undo the `ADD_NODE` and `ADD_INSTANCE` steps, you would use the following commands:

```
DELETE GUEST WHERE SRCNAME = node_name STEPNAME=ADD_INSTANCE
DELETE GUEST WHERE SRCNAME = node_name STEPNAME=EXTEND_DBHOME
DELETE GUEST WHERE SRCNAME = node_name STEPNAME=ADD_NODE
```

4.3.11 ILOM

You can alter or list the ILOMS in the loaded `es.xml` file.

- [ALTER ILOM](#)
- [LIST ILOMS](#)

4.3.11.1 ALTER ILOM

This command alters attributes of an ILOM.

Syntax

```
ALTER ILOM { DNSSERVERS='dns_servers' | NTPSERVERS= 'ntp_servers' |
    TIMEZONE = timezone }
WHERE {
    ID = ilom_id |
    HOSTNAME = hostname |
    ILOMNAME = ilom_hostname }
```

Arguments

You can specify the following attributes for the ILOM:

- `DNSSERVERS` : Specifies a comma-separated list of DNS servers that is enclosed in single quotes.
- `NTPSERVERS` : Specifies a comma-separated list of NTP servers that is enclosed in single quotes.
- `TIMEZONE` : Specifies a valid time zone for the ILOM.

The following arguments are available in the `WHERE` clause:

- `ID` : Specifies the `es.xml` ID for the target ILOM
- `HOSTNAME` : Specifies the current host name of the target machine
- `ILOMNAME` : Specifies the current host name of the ILOM

4.3.11.2 LIST ILOMS

This command lists the ILOMS in the `es.xml` file.

Syntax

```
LIST ILOMS
[ WHERE {
    ID = ilom_id |
```

```

HOSTNAME = hostname |
ILOMNAME = ilom_hostname } ]

```

Arguments

The following arguments are available in the `WHERE` clause:

- `ID` : Specifies the `es.xml` ID for the target ILOM
- `HOSTNAME` : Specifies the current host name of the target machine
- `ILOMNAME` : Specifies the current host name of the ILOM

Usage Notes

If you do not specify a `WHERE` clause, then this command lists all the ILOMs in the `es.xml` file.

4.3.12 MACHINE

You can alter or list the machines in the `es.xml` file.

- [ALTER MACHINE](#)
- [LIST MACHINES](#)

4.3.12.1 ALTER MACHINE

This command alters attributes of a compute node or storage cell in an engineered system.

Syntax

```

ALTER MACHINE { DNSSERVERS='dns_servers' |
GATEWAYADAPTER=gateway_adapter |
  HOSTNAMEADAPTER=hostname_adapter | NTPSERVERS='ntp_servers' |
  TIMEZONE = timezone }
WHERE {
  ID = machine_id |
  HOSTNAME = hostname |
  CLUSTERNUMBER = cluster_number COMPUTENUMBER = compute_number |
  CLUSTERNUMBER = cluster_number STORAGENUMBER = storage_number |
  CLUSTERNAME = cluster_name COMPUTENUMBER = compute_number |
  CLUSTERNAME = cluster_name STORAGENUMBER = storage_number |
  CLUSTERID = cluster_id COMPUTENUMBER = compute_number |
  CLUSTERID = cluster_id STORAGENUMBER = storage_number }

```

Arguments

You can specify the following attributes for the machine:

- `DNSSERVERS` : Specifies a comma-separated list of DNS servers that is enclosed in single quotes.
- `GATEWAYADAPTER` : Specifies the network adapter to use as the machine gateway. Valid values are `ADMIN` or `CLIENT`.

- **HOSTNAMEADAPTER** : Specifies the network adapter to use as the host name. Valid values are `ADMIN` or `CLIENT`.
- **NTPSERVERS** : Specifies a comma-separated list of NTP servers that is enclosed in single quotes.
- **TIMEZONE** : Specifies a valid time zone for the machine.

The following arguments are available in the `WHERE` clause:

- **ID** : Specifies the `es.xml` ID for the target machine
- **HOSTNAME** : Specifies the current host name of the target machine
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster
- **COMPUTENUMBER** : Specifies the compute node number in the cluster, starting at 1
- **STORAGENUMBER** : Specifies the storage cell number in the cluster, starting at 1

4.3.12.2 LIST MACHINES

This command lists the machine available in the `es.xml` file for an engineered system.

Syntax

```
LIST MACHINES
{ WHERE
  TYPE = type |
  ID = machine_id [ TYPE = type ] |
  CLUSTERNUMBER = cluster_number [ TYPE = type ] |
  CLUSTERNAME = cluster_name [ TYPE = type ] |
  CLUSTERID = cluster_id [ TYPE = type ] }
```

Arguments

The following arguments are available in the `WHERE` clause:

- **TYPE** : Specifies the machine type. The valid values are `COMPUTE`, `CELL`, `STORAGE`, `DOM0`, `GUEST`, or `DOMU`.
- **ID** : Specifies the `es.xml` ID for an individual machine
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster

Usage Notes

If you do not specify a `WHERE` clause, then this command lists all the machines defined in the `es.xml` file.

4.3.13 NETWORK

You can add, alter, delete, or list the networks defined in the XML configuration file.

- [ADD NETWORK](#)
- [ALTER NETWORK](#)
- [DELETE NETWORK](#)
- [LIST NETWORKS](#)

4.3.13.1 ADD NETWORK

This command adds a network record to a machine in the existing configuration.

Syntax

```
ADD NETWORK
  NETWORKTYPE=network_type HOSTNAME=hostname IP=ip_addr
  NETMASK=netmask DOMAINNAME=domain_name MASTER=master
  [ GATEWAY=gateway] [ SSHENABLED=ssh_enabled ] [ MAC=mac_addr ]
  [ LINKSPEED=nic_speed [ SLAVE='slave' ] [ PKEY=pkey ]
  [ PKEYNAME=pkey_name ] [ STATUS = status ] [ LACP=lACP ]
  [ VLANID=vlan_id ] [ INTERFACENAME=interface_name ]
  [ DESCRIPTION='network_description' ]
[ VSWITCHNETWORKPARAMS='vswitchnetwork_parameters' ]
  [ NATHOSTNAME=nat_host_name ] [ NATIP=nat_ip_addr ]
[ NATDOMAINNAME=nat_domain_name]
  [ NATNETMASK=nat_mask ]
WHERE {
  HOSTNAME = hostname |
  CLUSTERNAME=cluster_name [ COMPUTENUMBER=compute_number |
  STORAGENUMBER=storage_number] |
  CLUSTERNUMBER=cluster_name [ COMPUTENUMBER=compute_number |
  STORAGENUMBER=storage_number]
```

Arguments

You can specify the following attributes when adding a network:

- **NETWORKTYPE:** (Mandatory) Specifies the type of network, which can be ADMIN, CLIENT, BACKUP, PRIVATE, OTHER, INGESTVIP, REPLICATIONVIP, or ILOM.
- **HOSTNAME:** (Mandatory) Specifies the short host name for the network
- **IP:** (Mandatory) Specifies the IP address of the network
- **NETMASK:** (Mandatory) Specifies the netmask of the network
- **DOMAINNAME:** (Mandatory) Specifies the domain name of the network
- **MASTER:** (Mandatory) Specifies the master adapter name
- **GATEWAY:** Specifies the network subnet gateway
- **SSHENABLED:** A Boolean variable where TRUE indicates SSH should be enabled for this adapter
- **MAC:** Specifies the MAC address
- **LINKSPEED:** (Only for X7 2 or later socket compute nodes), either 10000 or 25000
- **SLAVE:** Specifies a space-separated list of slave devices for a bonded network, enclosed in single quotes

- **PKEY:** (For Private networks only) Specifies the pkey
- **PKEYNAME:** (For Private networks only) Specifies the pkey name
- **STATUS:** Specifies the status of the network
- **LACP:** A Boolean variable where `TRUE` indicates that LACP should be enabled
- **VLANID:** Specifies the VLAN Id if network is a VLAN network and not private
- **INTERFACENAME:** The name for the private network on the VLAN for a KVM guest and cell, `clre[0-1]` or `stre[0-1]`.
- **DESCRIPTION:** The text description for this network, surrounded by single quotes.
- **VSWITCHNETWORKPARAMS:** Parameters required for vswitches, surrounded by single quotes.
- **NATHOSTNAME:** Specifies the NAT host name
- **NATIP:** Specifies the NAT IP address
- **NATDOMAINNAME:** Specifies the NAT domain name
- **NATNETMASK:** Specifies the NAT netmask

The following arguments are available in the `WHERE` clause:

- **HOSTNAME:** Specifies the host name for an existing compute node
- **CLUSTERNUMBER:** Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME:** Specifies the name of an existing cluster
- **COMPUTENUMBER:** Specifies the number of an existing compute node in the cluster, starting at 1
- **STORAGENUMBER:** Specifies the umber of an existing storage cell in the cluster, starting at 1

4.3.13.2 ALTER NETWORK

This command modifies a network record for a machine in the existing configuration.

Syntax

```
ALTER NETWORK {
  [ HOSTNAME=hostname ] [ IP=ip_addr ] [ NETMASK=netmask ]
  [ DOMAINNAME=domain_name ] [ MASTER=master ] [ GATEWAY=gateway ]
  [ SSHENABLED=ssh_enabled ] [ MAC=mac_addr ] [ LINKSPEED=nic_speed ]
  [ SLAVE='slave' ] [ PKEY=pkey ] [ PKEYNAME=pkey_name ]
  [ STATUS = status ] [ LACP=lacp ] [ VLANID=vlan_id ]
  [ INTERFACENAME=interface_name ]
  [ DESCRIPTION='network_description' ]
  [ VSWITCHNETWORKPARAMS='vswitchnetwork_parameters' ]
  [ NATHOSTNAME=nat_host_name ]
  [ NATIP=nat_ip_addr ] [ NATDOMAINNAME=nat_domain_name ]
  [ NATNETMASK=nat_mask ] }
WHERE {
  ID = network_id |
  NETWORKHOSTNAME = network_hostname |
  NATHOSTNAME = nat_hostname |
```

```

    HOSTNAME = hostname NETWORKTYPE = network_type
[PRIVATEID=private_id]
    CLUSTERNAME = cluster_name { COMPUTENUMBER=compute_number |
STORAGENUMBER=storage_number }
    NETWORKTYPE = network_type [PRIVATEID=private_id] |
    CLUSTERNUMBER = cluster_number { COMPUTENUMBER=compute_number |
    STORAGENUMBER=storage_number }
    NETWORKTYPE = network_type [PRIVATEID=private_id] }

```

Arguments

You can specify the following attributes when altering a network:

- **HOSTNAME:** Specifies the short host name for the network
- **IP:** Specifies the IP address of the network
- **NETMASK:** Specifies the netmask of the network
- **DOMAINNAME:** Specifies the domain name of the network
- **MASTER:** Specifies the master adapter name
- **GATEWAY:** Specifies the network subnet gateway
- **SSHENABLED:** A Boolean variable where **TRUE** indicates SSH should be enabled for this adapter
- **MAC:** Specifies the MAC address
- **LINKSPEED:** (Only for X7 2 socket compute nodes), either 10000 or 25000
- **SLAVE:** Specifies a space-separated list of slave devices for a bonded network, enclosed in single quotes
- **PKEY:** (For Private networks only) Specifies the pkey
- **PKEYNAME:** (For Private networks only) Specifies the pkey name
- **STATUS:** Specifies the status of the network
- **LACP:** A Boolean variable where **TRUE** indicates that LACP should be enabled
- **VLANID:** Specifies the VLAN Id if network is a VLAN network and not private
- **INTERFACENAME:** The name for the private network on the VLAN for a KVM guest and cell, `clre[0-1]` or `stre[0-1]`.
- **DESCRIPTION:** The text description for this network, surrounded by single quotes.
- **VSWITCHNETWORKPARAMS:** Parameters required for vswitches, surrounded by single quotes.
- **NATHOSTNAME:** Specifies the NAT host name
- **NATIP:** Specifies the NAT IP address
- **NATDOMAINNAME:** Specifies the NAT domain name
- **NATNETMASK:** Specifies the NAT netmask

The following arguments are available in the **WHERE** clause:

- **ID:** Specifies the ID for the target network in the XML configuration file.

- NETWORKHOSTNAME: Specifies the host name for a network object, not the machine host name
- NATHOSTNAME: Specifies the NAT host name for an existing compute node
- HOSTNAME: Specifies the host name for an existing compute node
- NETWORKTYPE: Specifies the type of network, which can be ADMIN, CLIENT, BACKUP, PRIVATE, OTHER, INGESTVIP, REPLICATIONVIP, or ILOM.
- PRIVATEID: Specifies the ID of the private network
- CLUSTERNUMBER: Specifies the cluster number in the XML configuration file, starting at 1
- CLUSTERNAME: Specifies the name of an existing cluster
- COMPUTENUMBER: Specifies the number of an existing compute node in the cluster, starting at 1
- STORAGENUMBER: Specifies the number of an existing storage cell in the cluster, starting at 1

4.3.13.3 DELETE NETWORK

This command removes a network record for a machine.

Syntax

```
DELETE NETWORK
WHERE {
  ID = network_id |
  HOSTNAME = network_hostname NETWORKTYPE = network_type |
  CLUSTERNAME = cluster_name { COMPUTENUMBER=compute_number |
    STORAGENUMBER=storage_number } NETWORKTYPE = network_type |
  CLUSTERNUMBER = cluster_number { COMPUTENUMBER=compute_number |
    STORAGENUMBER=storage_number } NETWORKTYPE = network_type |
  CLUSTERID = cluster_id { COMPUTENUMBER=compute_number |
    STORAGENUMBER=storage_number } NETWORKTYPE = network_type }
```

Arguments

The following arguments are available in the WHERE clause:

- ID : Specifies the ID for the network object in the XML configuration file.
- HOSTNAME : Specifies the host name for an existing compute node
- NETWORKTYPE : Specifies the type of network, which can be ADMIN, CLIENT, BACKUP, PRIVATE, OTHER, INGESTVIP, REPLICATIONVIP, or ILOM.
- CLUSTERNAME : Specifies the name of an existing cluster
- CLUSTERNUMBER : Specifies the cluster number in the XML configuration file, starting at 1
- CLUSTERID : Specifies the cluster ID in the XML configuration file
- COMPUTENUMBER : Specifies the number of an existing compute node in the cluster, starting at 1

- **STORAGENUMBER** : Specifies the number of an existing storage cell in the cluster, starting at 1

4.3.13.4 LIST NETWORKS

This command lists the network details.

Syntax

```
LIST NETWORKS
[ WHERE {
  ID = network_id |
  HOSTNAME = network_hostname [ NETWORKTYPE = network_type ] |
  CLUSTERNAME = cluster_name { COMPUTENUMBER=compute_number |
  STORAGENUMBER=storage_number } [ NETWORKTYPE = network_type ] |
  CLUSTERNUMBER = cluster_number { COMPUTENUMBER=compute_number |
  STORAGENUMBER=storage_number } [ NETWORKTYPE = network_type ] }
CLUSTERID = cluster_ID { COMPUTENUMBER=compute_number |
  STORAGENUMBER=storage_number } [ NETWORKTYPE = network_type ] } ]
```

Arguments

The following arguments are available in the WHERE clause:

- **ID** : Specifies the ID for the network object in the XML configuration file.
- **HOSTNAME** : Specifies the host name for an existing compute node
- **NETWORKTYPE** : Specifies the type of network, which can be ADMIN, CLIENT, BACKUP, PRIVATE, or ILOM.
- **CLUSTERNAME** : Specifies the name of an existing cluster
- **CLUSTERNUMBER** : Specifies the cluster number in the XML configuration file, starting at 1
- **CLUSTERID** : Specifies the cluster ID in the XML configuration file
- **COMPUTENUMBER** : Specifies the number of an existing compute node in the cluster, starting at 1
- **STORAGENUMBER** : Specifies the number of an existing storage cell in the cluster, starting at 1

4.3.14 SCAN

You can add, alter, delete, or list the SCANS for a cluster.

- [ADD SCAN](#)
- [ALTER SCAN](#)
- [DELETE SCAN](#)
- [LIST SCANS](#)

4.3.14.1 ADD SCAN

This command adds a SCAN to the an individual cluster.

Syntax

```
ADD SCAN SCANNAME=scan_name [ SCANPORT=scan_port ]
SCANIPS='scan_ip_addrs'
WHERE
{ CLUSTERNAME = cluster_name |
  CLUSTERNUMBER = cluster_number |
  CLUSTERID = cluster_id }
```

Arguments

You can specify the following attributes for the new SCAN:

- **SCANNAME** : Specifies the DNS name for the SCAN
- **SCANPORT** : Specifies the port for the SCAN. The default value is 1521.
- **SCANIPS** : A comma-separated list of IP addresses for the SCAN, enclosed in single quotes.

The following arguments are available in the **WHERE** clause:

- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERNUMBER** : Specifies the cluster number in the XML configuration file, starting at 1
- **CLUSTERID** : Specifies the ID of the cluster in the XML configuration file

4.3.14.2 ALTER SCAN

This command modifies the attributes of a SCAN for an individual cluster.

Syntax

```
ALTER SCAN { SCANNAME=scan_name | SCANPORT=scan_port |
SCANIPS='scan_ip_addrs' }
WHERE
{ CLUSTERNAME = cluster_name |
  CLUSTERNUMBER = cluster_number |
  CLUSTERID = cluster_id }
```

Arguments

You can modify the following attributes for the SCAN:

- **SCANNAME** : Specifies the DNS name for the SCAN
- **SCANPORT** : Specifies the port for the SCAN.
- **SCANIPS** : A comma-separated list of IP addresses for the SCAN, enclosed in single quotes.

The following arguments are available in the **WHERE** clause:

- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERNUMBER** : Specifies the cluster number in the XML configuration file, starting at 1

- **CLUSTERID** : Specifies the ID of the cluster in the XML configuration file

4.3.14.3 DELETE SCAN

This command removes a SCAN for an individual cluster.

Syntax

```
DELETE SCAN
WHERE
{  CLUSTERNAME = cluster_name |
  CLUSTERNUMBER = cluster_number |
  CLUSTERID = cluster_id }
```

Arguments

The following arguments are available in the `WHERE` clause:

- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERNUMBER** : Specifies the cluster number in the XML configuration file, starting at 1
- **CLUSTERID** : Specifies the ID of the cluster in the XML configuration file

4.3.14.4 LIST SCANS

This command lists the SCAN details for all clusters or for an individual cluster.

Syntax

```
LIST SCANS
[ WHERE
{  SCANID = scan_id |
  CLUSTERNAME = cluster_name |
  CLUSTERNUMBER = cluster_number |
  CLUSTERID = cluster_id } ]
```

Arguments

The following arguments are available in the `WHERE` clause:

- **SCANID** : Specifies the ID of the SCAN object in the XML configuration file
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERNUMBER** : Specifies the cluster number in the XML configuration file, starting at 1
- **CLUSTERID** : Specifies the ID of the cluster in the XML configuration file

Usage Notes

If you do not specify a `WHERE` clause, then this command lists details about all the SCANS for all the clusters in the XML configuration file.

4.3.15 SSHKEYS

You can generate and deploy SSH keys for `root` user authentication on Oracle Exadata Rack.

- **DEPLOY SSHKEYS**
This command configures SSH key-based authentication on all hosts in the engineered system XML file (`es.xml`), or the host names specified in the command.
- **GENERATE SSHKEYS**
This command generates SSH key pairs for all hosts in the engineered system XML file (`es.xml`), or the host names specified in the command. The keys are saved in the `OEDA WorkDir`.
- **SET SSHKEYS**
This command enables and disables SSH key-based authentication for remote commands run by OEDACLI.

4.3.15.1 DEPLOY SSHKEYS

This command configures SSH key-based authentication on all hosts in the engineered system XML file (`es.xml`), or the host names specified in the command.

Syntax

```
DEPLOY SSHKEYS [ PASSWORD=root-password ] [ ENABLEPASSWORDLOGIN={true|false} ] [ HOSTNAMES="host1[,host2[,host3 ...]]" ]
```

Arguments

- `PASSWORD` optionally specifies the non-default `root` password for connecting to the hosts to set up SSH key access.
- `ENABLEPASSWORDLOGIN` optionally specifies whether to enable or disable password-based authentication for the `root` user after SSH key access is set up.
- `HOSTNAMES` optionally specifies the list of host names on which to deploy SSH keys.

Usage Notes

If the SSH key pairs are not located in the `OEDA WorkDir`, then the command also generates them.

Example 4-20 Deploy SSH keys for all hosts of an XML file

```
oedacli> LOAD FILE NAME=es.xml
SUCCESS - file loaded
edacli> DEPLOY SSHKEYS ENABLEPASSWORDLOGIN=false
```

Example 4-21 Deploy SSH keys for specified hosts

```
oedacli> DEPLOY SSHKEYS ENABLEPASSWORDLOGIN=false PASSWORD="password"
HOSTNAMES="dbm01adm05.example.com,dbm01adm06.example.com"
```

4.3.15.2 GENERATE SSHKEYS

This command generates SSH key pairs for all hosts in the engineered system XML file (es.xml), or the host names specified in the command. The keys are saved in the OEDA WorkDir .

Syntax

```
GENERATE SSHKEYS [ HOSTNAMES="host1[,host2[,host3 ...]]" ]
```

Arguments

- **HOSTNAMES** optionally specifies the list of host names for which to generate SSH keys.

Example 4-22 Generate SSH keys for all hosts of an XML file

```
oedacli> LOAD FILE NAME=es.xml
SUCCESS - file loaded
edacli> GENERATE SSHKEYS
Creating SSH Key Pairs for hosts [dbm01adm01.example.com,
dbm01adm02.example.com, ...]...
Generated SSH Keys for [dbm01adm01.example.com,
dbm01adm02.example.com, ...]
```

Example 4-23 Generate SSH keys for specified hosts

```
oedacli> GENERATE SSHKEYS
HOSTNAMES="dbm02adm05.example.com,dbm02adm06.example.com"
Creating SSH Key Pairs for hosts [dbm02adm06.example.com,
dbm02adm05.example.com]...
Generated SSH Keys for [dbm02adm06.example.com, dbm02adm05.example.com]
```

4.3.15.3 SET SSHKEYS

This command enables and disables SSH key-based authentication for remote commands run by OEDACLI.

Prerequisites

Before you can use this command, you must configure SSH key-based authentication on the target system. See [DEPLOY SSHKEYS](#).

Syntax

```
SET SSHKEYS ENABLE={true|false}
```

Arguments

- **ENABLE** specifies whether or not to enable SSH key-based authentication for remote commands run by OEDACLI.

4.3.16 SWITCH

You can alter or list the switches.

- [ALTER SWITCH](#)
- [LIST SWITCHES](#)

4.3.16.1 ALTER SWITCH

This command alters attributes of an individual switch.

Syntax

```
ALTER SWITCH { DNSSERVERS='dns_servers' | NTPSERVERS= 'ntp_servers' |  
    TIMEZONE = timezone }  
WHERE {  
    ID = switch_id |  
    HOSTNAME = switch_name }
```

Arguments

You can modify the following attributes for a switch:

- **DNSSERVERS** : Specifies a comma-separated list of DNS servers that is enclosed in single quotes.
- **NTPSERVERS** : Specifies a comma-separated list of NTP servers that is enclosed in single quotes.
- **TIMEZONE** : Specifies a valid time zone for the switch.

The following arguments are available in the **WHERE** clause:

- **ID** : Specifies the switch ID in the XML configuration file
- **HOSTNAME** : Specifies the host name of the switch

4.3.16.2 LIST SWITCHES

This command lists the switch details for all switches or for an individual switch.

Syntax

```
LIST SWITCHES  
[ WHERE {  
    ID = switch_id |  
    HOSTNAME = switch_name } ]
```

Arguments

The following arguments are available in the **WHERE** clause:

- **ID** : Specifies the switch ID in the XML configuration file
- **HOSTNAME** : Specifies the host name of the switch

Usage Notes

If you do not specify a WHERE clause, then this command lists the details for all switches.

4.3.17 VIP

You can add, alter, delete, or list the VIPs for the clusters.

- [ADD VIP](#)
- [ALTER VIP](#)
- [DELETE VIP](#)
- [LIST VIPS](#)

4.3.17.1 ADD VIP

This command adds a VIP for an individual node in a cluster.

Purpose

This command does not add additional VIPs. You should only use this command after previously issuing a `DELETE VIP` command.

Syntax

```
ADD VIP NAME=vip_name DOMAINNAME=domain_name IP=vip_ip_addr
WHERE {
  HOSTNAME = hostname |
  CLUSTERNAME = cluster_name COMPUTENUMBER = compute_number |
  CLUSTERNUMBER = cluster_number COMPUTENUMBER = compute_number |
  CLUSTERID = cluster_id COMPUTENUMBER = compute_number }
```

Arguments

You can specify the following attributes when adding a VIP:

- **NAME** : Specifies the short DNS name for the VIP
- **DOMAINNAME** : Specifies the domain name for the VIP
- **IP** : Specifies the IP address for the VIP

The following arguments are available in the WHERE clause:

- **HOSTNAME** : Specifies the host name for the compute node in the cluster
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster
- **COMPUTENUMBER** : Specifies the compute node number in the cluster, starting at 1

Usage Notes

When adding a VIP to a compute node, the node must not already have a VIP defined.

4.3.17.2 ALTER VIP

This command alters the attributes of a VIP for an individual node in a cluster.

Syntax

```
ALTER VIP { NAME=vip_name | DOMAINNAME=domain_name |  
            IP=vip_ip_addr }  
WHERE {  
    HOSTNAME = hostname |  
    CLUSTERNAME = cluster_name COMPUTENUMBER = compute_number |  
    CLUSTERNUMBER = cluster_number COMPUTENUMBER = compute_number |  
    CLUSTERID = cluster_id COMPUTENUMBER = compute_number }
```

Arguments

You can modify the following attributes for the VIP:

- **NAME** : Specifies the short DNS name for the VIP
- **DOMAINNAME** : Specifies the domain name for the VIP
- **IP** : Specifies the IP address for the VIP

The following arguments are available in the **WHERE** clause:

- **HOSTNAME** : Specifies the host name for the compute node in the cluster
- **CLUSTERNUMBER** : Specifies the cluster number in the `es.xml`, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the `es.xml` ID of the cluster
- **COMPUTENUMBER** : Specifies the compute node number in the cluster, starting at 1

4.3.17.3 DELETE VIP

This command removes a VIP from an individual node in a cluster.

Syntax

```
DELETE VIP  
WHERE {  
    ID = vip_id |  
    HOSTNAME = hostname |  
    CLUSTERNAME = cluster_name COMPUTENUMBER = compute_number |  
    CLUSTERNUMBER = cluster_number COMPUTENUMBER = compute_number |  
    CLUSTERID = cluster_id COMPUTENUMBER = compute_number }
```

Arguments

The following arguments are available in the **WHERE** clause:

- **ID** : Specifies the ID for the VIP in the XML configuration file
- **HOSTNAME** : Specifies the host name for the compute node in the cluster

- **CLUSTERNUMBER** : Specifies the cluster number in the XML configuration file, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the ID of the cluster in the XML configuration file
- **COMPUTENUMBER** : Specifies the compute node number in the cluster, starting at 1

4.3.17.4 LIST VIPS

This command lists the VIPs for all clusters or for an individual node in a cluster.

Syntax

```
LIST VIPS
[ WHERE {
  HOSTNAME = hostname |
  CLUSTERNAME = cluster_name [ COMPUTENUMBER = compute_number ] |
  CLUSTERNUMBER = cluster_number [ COMPUTENUMBER = compute_number ] |
  CLUSTERID = cluster_id [ COMPUTENUMBER = compute_number ] }
```

Arguments

The following arguments are available in the `WHERE` clause:

- **ID** : Specifies the ID for the VIP in the XML configuration file
- **HOSTNAME** : Specifies the host name for the compute node in the cluster
- **CLUSTERNUMBER** : Specifies the cluster number in the XML configuration file, starting at 1
- **CLUSTERNAME** : Specifies the name of the cluster
- **CLUSTERID** : Specifies the ID of the cluster in the XML configuration file
- **COMPUTENUMBER** : Specifies the compute node number in the cluster, starting at 1

Usage Notes

If you do not specify a `WHERE` clause, then this command lists the VIPs for all clusters.

5

Installing Oracle Exadata Database Machine or Oracle Exadata Storage Expansion Rack at the Site

This chapter describes how to move, install and configure the hardware through to powering on the system.

Note:

For ease of reading, the name "Oracle Exadata Rack" is used when information refers to both Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack.

- [Reviewing Safety Guidelines](#)
Before Oracle Exadata Rack arrives, the following safety precautions should be reviewed to ensure the site is safe, as well as ready for delivery.
- [Unpacking Oracle Exadata Rack](#)
- [Placing Oracle Exadata Rack in Its Allocated Space](#)
- [Acclimating the Oracle Exadata Database Machine](#)
Oracle recommends an acclimation period of 24 hours for the Oracle Exadata Database Machine.
- [Powering on the System the First Time](#)

5.1 Reviewing Safety Guidelines

Before Oracle Exadata Rack arrives, the following safety precautions should be reviewed to ensure the site is safe, as well as ready for delivery.

Failing to observe these precautions can result in injury, equipment damage, or malfunction.

- Do not block ventilation openings.
- Do not install Oracle Exadata Rack in a location that is exposed to direct sunlight or near a device that may become hot.
- Do not install Oracle Exadata Rack in a location that is exposed to excessive dust, corrosive gases, or air with high salt concentrations.
- Do not install Oracle Exadata Rack in a location that is exposed to frequent vibrations. Install Oracle Exadata Rack on a flat, level surface.
- Use a power outlet that uses proper grounding. When using shared grounding, the grounding resistance must not be greater than 10 ohms. Ensure that your facility

administrator or a qualified electrical engineer verifies the grounding method for the building, and performs the grounding work.

- Be sure that each grounding wire used for Oracle Exadata Rack is used exclusively for Oracle Exadata Rack. Also be sure to observe the precautions, warnings, and notes about handling that appear on labels on the equipment.
- Do not place cables under the equipment or stretch the cables too tightly.
- Do not disconnect power cords from the equipment while its power is on.
- If you cannot reach the connector lock when disconnecting LAN cables, then press the connector lock with a flathead screwdriver to disconnect the cable. You could damage the system board if you force your fingers into the gap rather than using a flathead screwdriver.
- Do not place anything on top of Oracle Exadata Rack or perform any work directly above the rack.
- Do not let the room temperature rise sharply, especially in winter. Sudden temperature changes can cause condensation to form inside Oracle Exadata Rack. Allow for a sufficient warm-up period prior to server operation.
- Do not install Oracle Exadata Rack near a photocopier, air conditioner, welding machine, or any other equipment that generates loud, electronic noises.
- Avoid static electricity at the installation location. Static electricity transferred to Oracle Exadata Rack can cause malfunctions. Static electricity is often generated on carpets.
- Confirm the supply voltage and frequency match the electrical ratings indicated for Oracle Exadata Rack.
- Do not insert anything into any Oracle Exadata Rack opening, unless doing so is part of a documented procedure. Oracle Exadata Rack contains high-voltage parts. If a metal object or other electrically-conductive object enters an opening in Oracle Exadata Rack, then it could cause a short circuit. This could result in personal injury, fire, electric shock, and equipment damage.
- When using single phase power distribution units (PDUs), note the following:
 - PDU A input 0 and PDU B input 2 must be on the same phase.
 - PDU A input 1 and PDU B input 1 must be on the same phase.
 - PDU A input 2 and PDU B input 0 must be on the same phase.

The inputs are labeled where they come out of the PDU. Connecting cables as described ensures the phases are balanced on both sides, A and B, in case of a failover.


 **See Also:**

- *Important Safety Information for Sun Hardware Systems (816-7190)* included with the rack
- *Oracle Engineered System Safety and Compliance Guide, Compliance Model No.: ESY27*

5.2 Unpacking Oracle Exadata Rack

The unpacking location should be determined during the site planning process. The following image shows the Oracle Exadata Rack shipping crate.

Figure 5-1 Oracle Exadata Rack in Shipping Crate



- [Contents of the Shipping Kit for X7 and Later Racks](#)
- [Contents of the Shipping Kit for X6 and Earlier Racks](#)
- [Removing Oracle Exadata Rack from the Shipping Crate](#)
Use these steps to unpack Oracle Exadata Rack and prepare it for moving to the installation site.

5.2.1 Contents of the Shipping Kit for X7 and Later Racks

The Oracle Exadata Rack shipping kit contains the following tools and equipment to install and service the Engineered System rack:

- Unpacking Tools (provided in rack ship kit)
 - 17 mm and 13 mm open-ended, double-sided wrench
 - Allen L-Key hex 6mm wrench, 5-3/4 inches in length
 - Allen L-Key hex 8mm wrench, 6-3/8 inches in length
- Rack Setup Tools and Accessories (provided in rack ship kit)
 - 32 M6 cage nuts and washers
 - 32 M6 pan head screws
 - 8 M5 screws
 - 8 self-tapping Torx screws
 - 4 PDU brackets
 - 4 M5 cage nuts
 - 4 grounded power cables
 - 18 mm and 16 mm open-ended, double-sided wrench
 - T-40 Torx screwdriver
 - T-25 Torx screwdriver
 - T-15 Torx screwdriver
 - 8 mm and 6 mm Allan wrench
 - Cage nut tool
 - Keys to the front door, rear door, and side panel locks



Note:

The following installation tools are not included in the shipping kit:

- No. 2 Phillips screwdriver
- Antistatic wrist strap

5.2.2 Contents of the Shipping Kit for X6 and Earlier Racks

The Oracle Exadata Rack shipping kit contains the following tools and equipment to install and service the Engineered System rack:

- 16 mm long No. 2 Phillips screw
- T30 Torx cranked wrench key
- T25 Torx cranked wrench key
- 6 mm hexagon Allen wrench key

- SW 12 mm single-headed wrench
- 2 square jumper brackets with 4 M5 Torx screws
- 2 cable management hooks with 4 spring nuts
- Side panel removal tool
- Keys to the front door, rear door, and side panel locks
- 32 M6 cage nuts
- 32 M6 screws
- Cage nut mounting tool
- SW 17 mm single-headed wrench is included on the shipping pallet

 **Note:**

The following items are not included in the shipping kit:

- No. 2 Phillips screwdriver
- Antistatic wrist strap


5.2.3 Removing Oracle Exadata Rack from the Shipping Crate

Use these steps to unpack Oracle Exadata Rack and prepare it for moving to the installation site.

 **WARNING:**

Oracle strongly recommends that you use professional movers when unpacking and installing Oracle Exadata Rack. Rocking or tilting the rack can cause it to fall over and cause serious personal injury or death.

1. Unpack Oracle Exadata Rack carefully from the packaging and shipping pallet as follows:

 **Note:**

Before unpacking the rack from the shipping carton, refer to the labels on the carton and to the instructions that they provide.

- a. Remove the shipping carton bands.
 - b. Remove the carton top.
 - c. Remove the carton sides and inner top.
2. Remove the shipping kit.
 3. Attach the ramps to the shipping pallet as follows:

- a. Remove the ramps from the pallet sides.
 - b. Obtain the parts bag from inside the cabinet.
 - c. Adjust the leveling bolts on the ramps and connect the ramps to the pallet wheel track.
 - d. Verify that leveling feet are up before moving rack.
4. Carefully roll Oracle Exadata Rack off the shipping pallet as follows:
 - a. Unfasten the exterior shipping brackets from the pallet.
 - b. Unfasten the interior shipping brackets from the pallet.

 **Note:**

Use care when removing the shipping brackets from underneath Oracle Exadata Rack. Access to the inside shipping brackets might be limited.

- c. Roll the cabinet down the ramps to the level floor. Oracle recommends having three people available to move the rack down the ramp: two people on both sides to help guide the rack and one person in back.

The rack's front casters are fixed, so you must steer using the rear casters. You can maneuver the rack safely by pushing it from behind.
5. Save the shipping brackets used to secure the rack to the shipping pallet. Do not dispose of these brackets, because you cannot order replacement brackets.
 6. Recycle the packaging properly. Follow local laws and guidelines to dispose of the material.

 **See Also:**

- *Sun Rack II User's Guide* at <https://docs.oracle.com/cd/E19657-01/html/E29153/index.html>
- *Oracle Rack Cabinet 1242 User's Guide* at http://docs.oracle.com/cd/E85660_01/html/E87280/index.html for X7 or later systems

5.3 Placing Oracle Exadata Rack in Its Allocated Space

This section describes how to position, stabilize, and ground Oracle Exadata Rack.

- [Moving Oracle Exadata Rack](#)
- [Stabilizing Oracle Exadata Rack](#)

After moving Oracle Exadata Rack to the installation site, stabilize the rack to ensure that it does not move or tip over.
- [Attaching a Ground Cable \(Optional\)](#)


5.3.1 Moving Oracle Exadata Rack

The following procedure describes how to move Oracle Exadata Rack:

1. Ensure the doors are closed and secured.
2. Ensure the leveling and stabilizing feet on the rack are raised and out of the way.
3. Push Oracle Exadata Rack from behind to the installation site.

When moving Oracle Exadata Rack to the installation site, the front casters do not turn; you must steer the unit by moving the rear casters. You can safely maneuver Oracle Exadata Rack by carefully pushing it. [Figure 5-2](#) shows the correct way to push the rack.

Figure 5-2 Carefully Push Oracle Exadata Rack from Behind



Oracle recommends having two people to move the rack: one person in front and one person in back to help guide the rack. When transporting configured racks from one location to another, take care to move them slowly, 0.65 meters per second (2.13 feet per second) or slower.

Carefully examine the transportation path. Avoid obstacles such as doorways or elevator thresholds that can cause abrupt stops or shocks. Go around obstacles by using ramps or lifts to enable smooth transport.

 **WARNING:**

- Never attempt to move Oracle Exadata Rack by pushing on the side panels. Pushing on the side panels can tip the rack over. This can cause serious personal injury or death as well as damage to the equipment.
- Never tip or rock Oracle Exadata Rack because the rack can fall over.

5.3.2 Stabilizing Oracle Exadata Rack

After moving Oracle Exadata Rack to the installation site, stabilize the rack to ensure that it does not move or tip over.

Oracle does not recommend using the shipping brackets to permanently mount the rack to the floor.

 **Caution:**

Shipping brackets are not for use for bracing or anchoring the rack during seismic events.

You can stabilize the rack by extending the rack leveling feet.

- [Stabilize Oracle Exadata Rack with Leveling Feet](#)
- [Stabilize Oracle Exadata Rack X7 and Later with Leveling Feet](#)
The rack contains four leveling feet that can be lowered to stabilize the rack.

5.3.2.1 Stabilize Oracle Exadata Rack with Leveling Feet


The Oracle Exadata Rack contains four leveling feet that can be lowered to stabilize the rack. The leveling feet can be used even when the rack is permanently secured to the floor. The four leveling feet should share the rack load with at least one other support. The following is an acceptable method to share the rack load:

- Casters and leveling feet

To adjust the leveling feet, do the following:


1. Locate the four leveling feet located at the bottom corners of the Oracle Exadata Rack. [Figure 5-3](#) shows the location of the leveling feet on the bottom of the rack.

Figure 5-3 Location of Leveling Feet on Oracle Exadata Rack



- 1: Distance from the edge of the mounting feet to the side of the rack is 33.75 mm (1.33 inches)
 - 2: Width from the outside edges of the leveling feet is 532.5 mm (20.96 inches)
 - 3: Width from the inside edges of the leveling feet is 429 mm (16.89 inches)
 - 4: Distance from the edge of the feet to the front rack surface is 73.75 mm (2.90 inches)
 - 5: Depth of the outside edges of the leveling feet is 1058.5 mm (41.67 inches)
 - 6: Distance from the edge of the leveling feet to the rear rack surface is 32.5 mm (1.28 inches)
 - 7: Distance from the center of front casters to the side of the rack is 86.7 mm (3.41 inches)
 - 8: Width between the center of the front casters is 426.6 mm (16.80 inches)
 - 9: Distance from the center of the rear casters to the rear of the rack is 173.7 mm (6.83 inches)
 - 10: Depth between the front and rear casters is 828.6 mm (32.62 inches)
 - 11: Distance between the rear casters and the rear of the rack is 162.4 mm (6.39 inches)
 - 12: Distance from the center of rear casters to the side of the rack is 96.4 mm (3.80 inches)
 - 13: Width between the center of the rear casters is 407.2 mm (16.03 inches)
2. Lower the leveling feet to the floor to share the load with the casters, as shown in [Figure 5-4](#) using the SW 12 mm wrench.

Figure 5-4 Securing Oracle Exadata Rack Using the Leveling Feet



5.3.2.2 Stabilize Oracle Exadata Rack X7 and Later with Leveling Feet


The rack contains four leveling feet that can be lowered to stabilize the rack.

The leveling feet can be used even when the rack is permanently secured to the floor. The four leveling feet share the load with the casters. This increases the footprint of the rack which improves stability and helps prevent rack movement.

To adjust the leveling feet, do the following:

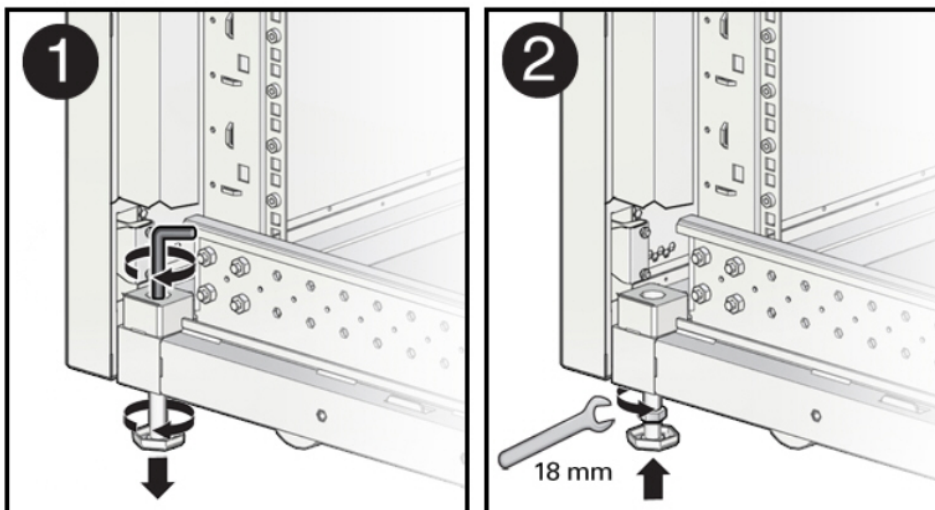
1. Locate the four leveling feet located at the bottom corners of Oracle Exadata Rack. The following figure shows the location of the leveling feet on the bottom of Oracle Exadata Rack. The image dimensions include the doors attached to the rack.

Figure 5-5 Location of Leveling Feet on Oracle Exadata Rack X7 or later



- 1: Distance from the edge of the feet to the Front door surface is 58 mm (2.3 inches)
 - 2: Depth of the outside edges of the leveling feet is 1113 mm (43.8 inches)
 - 3: Distance from the edge of the leveling feet to the rear door surface is 26 mm (1.0 inches)
 - 4: Distance from the center of rear casters to the side of the rack is 96.5 mm (3.8 inches)
 - 5: Width between the center of the rear casters is 407 mm (16.0 inches)
 - 6: Distance from the center of the rear casters to the rear door surface is 142 mm (5.6 inches)
 - 7: Depth between the front and rear casters is 889 mm (35.0 inches)
 - 8: Distance between the center of the front casters and the Front door surface is 166 mm (6.5 inches)
 - 9: Width from the outside edges of the leveling feet is 583 mm (23.0 inches)
 - 10: Width from the inside edges of the leveling feet is 513 mm (20.2 inches)
 - 11: Distance from the edge of the mounting feet to the side of the rack is 8.5 mm (0.3 inches)
 - 12: Width between the center of the front casters is 410 mm (16.1 inches)
 - 13: Distance from the center of front casters to the side of the rack is 95 mm (3.7 inches)
2. Stabilize the rack by lowering the four leveling feet to the floor, using a 6-mm hex wrench or your fingers, as shown in the following image.

Figure 5-6 Securing Oracle Exadata Rack Using the Leveling Feet



 **See Also:**

Oracle Rack Cabinet 1242 User's Guide at http://docs.oracle.com/cd/E85660_01/html/E87280/index.html for the rack specifications

5.3.3 Attaching a Ground Cable (Optional)

The Oracle Exadata Rack power distribution units (PDUs) achieve earth ground through their power cords. Final chassis ground is achieved by way of the ground prong when you connect the power cord to a socket. For additional grounding, attach a chassis earth ground cable to Oracle Exadata Rack. The additional ground point enables electrical current leakage to dissipate more efficiently.

 **WARNING:**


The PDU power input lead cords and the ground cable must reference a common earth ground. If they do not, then a difference in ground potential can be introduced. If you are unsure of your facility's PDU receptacle grounding, then do not install a ground cable until you confirm that there is a proper PDU receptacle grounding. If a difference in ground potential is apparent, then you must take corrective action.

 **Note:**

A grounding cable is not shipped with the system.

1. Ensure the installation site has properly grounded the power source in the data center. The facility PDU must have earth ground.
2. Ensure all grounding points, such as raised floors and power receptacles, reference the facility ground.
3. Ensure that direct, metal-to-metal contact is made for this installation. During manufacturing, the ground cable attachment area might have been painted or coated.
4. Attach the ground cable to one of the attachment points located at the bottom rear of the system frame as shown in [Figure 5-7](#). The attachment point is an adjustable bolt that is inside the rear of the Oracle Exadata Rack cabinet on the right side.

Figure 5-7 Earth Ground Attachment Bolt Location



 **See Also:**

Oracle Exadata Database Machine Extending and Multi-Rack Cabling Guide if Oracle Exadata Rack will be connected to another Oracle Exadata Rack.

5.4 Acclimating the Oracle Exadata Database Machine

Oracle recommends an acclimation period of 24 hours for the Oracle Exadata Database Machine.

This applies to all hardware, including racks and cables. Acclimation to the surrounding environment prevents damage as a result of condensation.

5.5 Powering on the System the First Time

Before powering on the system for the first time, it is necessary to inspect the machine, and connect the power cords.

- [Inspecting the Machine After it is in Place](#)
This procedure describes how to visually examine the Oracle Exadata Rack physical system after it is in place, but before power is supplied.
- [Connecting Power Cords](#)
- [Powering on Oracle Exadata Rack](#)
This procedure describes how to power on Oracle Exadata Rack.

5.5.1 Inspecting the Machine After it is in Place

This procedure describes how to visually examine the Oracle Exadata Rack physical system after it is in place, but before power is supplied.

1. Check the rack for damage.
2. Check the rack for loose or missing screws.
3. Check Oracle Exadata Rack for the ordered configuration. Refer to the Customer Information Sheet (CIS) on the side of the packaging.
4. Check that all cable connections are secure and firmly in place as follows:
 - a. Check the power cables. Ensure that the correct connectors have been supplied for the data center facility power source.
 - b. Check the network data cables.
5. Check the site location tile arrangement for cable access and airflow.
6. Check the data center airflow that leads in to the front of Oracle Exadata Rack.

Related Topics


- [Ventilation and Cooling Requirements](#)
Always provide adequate space in front and behind the rack to allow for proper ventilation.

5.5.2 Connecting Power Cords

The following procedure describes how to connect power cords to Oracle Exadata Rack:

1. Open the rear cabinet door.
2. Ensure that the correct power connectors have been supplied.
3. Unfasten the power cord cable ties. The ties are for shipping only and are no longer needed.
4. Route the power cords to the facility receptacles as shown in [Figure 5-8](#).
5. Secure the power cords in bundles as shown in [Figure 5-8](#).
6. Plug the PDU power cord connectors into the facility receptacles. Ensure the breaker switches are in the OFF position before connecting the power cables.

Figure 5-8 Power Cord Routing Example



5.5.3 Powering on Oracle Exadata Rack


This procedure describes how to power on Oracle Exadata Rack.

 **Note:**

Oracle Exadata Database Machine X2-8 Full Rack database servers may take up to 15 minutes to start through the normal BIOS POST tests.

1. Switch on the PDU B circuit breakers, one at a time. PDU B is on the right side of the rack when viewed from the rear. The circuit breakers are on the rear of the Oracle Exadata Rack cabinet as shown in the following figure. Press the ON (I) side of the toggle switch. When the breaker is in the ON position, the breakers are flush with the side of the PDU.

Figure 5-9 PDU Switch Locations



2. Verify the expected power LEDs are on. The LEDs are located as follows:
 - Database servers: Right-hand LED
 - Storage Servers: Top LED
 - Cisco Ethernet switch: Left LED (viewed from front) is green, and the other LEDs are red
 - InfiniBand Network Fabric switches: Left LED (viewed from front) labeled PS0, or right LED (viewed from front) labeled PS1
 - RoCE Network Fabric switches: Left LED (viewed from front) labeled P1, or right LED (viewed from front) labeled P2
 - KVM switch on Oracle Exadata Database Machine X2-2: Lower LED B
 - KMM switch on Oracle Exadata Database Machine X2-2: Only supplied by PDU B
3. Connect power to PDU A. Ensure the breaker switches are in the OFF position before connecting the power cables.
4. Switch on the PDU A circuit breakers, one at a time.
5. Verify the expected LEDs are on.
6. Perform a visual check of all cable connections in the rack. Do not press every connector to verify connection.
7. Verify the OK LED is blinking standby for all systems. The OK LED blinks on for 0.1 seconds, once every 3 seconds when in standby mode. The OK LED does not blink when the ILOM is starting. The LED remains dark until it goes to standby mode after 2 to 3 minutes.

Related Topics

- [Reviewing Safety Guidelines](#)
Before Oracle Exadata Rack arrives, the following safety precautions should be reviewed to ensure the site is safe, as well as ready for delivery.

6

Configuring Oracle Exadata Database Machine

This chapter describes how to configure the system, accounts, and software for Oracle Exadata Database Machine.

Note:

- For ease of reading, the name "Oracle Exadata Rack" is used when information refers to both Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack.
- The procedures in this chapter are applicable to Oracle Exadata Storage Expansion Rack as well as Oracle Exadata Database Machine.
- The procedures in this chapter use the information from Oracle Exadata Database Machine Deployment Assistant (OEDA). Be sure to run OEDA before starting the procedures in this chapter.

- [Verifying the Network Configuration Prior to Configuring the Rack](#)
Use the `checkip.sh` script to ensure there are no IP address conflicts between the existing network and your new Oracle Exadata Rack.
- [Configuring the RDMA Network Fabric Switch](#)
You must perform an initial configuration of the RDMA Network Fabric switch.
- [Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack](#)
- [Configuring the Cisco Ethernet Switch](#)
You must perform an initial configuration of the Cisco Ethernet switch.
- [Configuring the Power Distribution Units](#)
The power distribution units (PDUs) are configured with a static IP address to connect to the network for monitoring.
- [Configuring the Threshold Settings for the Power Distribution Units](#)
- [Checking Exadata Storage Servers](#)
After powering up the rack, check the Oracle Exadata Storage Servers.
- [Checking Oracle Exadata Database Servers](#)
Before installing the software, check the configuration of the Oracle Exadata Database Servers.
- [Performing Additional Checks and Configuration](#)
Perform these additional checks and steps to verify the configuration of your Oracle Exadata Rack.
- [Verifying the RoCE Network Fabric Configuration](#)
This procedure describes how to verify the RoCE Network Fabric configuration.

- [Verifying the InfiniBand Network Fabric Network](#)
This procedure describes how to verify the InfiniBand Network Fabric network.
- [Imaging a New System](#)
You have different options for preparing an Oracle Exadata system for deployment.
- [Performing Initial Elastic Configuration of Oracle Exadata Database Machine](#)
This topic provides background information on elastic configuration and describes how to perform the initial configuration of Oracle Exadata Database Machine.
- [Adding Additional Elastic Nodes to an Existing Rack](#)
You can add nodes to an existing rack using elastic configuration.
- [Using the OEDA setuprootssh Utility](#)
You can set up key-based authentication for the `root` user by using the `setuprootssh.sh` utility included with OEDA.
- [Loading the Configuration Information and Installing the Software](#)
You can use Oracle Exadata Deployment Assistant to configure your rack, or you can do it manually.
- [Using the OEDA changePassword Utility](#)
You can change the password of the `grid` and `oracle` users across all clusters by using the `changePassword.sh` utility included with OEDA.
- [Installing Oracle Enterprise Manager Cloud Control](#)
- [Adding a VM Cluster to Oracle Exadata Database Machine Using OEDA](#)
You can use Oracle Exadata Deployment Assistant (OEDA) to add a virtual machine (VM) cluster to an existing Oracle Exadata Database Machine deployment.

Related Topics

- [Oracle Sun Database Machine Setup/Configuration Best Practices \(My Oracle Support Doc ID 1274318.1\)](#)

6.1 Verifying the Network Configuration Prior to Configuring the Rack

Use the `checkip.sh` script to ensure there are no IP address conflicts between the existing network and your new Oracle Exadata Rack.

The `checkip.sh` script performs a pre-installation check to verify that the IP addresses and host names that you specified in Oracle Exadata Deployment Assistant (OEDA) are defined in the DNS, that the NTP servers and gateways are available, and that private addresses are not pingable. Running this script before the hardware arrives help to avoid additional delays that would be caused by misconfigured network services, such as Domain Name System (DNS) and NTP.

The `checkip.sh` script is created in a format that matches the operating system of the client on which you ran OEDA. Because this script is run before the engineered system rack has arrived, you typically do not run this script on an engineered system server, but on a client. The client must have access to the same network where the engineered system will be deployed. The script is also available in the ZIP file generated by OEDA.

1. On the client where OEDA was run, copy the `checkip.sh` script generated by OEDA and the XML file `CustomerName_hostname.xml` to the same directory (one directory level up) as the OEDA `config.sh` script.
2. Run the `checkip.sh` script on the client machine or existing server.

Use a command similar to the following, where `configuration_file` is the name of the configuration generated by the Oracle Exadata Deployment Assistant for the rack being installed.

```
# ./checkip.sh -cf configuration_file
```

If the command is run from a Microsoft Windows machine, then the command is `checkip.cmd`.

If this engineered system rack is an addition for an existing installation, then run the `checkip.sh` script from an existing engineered system server. This enables the script to identify in-use IP addresses in the fabric. Not identifying existing IP addresses may cause IP collisions after installation of the new engineered system rack. To create a `checkip.sh` that can run on an existing server, you must run OEDA on a server or client that uses the same operating system as the existing engineered system server.

The output from the script is a file that contains status messages such as `GOOD` or `ERROR`.

If there are conflicts that you are unable to resolve, then work with your assigned Oracle representative to correct the problems.

6.2 Configuring the RDMA Network Fabric Switch

You must perform an initial configuration of the RDMA Network Fabric switch.

- [Configuring Sun Datacenter InfiniBand Switch 36 Switch](#)
- [Configuring the Cisco Nexus C9336C-FX2 Switch](#)
The RoCE Network Fabric switch supplied with the engineered system rack is minimally configured during installation.
- [Configuring the Cisco Nexus C9336C-FX2 Switches to Enable Exadata Secure RDMA Fabric Isolation](#)
You can configure your RoCE Network Fabric switches to enable Exadata Secure RDMA Fabric Isolation.

6.2.1 Configuring Sun Datacenter InfiniBand Switch 36 Switch

The following procedure describes how to configure the Sun Datacenter InfiniBand Switch 36 switches.

1. Log in to the first database server as the `root` user. The first database server is the lowest database server in the rack, which is rack position U16.

 **Note:**

If you do not have the password for the `root` user, then contact Oracle Support Services.

2. Use SSH to log in to the Sun Datacenter InfiniBand Switch 36 switch as the `ilom-admin` user. Find the default IP address for the Sun Datacenter InfiniBand Switch 36 switch for your system in "[Default IP Addresses](#)". The following is an example of the command:

```
ssh ilom-admin@192.168.1.201
```

3. Use `spsh` to open the Integrated Lights Out Manager (ILOM) command line interface.
4. Configure the network interface using the following commands:

```
set /SP/network pendingipdiscovery=static
pendingipaddress=pending_ip \
pendingipgateway=pending_gw pendingipnetmask=pending_nm

set /SP/network commitpending=true
```

In the preceding commands, *pending_ip*, *pending_gw*, and *pending_nm* are IP addresses defined by the network administrator.

5. Use the `exit` command to exit the interface.
6. Edit the `/etc/hosts` file to set the IP address and host name using an editor such as `vi`. The following is an example of the updated file:

```
#Do not remove the following link, or various programs
#that require network functionality will fail.
127.0.0.1    localhost.localdomain localhost
10.7.7.32   dm01sw-ib1-ib2.example.com trnasw-ib2
```

It is required that the second line has both fully-qualified and non-fully-qualified names.

7. Use the ILOM interface to configure the IP address, host name, NTP servers, and DNS servers.
8. Restart the switch.
9. Examine the firmware version using the following command:

```
# version
```

Refer to My Oracle Support Note 888828.1 for the current firmware version.

10. Check the health of the switch using the following command:

```
# showunhealthy

OK - No unhealthy sensors
```

11. Run the environment test using the following command:

```
# env_test

NM2 Environment test started:
Starting Voltage test:
Voltage ECB OK
Measured 3.3V Main = 3.28 V
Measured 3.3V Standby = 3.42 V
Measured 12V =12.06 V
Measured 5V =5.03 V
Measured VBAT =3.06 V
Measured 2.5V =2.53 V
Measured 1.8V =1.79 V
Measured I4 1.2V =1.22 V
Voltage test returned OK
Starting PSU test:
PSU 0 present
PSU 1 present
PSU test returned OK
Starting Temperature test:
Back temperature 30.50
Front temperature 33.88
ComEx temperature 34.12
I4 temperature 56,
maxtemperature 57
Temperature test returned OK
Starting FAN test:
Fan 0 not present
Fan 1 running at rpm 12946
Fan 2 running at rpm 12684
Fan 3 running at rpm 12558
Fan 4 not present
FAN test returned OK
Starting Connector test:
Connector test returned OK
Starting I4 test:
I4 OK
All I4s OK
I4 test returned OK
NM2 Environment test PASSED
```

12. Enable the InfiniBand Subnet Manager using the following command:

```
# enablesm
```

 **Note:**

If you get an error indicating the InfiniBand Subnet Manager is already running, then restart it as follows:

```
# disablesm  
  
# enablesm
```

13. Verify the IP address is correct using the following command:

```
# ifconfig eth0  
  
eth0    Link encap:Ethernet HWaddr 00:E0:4B:2A:07:2B  
        inet addr:172.16.10.32 Bcast:172.16.10.255  
        Mask:255.255.255.0  
        inet6 addr:fe80::2e0:3a00:fe2a:61e/64 Scope:Link  
        UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1  
        RX packets:11927 errors:0 dropped:0 overruns:0 frame:0  
        TX packets:89 errors:0 dropped:0 overruns:0 carrier:0  
        collisions:0 txqueuelen:1000  
        RX bytes:720262 (703.3 KiB)TX bytes:11402 (11.1 KiB)
```

14. Verify the host name is correct using the following command:

```
# hostname  
  
dm01sw-ib2.example.com
```

15. Set the Subnet Manager Master as described in "[Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack](#)". This step is needed for Oracle Exadata Database Machine Half Rack and Oracle Exadata Database Machine Full Rack.
16. Log out from the InfiniBand switch using the following command:

```
# exit
```

 **Note:**

If the network settings did not appear after restart, then power cycle the switch by removing both power cords for one minute.

Related Topics

- [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID 888828.1\)](#)

6.2.2 Configuring the Cisco Nexus C9336C-FX2 Switch

The RoCE Network Fabric switch supplied with the engineered system rack is minimally configured during installation.

During initial system configuration, you can reset and configure the switch.

1. Connect from the RoCE Network Fabric switch serial console to a laptop or similar device using the available RJ45 cable.
2. Ensure the terminal session is recorded on the laptop by logging the output.
The output can be used as a reference that the switch has been configured correctly.
3. Power on the switch.
4. Log in as the `admin` user.

```
User Access Verification
dbm0sw-rocea0 login: admin
Password: *****
```



Note:

If you do not have the password for the `admin` user, then contact Oracle Support Services.

5. Erase the existing configuration.

```
dbm0sw-rocea0# write erase
```

```
Warning: This command will erase the startup-configuration.
```

```
Do you wish to proceed anyway? (y/n) [n] y
```

6. Restart the system so you can perform the automated setup.

```
dbm0sw-rocea0# reload
```

```
This command will reboot the system. (y/n)? [n] y
```

```
2017 Aug 31 01:09:00 dbm0sw-rocea0 %$ VDC-1 %$ %PLATFORM-2-
PFM_SYSTEM_RESET: Manual system restart from Command Line Interface
```

```
CISCO SWITCH Ver7.59
Device detected on 0:1:2 after 0 msec
...
```

7. Switch to normal setup and, when asked if you want to enforce secure password standard, enter `no`, then enter a new password for the `admin` user.

```
Running S93thirdparty-script...

Populating conf files for hybrid sysmgr ...
Starting hybrid sysmgr ...
inserting /isan/lib/modules/klm_cisco_nb.o ... done

Abort Auto Provisioning and continue with normal setup ? (yes/no)
[n]: yes

      ---- System Admin Account Setup ----

Do you want to enforce secure password standard (yes/no) [y]: no

Enter the password for "admin":
Confirm the password for "admin":
```

8. When the Basic System Configuration Dialog appears, choose to enter the basic configuration dialog.

```
      ---- Basic System Configuration Dialog VDC: 1 ----

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Please register Cisco Nexus9000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus9000 devices must be registered to receive
entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): yes
```

9. In the basic configuration, you can use the default inputs until asked to enter the switch name.

In this example, the switch has a name of `test123sw-rocea0`.

```
Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : test123sw-rocea0
```

10. Respond `yes` when asked to configure Out-of-band management configuration, and specify appropriate network addresses when prompted.

```
Continue with Out-of-band (mgmt0) management configuration?
(yes/no) [y]: yes
Mgmt0 IPv4 address : 100.104.10.21
```

```
Mgmt0 IPv4 netmask : 255.255.248.0
Configure the default gateway? (yes/no) [y]:
IPv4 address of the default gateway : 100.104.10.1
```

11. Respond **yes** when asked to configure advanced IP options.

```
Configure advanced IP options? (yes/no) [n]: yes
```

12. Respond **yes** when asked to configure static route (this can be changed later).

```
Configure static route? (yes/no) [n]: yes
```

13. Enter the destination prefix and mask, and other values as prompted.

```
Destination prefix : 10.100.100.0
Destination prefix mask : 255.255.255.0
Next hop IPv4 address : 10.100.100.1
```

14. Configure the DNS IPv4 addresses.

```
Configure the DNS IPv4 address? (yes/no) [n]: yes
DNS IP address: 10.100.100.2
```

15. Skip configuring the default domain name (this will be configured later).

```
Configure the default domain name? (yes/no) [n]: no
```

16. Accept the default responses until asked to configure SSH and the NTP server.

```
Enable the telnet service? (yes/no) [n]: no
Enable the ssh service? (yes/no) [y]: yes
Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
Number of rsa key bits <1024-2048> [1024]: 1024

Configure the NTP server? (yes/no) [n]: yes
NTP server IPv4 address : 10.100.100.3
```

17. Accept the default responses until asked to specify the CoPP system profile. Enter **strict**.

```
Configure default interface layer (L3/L2) [L2]:
Configure default switchport interface state (shut/noshut)
[noshut]:
Configure CoPP system profile (strict/moderate/lenient/dense)
[strict]: strict
```

18. After reviewing the configuration, save the configuration.

```
The following configuration will be applied:
no password strength-check
switchname test123sw-rocea0
ip route 100.104.8.0 255.255.248.0 100.104.10.1
```

```
vrf context management
ip route 0.0.0.0/0 100.104.10.1
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
ntp server 100.104.10.1
system default switchport
no system default switchport shutdown
copp profile strict
interface mgmt0
ip address 100.104.10.21 255.255.248.0
no shutdown
```

Would you like to edit the configuration? (yes/no) [n]:

Use this configuration and save it? (yes/no) [y]: **yes**

```
[#####] 100%
Copy complete.
```

19. Enable the scp server feature on the switch.

```
test123sw-rocea0# feature scp-server
```

20. Save the running configuration to flash.

```
test123sw-rocea0# copy running-config startup-config
[#####] 100%
Copy complete.
```

21. Apply the golden configuration on the switch.

- Starting with Oracle Exadata System Software release 20.1.0, use the procedure described in Applying Golden Configuration Settings on RoCE Network Fabric Switches, in *Oracle Exadata Database Machine Maintenance Guide*.
- Otherwise, use the following procedure to apply the golden configuration on the switch:
 - a. Delete the configuration file on the switch for the target configuration.

 **Note:**

If you do not remove the file you are replacing, then when you attempt to overwrite the file you will get a 'permission denied' error.

Log in to the switch, enter configuration mode, then run a command similar to the following:

```
test123sw-rocea0# delete bootflash:roce_leaf_switch.cfg
Do you want to delete "/roce_leaf_switch.cfg" ? (yes/no/abort)
```

```
[y] y
test123sw-rocea0#
```

- b. Log in to a server that has SSH access to the switch, and contains the latest RDMA Network Fabric patch ZIP file.

If you do not have the patch ZIP file available, you can download the latest patch for your Oracle Exadata System Software release. Refer to My Oracle Support Doc ID 888828.1 and search for the term 'RDMA network switch' for information about the latest available patches.

- c. Unzip the RDMA Network Fabric patch ZIP file and change directories to the location of the `patchmgr` utility.
- d. Locate the golden configuration files in the RDMA Network Fabric patch bundle.

The files are located within the `roce_switch_templates` directory.

The golden configuration files are as follows:

- Single rack leaf: `roce_leaf_switch.cfg`
 - Multi-rack leaf: `roce_leaf_switch_multi.cfg`
 - Multi-rack spine: `roce_spine_switch_multi.cfg`
 - Single rack leaf with Secure Fabric support: `roce_sf_leaf_switch.cfg`
 - Multi-rack leaf with Secure Fabric support:
`roce_sf_leaf_switch_multi.cfg`
 - Single rack leaf configured with 23 host ports:
`roce_leaf_switch_23hosts.cfg`
 - Multi-rack leaf configured with 23 host ports:
`roce_leaf_switch_23hosts_multi.cfg`
- e. Copy the golden configuration file to the switch.

In the following example, `100.104.10.21` represents the IP address of the switch you are configuring.

```
# scp roce_leaf_switch.cfg admin@100.104.10.21:/
User Access Verification
Password:
roce_leaf_switch.cfg 100% 23KB 23.5KB/s 00:00
```

- f. Apply the golden configuration file on the switch.

Use the `run-script` command while connected directly to the switch.

```
test123sw-rocea0# run-script bootflash:roce_leaf_switch.cfg |
grep 'none'
```

 **Note:**

This command may take up to 1-2 minutes on a single-rack switch and up to 3-4 minutes on a multi-rack switch.

- g. Verify the switch configuration.

Use the `patchmgr` utility on the server that has SSH access to the switch, and contains the latest RDMA Network Fabric patch bundle.

In the following command, `roceswitch.lst` is a file that contains the switch host name or IP address.

```
# ./patchmgr --roceswitches roceswitch.lst --verify-config
```

- 22. Backup up the switch configuration.

Follow the steps in [Backing Up Settings on the ROCE Switch](#), in *Oracle Exadata Database Machine Maintenance Guide*.

- 23. Optional: Set the clock, using the same procedure as in [Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch](#).

6.2.3 Configuring the Cisco Nexus C9336C-FX2 Switches to Enable Exadata Secure RDMA Fabric Isolation

You can configure your RoCE Network Fabric switches to enable Exadata Secure RDMA Fabric Isolation.

Secure Fabric enables network separation between different clusters on systems with RoCE Network Fabric. This capability is conceptually similar to InfiniBand Network Fabric partitioning. Secure Fabric support is available starting with Oracle Exadata System Software release 20.1.0.

To use Secure Fabric, you must enable Secure Fabric support in the RoCE Network Fabric switch hardware before initial system deployment using Oracle Exadata Deployment Assistant (OEDA):

For each RoCE Network Fabric leaf switch, perform the basic configuration steps outlined in [Configuring the Cisco Nexus C9336C-FX2 Switch](#).

Then, on each RoCE Network Fabric leaf switch, you must apply a specific golden configuration to enable Secure Fabric support. Use the procedure described in [Applying Golden Configuration Settings on RoCE Network Fabric Switches](#), in *Oracle Exadata Database Machine Maintenance Guide*. However, when you specify the configuration type for each leaf switch, ensure that you specify the configuration type that enables Secure Fabric support.

After you complete the switch configuration, the leaf switch ports become trunk ports, which can carry network traffic with multiple VLAN IDs.

6.3 Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack

Oracle Exadata Database Machine X3-2 systems and Oracle Exadata Database Machine X2-2 systems have three Sun Datacenter InfiniBand Switch 36 switches. Starting with Oracle Exadata Database Machine X4-2, Oracle Exadata Database Machine Systems have two Sun Datacenter InfiniBand Switch 36 switches.

 **Note:**

This procedure does not apply to Oracle Exadata Database Machine X8M racks with RoCE Network Fabric.

The switch located in rack unit 1 (U1) is referred to as the spine switch. The other two switches are referred to as the leaf switches. The location of the leaf switches is as follows:

- Oracle Exadata Database Machine Two-Socket Systems (X3-2 and later): rack unit 20 (U20) and rack unit 22 (U22)
- Oracle Exadata Database Machine X2-2 racks: rack unit 20 (U20) and rack unit 24 (U24)
- Oracle Exadata Database Machine Eight-Socket Systems (X2-8 and later) Full Racks: Rack unit 21 (U21) and rack unit 23 (U23)

The spine switch is the Subnet Manager Master for the InfiniBand Network Fabric subnet. The Subnet Manager Master has priority 8, and can be verified using the following procedure:

1. Log in to the spine switch as the `root` user.
2. Run the `setsmpriority list` command.

The command should show that `smpriority` has a value of 8. If `smpriority` has a different value, then do the following:

- a. Use the `disablesm` command to stop the Subnet Manager.
- b. Use the `setsmpriority 8` command to set the priority to 8.
- c. Use the `enablesm` command to restart the Subnet Manager.

The leaf switches are the Standby Subnet Managers with a priority of 5. This can be verified using the preceding procedure, substituting a value of 5 in the `setsmpriority` command above.

 **Note:**

Oracle Exadata Database Machine Half Rack with Sun Fire X4170 Oracle Database Servers include two Sun Datacenter InfiniBand Switch 36 switches, which are set to priority 5.

To determine the Subnet Manager Master, log in as the `root` user on any InfiniBand Network Fabric switch, and run the `getmaster` command. The location of the Subnet Manager Master is displayed. The following is an example of the output from the `getmaster` command:

```
# getmaster
20100701 11:46:38 OpenSM Master on Switch : 0x0021283a8516a0a0 ports 36
Sun DCS 36
QDR switch dm01sw-ib1.example.com enhanced port 0 lid 1 lmc 0
```

The preceding output shows the proper configuration. The Subnet Master Manager is running on spine switch `dm01sw-ib1.example.com`.

If the spine switch is not the Subnet Manager Master, then do the following procedure to set the Subnet Manager Master:

1. Use the `getmaster` command to identify the current location of the Subnet Manager Master.
2. Log in as the `root` user on the leaf switch that is the Subnet Manager Master.
3. Disable Subnet Manager on the switch. The Subnet Manager Master relocates to another switch.
4. Use the `getmaster` command to identify the current location of the Subnet Manager Master. If the spine switch is not Subnet Manager Master, then repeat steps 2 and 3 until the spine switch is the Subnet Manager Master.
5. Enable Subnet Manager on the leaf switches that were disabled during this procedure.

 **Note:**

- If the InfiniBand Network Fabric network consists of four or more racks cabled together, then only the spine switches should run Subnet Manager. The leaf switches should have Subnet Manager disabled on them.
- Oracle Exadata Database Machine Half Racks with Sun Fire X4170 Oracle Database Servers, and Oracle Exadata Database Machine Quarter Racks have two Sun Datacenter InfiniBand Switch 36 switches, and both are set to priority 5. The master is the one with the lowest GUID.

 **See Also:**

- ["Enable the Subnet Manager"](#) in *Sun Datacenter InfiniBand Switch 36 User's Guide*
- ["Disable the Subnet Manager"](#) in *Sun Datacenter InfiniBand Switch 36 User's Guide*
- *Oracle Exadata Database Machine System Overview* for hardware component information
- Cabling tables in *Oracle Exadata Database Machine System Overview*

6.4 Configuring the Cisco Ethernet Switch

You must perform an initial configuration of the Cisco Ethernet switch.

- [Configuring the Cisco Catalyst 4948 Ethernet Switch](#)
The Cisco Catalyst 4948 Ethernet switch supplied with Oracle Exadata Rack is minimally configured during installation.
- [Configuring the Cisco Nexus 93108-1G or 9348 Ethernet Switch](#)
The Cisco Nexus 93108-1G or 9348 Ethernet switch supplied with Oracle Exadata Rack is minimally configured during installation.

6.4.1 Configuring the Cisco Catalyst 4948 Ethernet Switch

The Cisco Catalyst 4948 Ethernet switch supplied with Oracle Exadata Rack is minimally configured during installation.

The minimal configuration disables IP routing, and sets the following:

- Host name
- IP address setup
- Subnet mask
- Default gateway
- Domain name
- Name server
- NTP server
- Time
- Time zone

Before configuring the switch, note the following:

- The Cisco Ethernet switch should not be connected until the running configuration has been verified, and any necessary changes have been made by the network administrator.
- The Cisco Ethernet switch should not be connected to the customer network until the IP addresses on all components have been configured in Oracle Exadata Rack. This is to prevent any duplicate IP address conflicts which are possible due to the default addresses set in the components when shipped.

Note that the Cisco 4948E-F switch supports multiple uplinks to the customer network by utilizing ports 49 - 52. This is a more complicated switch setup due to the redundant connectivity, and should be performed by the customer's network administrator.

The following procedure describes how to configure the Cisco Ethernet switch. Configuration should be done with the network administrator.

1. Connect a serial cable from the Cisco switch console to a laptop or similar device. An Oracle supplied rollover cable is pre-installed on the Cisco serial console port. Obtain the appropriate adapter and connect it at the end of the rollover cable. An Oracle P/N 530-3100 RJ45-DB9 adapter as used on ILOM ports will also work, connected at the end of the network cable.
2. Ensure the terminal session is recorded on the laptop by logging the output. The output can be used as a reference that the switch has been configured correctly.

The default serial port speed is 9600 baud, 8 bits, no parity, 1 stop bit, and no handshake.

```
Switch con0 is now available
Press RETURN to get started.
```

3. Change to the enable mode.

```
Switch> enable
Password: *****
Switch#
```

 **Note:**

If you do not have the password, then contact Oracle Support Services.

4. Check the current version on the switch.

```
Switch# show version
Cisco IOS Software, Catalyst 4500 L3 Switch Software (cat4500e-
IPBASEK9-M), Version 15.2(3)E2, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2014 by Cisco Systems, Inc.
Compiled Tue 11-Mar-14 18:28 by prod_rel_team

ROM: 12.2(44r)SG12
zdlralsw-ip uptime is 1 minute
System returned to ROM by reload
System image file is "bootflash:cat4500e-ipbasek9-mz.152-3.E2.bin"
Hobgoblin Revision 22, Fortooin Revision 1.40
...

Configuration register is 0x2102

Switch#
```

The version of the Cisco 4948E-F switch firmware purchased and shipped by Oracle with Oracle Exadata Database Machine X6 is IPBASEK9-MZ, which includes telnet and ssh support. Currently the full release version string is cat4500e-ipbasek9-mz.152-3.E2.bin.

5. Configure the network for a single VLAN. The following example assumes you are using IPv4 addressing.

```
Switch# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)# interface vlan 1
Switch(config-if)# ip address 10.7.7.34 255.255.255.0
Switch(config-if)# end
Switch# *Sep 15 14:12:06.309:%SYS-5-CONFIG_I:Configured from
console by console
Switch# write memory
```

```
Building configuration...
Compressed configuration from 2474 bytes to 1066 bytes [OK ]
```

6. If IP routing is not used on the switch, this step is required. Disable the default IP routing setting, and configure the default gateway.

```
Switch#configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)#no ip routing
Switch(config)#ip default-gateway 10.7.7.1
Switch(config)#end
*Sep 15 14:12:46.309:%SYS-5-CONFIG_I:Configured from console by
console
Switch#write memory
Building configuration...
Compressed configuration from 2492 bytes to 1070 bytes [OK ]
```

7. If IP routing is required on the switch, then leave the IP routing setting as the default, and configure the default gateway. Replace 10.7.7.1 with the IP address of the gateway for the installation:

```
Switch#configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)#ip route 0.0.0.0 0.0.0.0 10.7.7.1
Switch(config)#end
*Sep 15 14:13:26.013:%SYS-5-CONFIG_I:Configured from console by
console
Switch#write memory
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes [OK ]
```

8. Set the host name of the switch.

This example sets the name to exa1sw-ip:

```
Switch#configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)#hostname exa1sw-ip
exa1sw-ip(config)#end
exa1sw-ip#write memory
Building configuration...
Compressed configuration from 3789 bytes to 1469 bytes [OK ]
```

The system host name is used as the prompt name.

9. Configure up to three DNS servers. Replace the domain name and IP addresses used in this example with the values for the installation:

```
exa1sw-ip#configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
exa1sw-ip(config)#ip domain-name example.com
exa1sw-ip(config)#ip name-server 10.7.7.3
exa1sw-ip(config)#ip name-server 198.51.100.5
exa1sw-ip(config)#ip name-server 10.8.160.1
exa1sw-ip(config)#end
```

```
*Sep 15 14:26:37.045:%SYS-5-CONFIG_I:Configured from console by
console
exalsw-ip#write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes [OK ]
```

If you do not have DNS service available, you must still set the domain-name so that you can configure the SSH keys.

10. (Optional) Set the password.

```
exalsw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
exalsw-ip(config)# enable password password
exalsw-ip(config)# enable secret password
exalsw-ip(config)# end
exalsw-ip# write memory
*Sep 15 14:25:05.893:%SYS-5-CONFIG_I:Configured from console by
console
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes [OK ]
```

11. Verify telnet access is disabled. Telnet is not secure, and should not be enabled unless there is a compelling reason. To enable telnet, set a password. To disable it, remove the password.

```
exalsw-ip#configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
exalsw-ip(config)#line vty 0 15
exalsw-ip(config)#login
% Login disabled on line 1, until 'password' is set
% Login disabled on line 2, until 'password' is set
...
% Login disabled on line 16, until 'password' is set
exalsw-ip(config)#end
```

If the login command returns output as shown above, then telnet access has been disabled. If instead you get a prompt, then telnet access is not yet disabled so should be disabled now.

```
exalsw-ip(config-line)#no password
exalsw-ip(config-line)#end
exalsw-ip#write memory
Building configuration...
Compressed configuration from 3786 bytes to 1468 bytes [OK ]
```

12. To configure a secure shell (SSH) on the Ethernet switch:

```
exalsw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
exalsw-ip(config)# crypto key generate rsa
% You already have RSA keys defined named exalsw-ip.example.com.
% Do you really want to replace them? [yes/no]: yes
```

Choose the size of the key modulus in the range of 360 to 2048 for your General Purpose Keys. Choosing a key modulus greater than 512 may take a few minutes.

```
How many bits in the modulus [512]: 768
```

```
% Generating 768 bit RSA keys, keys will be non-exportable...[OK]
```

```
exalsw-ip(config)# username admin password 0 welcome1
```

```
exalsw-ip(config)# line vty 0 15
```

```
exalsw-ip(config-line)# transport input ssh
```

```
exalsw-ip(config-line)# exit
```

```
exalsw-ip(config)# aaa new-model
```

```
exalsw-ip(config)# ip ssh time-out 60
```

```
exalsw-ip(config)# ip ssh authentication-retries 3
```

```
exalsw-ip(config)# ip ssh version 2
```

```
exalsw-ip(config)# end
```

```
*Sep 15 14:26:37.045: %SYS-5-CONFIG_I: Configured from console by console
```

```
exalsw-ip# write memory
```

```
Building configuration...
```

```
Compressed configuration from 2603 bytes to 1158 bytes[OK]
```

13. Set the clock and time zone. The switch keeps internal time in Coordinated Universal Time (UTC) format.

- To use UTC, use the following command:

```
no clock timezone global configuration
```

- To use a time zone, use the following command:

```
clock timezone zone hours-offset [minutes-offset]
```

In the preceding command, *zone* is the time zone to display when standard time is in effect, *hours-offset* is the hours offset from UTC, and *minutes-offset* is the minutes offset from UTC.

- Daylight savings time (or summer time) is disabled by default. To set summer time hours, use the following command:

```
clock summer-time zone recurring [week day month hh:mm week day month \  
hh:mm [offset]]
```

In the preceding command, *zone* is the time zone to be displayed when summer time is in effect (EDT, for example), *week* is the week of the month (1 to 5 or last), *day* is the day of the week (Sunday, Monday, ...), *month* is the month (January, February, ...), *hh:mm* is the hours and minutes in 24-hour format, and *offset* is the number of minutes to add during summer time. The default offset is 60 minutes.

- To manually set the clock to any time use the following command, where the time specified is relative to the configured time zone:

```
clock set hh:mm:ss month day year
```

In the preceding command, *hh:mm:ss* is the time in 24-hour format, *day* is the day by date in the month, *month* is the name of the month, and *year* is the 4-digit year.

The ordering of commands is important when setting the local time and time zone. For example, to set the local time to US Eastern time:

```
exalsw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
exalsw-ip(config)# clock timezone EST -5
exalsw-ip(config)# clock summer-time EDT recurring
exalsw-ip(config)# end
exalsw-ip# clock set 21:00:00 August 09 2018
exalsw-ip# write memory
Building configuration...
Compressed configuration from 3784 bytes to 1465 bytes [OK ]
exalsw-ip# show clock
21:00:06.643 EST Mon Aug 9 2018
```

14. After setting the local time zone, you can configure up to two NTP servers. Replace the IP addresses used in this example with the values for the installation:

```
exalsw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
exalsw-ip(config)# ntp server 10.7.7.32 prefer
exalsw-ip(config)# ntp server 198.51.100.19
exalsw-ip(config)# end
*Sep 15 14:51:08.665:%SYS-5-CONFIG_I:Configured from console by
console
exalsw-ip# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes [OK ]
exalsw-ip# show ntp status
<output will vary per network>
.
exalsw-ip# show clock
21:00:23.175 EST Mon Aug 9 2018
```

The NTP server is synchronized to local time when you connect the Cisco switch to the network and it has access to NTP.

Symbols that precede the `show clock` display indicate that the time is the following:

- * Not authoritative
- . Authoritative, but NTP is not synchronized.
- Authoritative (blank space).

15. Verify the Ethernet configuration using the following command:

```
exalsw-ip# show running-config
Building configuration...
Current configuration : 3923 bytes
!
version 15.2
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
.
.
.
```

 **Note:**

If any setting is incorrect, then repeat the appropriate step. To erase a setting, enter `no` in front of the same command. For example, to erase the default gateway, use the following commands:

```
exalsw-ip#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
exalsw-ip(config)# no ip default-gateway 10.7.7.1
exalsw-ip(config)# end
exalsw-ip#
*Sep 15 14:13:26.013: %SYS-5-CONFIG_I: Configured from
console by console
exalsw-ip(config)# write memory
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes[OK]
```

16. Save the current configuration.

```
exalsw-ip#copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
Compressed configuration from 2654 bytes to 1189 bytes[OK]
```

17. Exit from the session using the following command:

```
exalsw-ip# exit

exalsw-ip con0 is now available

Press RETURN to get started.
```

18. Disconnect the cable from the Cisco console.

The Cisco switch must not be connected to the management network at this stage. The switch will be connected later after Oracle has configured the systems with the necessary IP addresses and you have worked with the field service engineer to make any additional changes necessary for connecting to the network.

19. To check the Cisco switch, attach a laptop computer to port 48, and ping the IP address of the internal management network to check the configuration.

Related Topics

- [Network Connection and IP Address Requirements for Oracle Exadata Database Machine](#)
- [Network Configuration Checklist](#)
Use this checklist to ensure that the network configuration requirements have been addressed.

6.4.2 Configuring the Cisco Nexus 93108-1G or 9348 Ethernet Switch

The Cisco Nexus 93108-1G or 9348 Ethernet switch supplied with Oracle Exadata Rack is minimally configured during installation.

Note that the Cisco Nexus 93108-1G or 9348 switch supports multiple uplinks to the customer network by utilizing the QSFP+ ports. This is a more complicated switch setup due to the redundant connectivity, and should be performed by the customer's network administrator.

Whether you are configuring the switch for the first time, or configuring a replacement switch, use the following procedures:

- [Performing the Initial Switch Configuration for the Cisco Nexus 93108-1G or 9348 Ethernet Switch](#)
During the initial configuration, you reset the switch and use the Basic System Configuration Dialog to configure the switch.
- [Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch](#)
After you have performed the initial configuration, you can adjust the time used by the switch.

6.4.2.1 Performing the Initial Switch Configuration for the Cisco Nexus 93108-1G or 9348 Ethernet Switch

During the initial configuration, you reset the switch and use the Basic System Configuration Dialog to configure the switch.

Before configuring the switch, note the following:

- The Cisco Ethernet switch should not be connected until the running configuration has been verified, and any necessary changes have been made by the network administrator.
- The Cisco Ethernet switch should not be connected to the customer network until the IP addresses on all components have been configured in Oracle Exadata Rack. This is to prevent any duplicate IP address conflicts which are possible due to the default addresses set in the components when shipped.

Configuration should be done with the network administrator.

1. Connect from the Cisco switch serial console to a laptop or similar device using the available RJ45 cable.
2. Ensure the terminal session is recorded on the laptop by logging the output.
The output can be used as a reference that the switch has been configured correctly. The default serial port speed is 9600 baud, 8 bits, no parity, 1 stop bit, and no handshake.
3. Power on the switch.
4. Log in as the admin user.

```
User Access Verification
exadatax7-adm0 login: admin
Password: *****
```

 **Note:**

If you do not have the password for the admin user, then contact Oracle Support Services.

5. Erase the existing configuration.

```
exadatax7-adm0# write erase
```

```
Warning: This command will erase the startup-configuration.
```

```
Do you wish to proceed anyway? (y/n) [n] y
```

6. Restart the system so you can perform the automated setup.

```
exadatax7-adm0# reload
```

```
This command will reboot the system. (y/n)? [n] y
```

```
2017 Aug 31 01:09:00 exadatax7-adm0 %$ VDC-1 %$ %PLATFORM-2-
PFM_SYSTEM_RESET: Manual system restart from Command Line Interface
```

```
CISCO SWITCH Ver7.59
Device detected on 0:1:2 after 0 msec
...
```

7. Switch to normal setup and, when asked if you want to enforce secure password standard, enter no, then enter a new password for the admin user.

```
Running S93thirdparty-script...
```

```
Populating conf files for hybrid sysmgr ...
Starting hybrid sysmgr ...
inserting /isan/lib/modules/klm_cisco_nb.o ... done
```

```
Abort Auto Provisioning and continue with normal setup ? (yes/no)
```

[n]: **yes**

---- System Admin Account Setup ----

Do you want to enforce secure password standard (yes/no) [y]: **no**

Enter the password for "admin":

Confirm the password for "admin":

8. When the Basic System Configuration Dialog appears, choose to enter the basic configuration dialog.

---- Basic System Configuration Dialog VDC: 1 ----

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Please register Cisco Nexus9000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus9000 devices must be registered to receive entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

9. In the basic configuration, you can use the default inputs until asked to enter the switch name.

In this example, the switch has a name of test123sw-adm0.

Create another login account (yes/no) [n]:

Configure read-only SNMP community string (yes/no) [n]:

Configure read-write SNMP community string (yes/no) [n]:

Enter the switch name : **test123sw-adm0**

10. Respond no when asked to configure Out-of-band management configuration.

Continue with Out-of-band (mgmt0) management configuration?
(yes/no) [y]: **no**

11. Respond yes when asked to configure advanced IP options.

Configure advanced IP options? (yes/no) [n]: **yes**

12. Respond no when asked to configure static route (this will be configured later).

Configure static route? (yes/no) [n]: **no**

13. Enter the destination prefix and mask, and other values as prompted.

```
Destination prefix : 10.100.100.0
Destination prefix mask : 255.255.255.0
Next hop IPv4 address : 10.100.100.1
```

14. Skip configuring the DNS IPv4 addresses (this will be configured later).

```
Configure the DNS IPv4 address? (yes/no) [n]: no
```

15. Skip configuring the default domain name (this will be configured later).

```
Configure the default domain name? (yes/no) [n]: no
```

16. Accept the default responses until asked to configure SSH and the NTP server.

```
Enable the telnet service? (yes/no) [n]: no
Enable the ssh service? (yes/no) [y]: yes
Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
Number of rsa key bits <1024-2048> [1024]: 1024
```

```
Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : 10.100.100.3
```

17. Accept the default responses until asked to specify the CoPP system profile. Enter lenient.

```
Configure default interface layer (L3/L2) [L2]:
Configure default switchport interface state (shut/noshut)
[noshut]:
Configure CoPP system profile (strict/moderate/lenient/dense)
[strict]: lenient
```

18. After reviewing the configuration, save the configuration.

```
The following configuration will be applied:
no password strength-check
switchname test123sw-adm0
...

Would you like to edit the configuration? (yes/no) [n]:

Use this configuration and save it? (yes/no) [y]: yes

[#####] 100%
Copy complete.
```

19. Add the VLAN 1 IP address.

```
test123sw-adm0(config)# feature interface-vlan
test123sw-adm0(config)# interface vlan 1
test123sw-adm0(config-if)# ip address 10.100.100.110/24
```

```
test123sw-adm0(config-if)# no shutdown
test123sw-adm0(config-if)# exit
```

20. Set the spanning tree port type for ports 1-47.

```
test123sw-adm0(config)# interface E1/1-47
test123sw-adm0(config-if)# spanning-tree port type edge
test123sw-adm0(config-if)# exit
```

21. Set switchport on all 48 ports and set port 48 to a network port (instead of a host port).

```
test123sw-adm0(config)# interface E1/1-48
test123sw-adm0(config-if)# switchport
test123sw-adm0(config-if)# exit
test123sw-adm0(config)# interface E1/48
test123sw-adm0(config-if)# spanning-tree port type network
test123sw-adm0(config-if)# ip route 0.0.0.0/0 10.100.100.1
```

22. Configure the DNS information.

```
test123sw-adm0(config)# ip domain-name example.com
test123sw-adm0(config)# ip name-server 10.100.100.2
test123sw-adm0(config)# exit
```

23. Save the current configuration.

```
test123sw-adm0# copy running-config startup-config
[#####] 100%
Copy complete.
```

24. Optional: Set the clock, as described in the next topic.

6.4.2.2 Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch

After you have performed the initial configuration, you can adjust the time used by the switch.

1. Log in as the admin user.
2. View the current time.

```
test123sw-adm0(config)# show clock
20:44:52.986 UTC Thu Aug 31 2017
Time source is NTP
```

3. Set the timezone appropriately.

```
test123sw-adm0(config)# clock timezone PST -8 0
```

- View the modified time.

```
test123sw-adm0(config)# show clock
12:46:22.692 PST Thu Aug 31 2017
Time source is NTP
```

- Save the configuration.

```
test123sw-adm0# copy running-config startup-config
[#####] 100%
Copy complete.
```

6.5 Configuring the Power Distribution Units

The power distribution units (PDUs) are configured with a static IP address to connect to the network for monitoring.

Ensure you have the following before connecting the PDU to the network:

- Static IP address
- Subnet mask
- Default gateway
- Ethernet cables
- Laptop computer with either Oracle Solaris or Microsoft Windows

For systems less than full racks, the PDU Ethernet connections may use available ports in the Cisco Management Switch. Note that for full rack systems, you cannot connect PDUs to the Cisco Management Switch because all the ports on the switch are already reserved for other components.

The following procedure describes how to configure and connect the PDU to the network:

- Power off the PDU.
- Disconnect the power leads from the power source.
- Select an unused LAN network connection from the available network connections on the laptop as follows:
 - For Oracle Solaris:
 - Log in as a super user.
 - Use the `dladm show-link` command to find an unused Ethernet interface.
 - Use the `ifconfig -a` command to determine which interface is being used.
 - Use the following command to plumb the unused interface:

```
# ifconfig interface plumb up
```

In the preceding command, *interface* is the network interface determined in step 3.c.

- e. Using the following command to assign an IPv4 address and netmask to the interface:

```
# ifconfig interface IPv4_address netmask + netmask \  
broadcast broadcast up
```

In the preceding command, *interface* is the network interface, *IPv4_address* is the IP address, *netmask* is the netmask address, and *broadcast* is the broadcast address.

 **Note:**

The network interfaces configured with the `ifconfig` command do not persist across system restarts.

- For Microsoft Windows:
 - a. Select **Settings** from the Start menu.
 - b. Select **Network Connections** from the Settings menu.
 - c. Select an unused local area connection from the Network Connections menu.
 - d. Double-click **Internet Protocol (TCP/IP)** in the Local Area Connection Properties window. The Internet Protocol (TCP/IP) window appears.
 - e. Select **Use the following IP address**.
 - f. Enter the IP address in the IP address field. Do not use 192.168.0.1 because that is the default PDU metering unit address.
 - g. Enter the IP address in the Subnet mask field.
 - h. Leave the Default gateway field blank.
 - i. Click **OK**.
- 4. Connect the PDU and the laptop to two ports on the Cisco Management Switch. The equipment currently using the Cisco Management Switch must be disconnected temporarily.
- 5. Connect the PDU power input lead to the power source. Only one PDU power input lead needs to be connected, as follows:
 - On PDUs with one power input lead, connect the one lead to the power source.
 - On PDUs with two power leads, connect the first power lead to the power source. The first power lead is labeled 0.
 - On PDUs with three power input leads, connect the middle power lead to the power source. The middle power lead is labeled 1.
- 6. Use a Web browser to access the PDU metering unit by entering the factory default IP address for the unit, 192.168.0.1, in the address line of the browser. The Current Measurement page should appear.
- 7. Click **Network Configuration** in the upper left of the page.

8. Log in as the `admin` user on the PDU metering unit. Change the password after configuring the network.
9. Confirm the DHCP Enable check box is not selected.
10. Enter the following network settings for the PDU metering unit:
 - IP address
 - Subnet mask address
 - Default gateway
11. Click **Submit** to set the network settings, and reset the PDU metering unit.
12. Disconnect the PDU and laptop from the Cisco Management Switch.
13. Reconnect the two cables that were originally connected to the Cisco Management Switch.
14. Connect an Ethernet cable to the PDU metering unit RJ-45 Ethernet port and to the network.
15. Log in to the PDU metering unit using a Web browser. Use the PDU metering unit's static IP address in the browser's address line. If the configuration was successful, then the Current Measurement page is displayed.
16. Connect the remaining PDU power input leads to the power source.

6.6 Configuring the Threshold Settings for the Power Distribution Units

The PDU current can be monitored directly or with Oracle Enterprise Manager Grid Control. Configure the threshold settings to monitor the PDUs. The threshold settings depend on the size of Oracle Exadata Rack and type PDU. The configurable threshold values for each metering unit module and phase are *Info low*, *Pre Warning*, and *Alarm*.

Note:

The thresholds must be set in the Param Configuration section of the PDU metering unit. Setting threshold values using the Oracle Enterprise Manager Grid Control plug-in has no effect.

- [PDU Thresholds for Oracle Exadata Database Machine X5-2 and Later](#)
- [PDU Thresholds for Oracle Exadata Database Machine X4-2](#)
- [PDU Thresholds for Oracle Exadata Database Machine X3-2](#)
- [PDU Thresholds for Oracle Exadata Database Machine X2-2 \(with X4170 M2 and X4270 M2 servers\)](#)
- [PDU Thresholds for Oracle Exadata Database Machine X2-2 \(with X4170 and X4275 servers\)](#)

- [PDU Thresholds for Oracle Exadata Database Machine X4-8 and Later](#)
This section contains the PDU thresholds for Oracle Exadata Database Machine Eight-Socket systems for X4-8 and later.
- [PDU Thresholds for Oracle Exadata Database Machine X3-8 Full Rack](#)
- [PDU Thresholds for Oracle Exadata Database Machine X2-8 Full Rack](#)
- [PDU Thresholds for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers](#)

See Also:

- *Oracle Enterprise Manager Exadata Management Getting Started Guide* for information about using Oracle Enterprise Manager Grid Control to collect PDU metrics and events
- *Sun Rack II Power Distribution Units User's Guide* at <https://docs.oracle.com/cd/E19657-01/html/E23956/index.html> for information about configuring and monitoring PDUs
- *Oracle Rack Cabinet 1242 Power Distribution Units User's Guide* at http://docs.oracle.com/cd/E85660_01/html/E87281/index.html for X7 and later systems
- *Oracle Exadata Database Machine Maintenance Guide* for information about monitoring components

6.6.1 PDU Thresholds for Oracle Exadata Database Machine X5-2 and Later

Starting with Oracle Exadata Database Machine X5-2, Oracle Exadata Configuration Assistant (OECA) is the only source for PDU thresholds. OECA is available on Oracle Technology Network at <http://www.oracle.com/technetwork/database/exadata/oeca-download-2817713.html>

If for some reason you are unable to utilize OECA, contact Oracle Support.

6.6.2 PDU Thresholds for Oracle Exadata Database Machine X4-2

This section contains the PDU thresholds for Oracle Exadata Database Machine X4-2.

Note:

The values listed here are reliable for use on systems in the configuration received upon deployment and delivery from the Oracle Factory. Any changes to this configuration should be modeled in Oracle Exadata Configuration Assistant (OECA), and the PDU thresholds listed in OECA should then be used. OECA is available on Oracle Technology Network.

- [PDU Thresholds for Oracle Exadata Database Machine X4-2 Full Rack](#)

- [PDU Thresholds for Oracle Exadata Database Machine X4-2 Half Rack](#)
- [PDU Thresholds for Oracle Exadata Database Machine X4-2 Quarter Rack](#)

6.6.2.1 PDU Thresholds for Oracle Exadata Database Machine X4-2 Full Rack

The following tables list the threshold values for Oracle Exadata Database Machine X4-2 Full Rack:

- [Table 6-1](#) Threshold Values for Oracle Exadata Database Machine X4-2 Full Rack with Single-phase, Low-voltage PDU
- [Table 6-2](#) Threshold Values for Oracle Exadata Database Machine X4-2 Full Rack with Three-phase, Low-voltage PDU
- [Table 6-3](#) Threshold Values for Oracle Exadata Database Machine X4-2 Full Rack with Single-phase, High-voltage PDU
- [Table 6-4](#) Threshold Values for Oracle Exadata Database Machine X4-2 Full Rack with Three-phase, High-voltage PDU

[Table 6-1](#) lists the threshold values for Oracle Exadata Database Machine X4-2 Full Rack using a single-phase, low-voltage PDU.

Table 6-1 Threshold Values for Oracle Exadata Database Machine X4-2 Full Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	20	24
A	Module 1, phase 2	0	17	22
A	Module 1, phase 3	0	19	24
B	Module 1, phase 1	0	19	24
B	Module 1, phase 2	0	17	22
B	Module 1, phase 3	0	20	24

[Table 6-2](#) lists the threshold values for Oracle Exadata Database Machine X4-2 Full Rack using a three-phase, low-voltage PDU.

Table 6-2 Threshold Values for Oracle Exadata Database Machine X4-2 Full Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	32	40
A and B	Module 1, phase 2	0	34	42

Table 6-2 (Cont.) Threshold Values for Oracle Exadata Database Machine X4-2 Full Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 3	0	32	40

Table 6-3 lists the threshold values for Oracle Exadata Database Machine X4-2 Full Rack using a single-phase, high-voltage PDU.

Table 6-3 Threshold Values for Oracle Exadata Database Machine X4-2 Full Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	18	23
A	Module 1, phase 2	0	16	20
A	Module 1, phase 3	0	18	23
B	Module 1, phase 1	0	18	23
B	Module 1, phase 2	0	16	20
B	Module 1, phase 3	0	18	23

Table 6-4 lists the threshold values for Oracle Exadata Database Machine X4-2 Full Rack using a three-phase, high-voltage PDU.

Table 6-4 Threshold Values for Oracle Exadata Database Machine X4-2 Full Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	18	21
A and B	Module 1, phase 2	0	18	21
A and B	Module 1, phase 3	0	16	21

6.6.2.2 PDU Thresholds for Oracle Exadata Database Machine X4-2 Half Rack

The following tables list the threshold values for Oracle Exadata Database Machine X4-2 Half Rack:

- [Table 6-5](#) Threshold Values for Oracle Exadata Database Machine X4-2 Half Rack with Single-phase, Low-voltage PDU
- [Table 6-6](#) Threshold Values for Oracle Exadata Database Machine X4-2 Half Rack with Three-phase, Low-voltage PDU
- [Table 6-7](#) Threshold Values for Oracle Exadata Database Machine X4-2 Half Rack with Single-phase, High-voltage PDU
- [Table 6-8](#) Threshold Values for Oracle Exadata Database Machine X4-2 Half Rack with Three-phase, High-voltage PDU

[Table 6-5](#) lists the threshold values for Oracle Exadata Database Machine X4-2 Half Rack using a single-phase, low-voltage PDU.

Table 6-5 Threshold Values for Oracle Exadata Database Machine X4-2 Half Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	20	24
A	Module 1, phase 2	0	10	13
A	Module 1, phase 3	0	0.1	0.2
B	Module 1, phase 1	0	0.1	0.2
B	Module 1, phase 2	0	10	13
B	Module 1, phase 3	0	20	24

[Table 6-6](#) lists the threshold values for Oracle Exadata Database Machine X4-2 Half Rack using a three-phase, low-voltage PDU.

Table 6-6 Threshold Values for Oracle Exadata Database Machine X4-2 Half Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	17	22
A and B	Module 1, phase 2	0	19	24
A and B	Module 1, phase 3	0	15	20

[Table 6-7](#) lists the threshold values for Oracle Exadata Database Machine X4-2 Half Rack using a single-phase, high-voltage PDU.

Table 6-7 Threshold Values for Oracle Exadata Database Machine X4-2 Half Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	18	23
A	Module 1, phase 2	0	9	12
A	Module 1, phase 3	0	0.1	0.2
B	Module 1, phase 1	0	0.1	0.2
B	Module 1, phase 2	0	9	12
B	Module 1, phase 3	0	18	23

Table 6-8 lists the threshold values for Oracle Exadata Database Machine X4-2 Half Rack using a three-phase, high-voltage PDU.

Table 6-8 Threshold Values for Oracle Exadata Database Machine X4-2 Half Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	11	15
A and B	Module 1, phase 2	0	9	12
A and B	Module 1, phase 3	0	7	10

6.6.2.3 PDU Thresholds for Oracle Exadata Database Machine X4-2 Quarter Rack

The following tables list the threshold values for Oracle Exadata Database Machine X4-2 Quarter Rack:

- [Table 6-9](#) Threshold Values for Oracle Exadata Database Machine X4-2 Quarter Rack with Single-phase, Low-voltage PDU
- [Table 6-10](#) Threshold Values for Oracle Exadata Database Machine X4-2 Quarter Rack with Three-phase, Low-voltage PDU
- [Table 6-11](#) Threshold Values for Oracle Exadata Database Machine X4-2 Quarter Rack with Single-phase, High-voltage PDU
- [Table 6-12](#) Threshold Values for Oracle Exadata Database Machine X4-2 Quarter Rack with Three-phase, High-voltage PDU

Table 6-9 lists the threshold values for Oracle Exadata Database Machine X4-2 Quarter Rack using a single-phase, low-voltage PDU.

Table 6-9 Threshold Values for Oracle Exadata Database Machine X4-2 Quarter Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	10	13
A	Module 1, phase 2	0	5	7
A	Module 1, phase 3	0	0.1	0.2
B	Module 1, phase 1	0	0.1	0.2
B	Module 1, phase 2	0	5	7
B	Module 1, phase 3	0	10	13

Table 6-10 lists the threshold values for Oracle Exadata Database Machine X4-2 Quarter Rack using a three-phase, low-voltage PDU.

Table 6-10 Threshold Values for Oracle Exadata Database Machine X4-2 Quarter Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	11	14
A and B	Module 1, phase 2	0	10	13
A and B	Module 1, phase 3	0	5	7

Table 6-11 lists the threshold values for Oracle Exadata Database Machine X4-2 Quarter Rack using a single-phase, high-voltage PDU.

Table 6-11 Threshold Values for Oracle Exadata Database Machine X4-2 Quarter Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	9	12
A	Module 1, phase 2	0	5	7
A	Module 1, phase 3	0	0.1	0.2

Table 6-11 (Cont.) Threshold Values for Oracle Exadata Database Machine X4-2 Quarter Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
B	Module 1, phase 1	0	0.1	0.2
B	Module 1, phase 2	0	5	7
B	Module 1, phase 3	0	9	12

Table 6-12 lists the threshold values for Oracle Exadata Database Machine X4-2 Quarter Rack using a three-phase, high-voltage PDU.

Table 6-12 Threshold Values for Oracle Exadata Database Machine X4-2 Quarter Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	9	11
A and B	Module 1, phase 2	0	3	4
A and B	Module 1, phase 3	0	3	4

6.6.3 PDU Thresholds for Oracle Exadata Database Machine X3-2

This section contains the PDU thresholds for Oracle Exadata Database Machine X3-2.

 **Note:**

The values listed here are reliable for use on systems in the configuration received upon deployment and delivery from the Oracle Factory. Any changes to this configuration should be modeled in Oracle Exadata Configuration Assistant (OECA), and the PDU thresholds listed in OECA should then be used. OECA is available on Oracle Technology Network.

- [PDU Thresholds for Oracle Exadata Database Machine X3-2 Full Rack](#)
- [PDU Thresholds for Oracle Exadata Database Machine X3-2 Half Rack](#)
- [PDU Thresholds for Oracle Exadata Database Machine X3-2 Quarter Rack](#)

6.6.3.1 PDU Thresholds for Oracle Exadata Database Machine X3-2 Full Rack

The following tables list the threshold values for Oracle Exadata Database Machine X3-2 Full Rack:

- [Table 6-13](#) Threshold Values for Oracle Exadata Database Machine X3-2 Full Rack with Single-phase, Low-voltage PDU
- [Table 6-14](#) Threshold Values for Oracle Exadata Database Machine X3-2 Full Rack with Three-phase, Low-voltage PDU
- [Table 6-15](#) Threshold Values for Oracle Exadata Database Machine X3-2 Full Rack with Single-phase, High-voltage PDU
- [Table 6-16](#)
Threshold Values for Oracle Exadata Database Machine X3-2 Full Rack with Three-phase, High-voltage PDU

[Table 6-13](#) lists the threshold values for Oracle Exadata Database Machine X3-2 Full Rack using a single-phase, low-voltage PDU.

Table 6-13 Threshold Values for Oracle Exadata Database Machine X3-2 Full Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	20	24
A	Module 1, phase 2	0	14	18
A	Module 1, phase 3	0	19	24
B	Module 1, phase 1	0	19	24
B	Module 1, phase 2	0	14	18
B	Module 1, phase 3	0	20	24

[Table 6-14](#) lists the threshold values for Oracle Exadata Database Machine X3-2 Full Rack using a three-phase, low-voltage PDU.

Table 6-14 Threshold Values for Oracle Exadata Database Machine X3-2 Full Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	30	38
A and B	Module 1, phase 2	0	32	40
A and B	Module 1, phase 3	0	30	38

[Table 6-15](#) lists the threshold values for Oracle Exadata Database Machine X3-2 Full Rack using a single-phase, high-voltage PDU.

Table 6-15 Threshold Values for Oracle Exadata Database Machine X3-2 Full Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	18	23
A	Module 1, phase 2	0	13	17
A	Module 1, phase 3	0	18	22
B	Module 1, phase 1	0	18	22
B	Module 1, phase 2	0	13	17
B	Module 1, phase 3	0	18	23

[Table 6-16](#) lists the threshold values for Oracle Exadata Database Machine X3-2 Full Rack using a three-phase, high-voltage PDU.

Table 6-16 Threshold Values for Oracle Exadata Database Machine X3-2 Full Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	17	21
A and B	Module 1, phase 2	0	17	21
A and B	Module 1, phase 3	0	15	19

6.6.3.2 PDU Thresholds for Oracle Exadata Database Machine X3-2 Half Rack

The following tables list the threshold values for Oracle Exadata Database Machine X3-2 Half Rack:

- [Table 6-17](#) Threshold Values for Oracle Exadata Database Machine X3-2 Half Rack with Single-phase, Low-voltage PDU
- [Table 6-18](#) Threshold Values for Oracle Exadata Database Machine X3-2 Half Rack with Three-phase, Low-voltage PDU
- [Table 6-19](#) Threshold Values for Oracle Exadata Database Machine X3-2 Half Rack with Single-phase, High-voltage PDU
- [Table 6-20](#) Threshold Values for Oracle Exadata Database Machine X3-2 Half Rack with Three-phase, High-voltage PDU

[Table 6-17](#) lists the threshold values for Oracle Exadata Database Machine X3-2 Half Rack using a single-phase, low-voltage PDU.

Table 6-17 Threshold Values for Oracle Exadata Database Machine X3-2 Half Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	20	24
A	Module 1, phase 2	0	9	11
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	9	11
B	Module 1, phase 3	0	20	24

Table 6-18 lists the threshold values for Oracle Exadata Database Machine X3-2 Half Rack using a three-phase, low-voltage PDU.

Table 6-18 Threshold Values for Oracle Exadata Database Machine X3-2 Half Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	16	21
A and B	Module 1, phase 2	0	19	24
A and B	Module 1, phase 3	0	14	18

Table 6-19 lists the threshold values for Oracle Exadata Database Machine X3-2 Half Rack using a single-phase, high-voltage PDU.

Table 6-19 Threshold Values for Oracle Exadata Database Machine X3-2 Half Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	18	23
A	Module 1, phase 2	0	8	10
A	Module 1, phase 3	0	0	0

Table 6-19 (Cont.) Threshold Values for Oracle Exadata Database Machine X3-2 Half Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	8	10
B	Module 1, phase 3	0	18	23

Table 6-20 lists the threshold values for Oracle Exadata Database Machine X3-2 Half Rack using a three-phase, high-voltage PDU.

Table 6-20 Threshold Values for Oracle Exadata Database Machine X3-2 Half Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	11	15
A and B	Module 1, phase 2	0	9	11
A and B	Module 1, phase 3	0	6	8

6.6.3.3 PDU Thresholds for Oracle Exadata Database Machine X3-2 Quarter Rack

The following tables list the threshold values for Oracle Exadata Database Machine X3-2 Quarter Rack:

- [Table 6-21](#) Threshold Values for Oracle Exadata Database Machine X3-2 Quarter Rack with Single-phase, Low-voltage PDU
- [Table 6-22](#) Threshold Values for Oracle Exadata Database Machine X3-2 Quarter Rack with Three-phase, Low-voltage PDU
- [Table 6-23](#) Threshold Values for Oracle Exadata Database Machine X3-2 Quarter Rack with Single-phase, High-voltage PDU
- [Table 6-24](#) Threshold Values for Oracle Exadata Database Machine X3-2 Quarter Rack with Three-phase, High-voltage PDU

Table 6-21 lists the threshold values for Oracle Exadata Database Machine X3-2 Quarter Rack using a single-phase, low-voltage PDU.

Table 6-21 Threshold Values for Oracle Exadata Database Machine X3-2 Quarter Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	10	13
A	Module 1, phase 2	0	5	6
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	5	6
B	Module 1, phase 3	0	10	13

Table 6-22 lists the threshold values for Oracle Exadata Database Machine X3-2 Quarter Rack using a three-phase, low-voltage PDU.

Table 6-22 Threshold Values for Oracle Exadata Database Machine X3-2 Quarter Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	11	14
A and B	Module 1, phase 2	0	10	13
A and B	Module 1, phase 3	0	4	6

Table 6-23 lists the threshold values for Oracle Exadata Database Machine X3-2 Quarter Rack using a single-phase, high-voltage PDU.

Table 6-23 Threshold Values for Oracle Exadata Database Machine X3-2 Quarter Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	9	11
A	Module 1, phase 2	0	4	6
A	Module 1, phase 3	0	0	0

Table 6-23 (Cont.) Threshold Values for Oracle Exadata Database Machine X3-2 Quarter Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	4	6
B	Module 1, phase 3	0	9	11

Table 6-24 lists the threshold values for Oracle Exadata Database Machine X3-2 Quarter Rack using a three-phase, high-voltage PDU.

Table 6-24 Threshold Values for Oracle Exadata Database Machine X3-2 Quarter Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	9	12
A and B	Module 1, phase 2	0	2	3
A and B	Module 1, phase 3	0	3	4

6.6.4 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers)

This section contains the PDU thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers).

Note:

The values listed here are reliable for use on systems in the configuration received upon deployment and delivery from the Oracle Factory. Any changes to this configuration should be modeled in Oracle Exadata Configuration Assistant (OECA), and the PDU thresholds listed in OECA should then be used. OECA is available on Oracle Technology Network.

- [PDU Thresholds for Oracle Exadata Database Machine X2-2 \(with X4170 M2 and X4270 M2 servers\) Full Rack](#)
- [PDU Thresholds for Oracle Exadata Database Machine X2-2 \(with X4170 M2 and X4270 M2 servers\) Half Rack](#)

- [PDU Thresholds for Oracle Exadata Database Machine X2-2 \(with X4170 M2 and X4270 M2 servers\) Quarter Rack](#)

6.6.4.1 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack

The following tables list the threshold values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack:

- [Table 6-25](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack with Single-phase, Low-voltage PDU
- [Table 6-26](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack with Three-phase, Low-voltage PDU
- [Table 6-27](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack with Single-phase, High-voltage PDU
- [Table 6-28](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack with Three-phase, High-voltage PDU

[Table 6-25](#) lists the threshold values for Oracle Exadata Database Machine Full Rack using a single-phase, low-voltage PDU.

Table 6-25 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	18	23
A	Module 1, phase 2	0	22	24
A	Module 1, phase 3	0	18	23
B	Module 1, phase 1	0	18	23
B	Module 1, phase 2	0	22	24
B	Module 1, phase 3	0	18	23

[Table 6-26](#) lists the threshold values for Oracle Exadata Database Machine Full Rack using a three-phase, low-voltage PDU.

Table 6-26 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	32	40

Table 6-26 (Cont.) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 2	0	34	43
A and B	Module 1, phase 3	0	33	42

Table 6-27 lists the threshold values for Oracle Exadata Database Machine Full Rack using a single-phase, high-voltage PDU.

Table 6-27 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	16	20
A	Module 1, phase 2	0	20	21
A	Module 1, phase 3	0	16	20
B	Module 1, phase 1	0	16	20
B	Module 1, phase 2	0	20	21
B	Module 1, phase 3	0	16	20

Table 6-28 lists the threshold values for Oracle Exadata Database Machine Full Rack using a three-phase, high-voltage PDU.

Table 6-28 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	18	21
A and B	Module 1, phase 2	0	18	21
A and B	Module 1, phase 3	0	17	21

6.6.4.2 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack

The following tables list the threshold values for Oracle Exadata Database Machine Half Rack:

- [Table 6-29](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack with Single-phase, Low-voltage PDU
- [Table 6-30](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack with Three-phase, Low-voltage PDU
- [Table 6-31](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack with Single-phase, High-voltage PDU
- [Table 6-32](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack with Three-phase, High-voltage PDU

[Table 6-29](#) lists the threshold values for Oracle Exadata Database Machine Half Rack using a single-phase, low-voltage PDU.

Table 6-29 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	18	23
A	Module 1, phase 2	0	13	17
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	13	17
B	Module 1, phase 3	0	18	23

[Table 6-30](#) lists the threshold values for Oracle Exadata Database Machine Half Rack using a three-phase, low-voltage PDU.

Table 6-30 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	18	23
A and B	Module 1, phase 2	0	17	22

Table 6-30 (Cont.) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 3	0	18	23

Table 6-31 lists the threshold values for Oracle Exadata Database Machine Half Rack using a single-phase, high-voltage PDU.

Table 6-31 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	16	20
A	Module 1, phase 2	0	12	15
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	12	15
B	Module 1, phase 3	0	16	20

Table 6-32 lists the threshold values for Oracle Exadata Database Machine Half Rack using a three-phase, high-voltage PDU.

Table 6-32 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	9	12
A and B	Module 1, phase 2	0	9	12
A and B	Module 1, phase 3	0	10	13

6.6.4.3 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack

The following tables list the threshold values for Oracle Exadata Database Machine Quarter Rack:

- [Table 6-33](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack with Single-phase, Low-voltage PDU
- [Table 6-34](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack with Three-phase, Low-voltage PDU
- [Table 6-35](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack with Single-phase, High-voltage PDU
- [Table 6-36](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack with Three-phase, High-voltage PDU

[Table 6-33](#) lists the threshold values for Oracle Exadata Database Machine Quarter Rack using a single-phase, low-voltage PDU.

Table 6-33 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	10	13
A	Module 1, phase 2	0	6	8
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	6	8
B	Module 1, phase 3	0	10	13

[Table 6-34](#) lists the threshold values for Oracle Exadata Database Machine Quarter Rack using a three-phase, low-voltage PDU.

Table 6-34 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	12	15
A and B	Module 1, phase 2	0	11	14
A and B	Module 1, phase 3	0	5	7

Table 6-35 lists the threshold values for Oracle Exadata Database Machine Quarter Rack using a single-phase, high-voltage PDU.

Table 6-35 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	9	12
A	Module 1, phase 2	0	5	7
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	5	7
B	Module 1, phase 3	0	9	12

Table 6-36 lists the threshold values for Oracle Exadata Database Machine Quarter Rack using a three-phase, high-voltage PDU.

Table 6-36 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	9	12
A and B	Module 1, phase 2	0	3	4

Table 6-36 (Cont.) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 3	0	4	5

6.6.5 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers)

This section contains the PDU thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers).

Note:

The values listed here are reliable for use on systems in the configuration received upon deployment and delivery from the Oracle Factory. Any changes to this configuration should be modeled in Oracle Exadata Configuration Assistant (OECA), and the PDU thresholds listed in OECA should then be used. OECA can be downloaded from Oracle Technology Network.

- [PDU Thresholds for Oracle Exadata Database Machine X2-2 \(with X4170 and X4275 servers\) Full Rack](#)
- [PDU Thresholds for Oracle Exadata Database Machine X2-2 \(with X4170 and X4275 servers\) Half Rack](#)
- [PDU Thresholds for Oracle Exadata Database Machine X2-2 \(with X4170 and X4275 servers\) Quarter Rack](#)

6.6.5.1 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack

The following tables list the threshold values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack:

- [Table 6-37](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack with Single-phase, Low-voltage PDU
- [Table 6-38](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack with Three-phase, Low-voltage PDU
- [Table 6-39](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack with Single-phase, High-voltage PDU
- [Table 6-40](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack with Three-phase, High-voltage PDU

Table 6-37 lists the threshold values for Oracle Exadata Database Machine Full Rack using a single-phase, low-voltage PDU.

Table 6-37 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	22	24
A	Module 1, phase 2	0	22	24
A	Module 1, phase 3	0	21	24
B	Module 1, phase 1	0	21	24
B	Module 1, phase 2	0	22	24
B	Module 1, phase 3	0	22	24

Table 6-38 lists the threshold values for Oracle Exadata Database Machine Full Rack using a three-phase, low-voltage PDU.

Table 6-38 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	36	45
A and B	Module 1, phase 2	0	39	44
A and B	Module 1, phase 3	0	38	45

Table 6-39 lists the threshold values for Oracle Exadata Database Machine Full Rack using a single-phase, high-voltage PDU.

Table 6-39 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	20	21
A	Module 1, phase 2	0	20	21
A	Module 1, phase 3	0	19	21

Table 6-39 (Cont.) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
B	Module 1, phase 1	0	19	21
B	Module 1, phase 2	0	20	21
B	Module 1, phase 3	0	20	21

[Table 6-40](#) lists the threshold values for Oracle Exadata Database Machine Full Rack using a three-phase, high-voltage PDU.

Table 6-40 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	20	21
A and B	Module 1, phase 2	0	21	21
A and B	Module 1, phase 3	0	18	21

6.6.5.2 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack

The following tables list the threshold values for Oracle Exadata Database Machine Half Rack:

- [Table 6-41](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack with Single-phase, Low-voltage PDU
- [Table 6-42](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack with Three-phase, Low-voltage PDU
- [Table 6-43](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack with Single-phase, High-voltage PDU
- [Table 6-44](#) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack with Three-phase, High-voltage PDU

[Table 6-41](#) lists the threshold values for Oracle Exadata Database Machine Half Rack using a single-phase, low-voltage PDU.

Table 6-41 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	21	24
A	Module 1, phase 2	0	13	17
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	13	17
B	Module 1, phase 3	0	21	24

[Table 6-42](#) lists the threshold values for Oracle Exadata Database Machine Half Rack using a three-phase, low-voltage PDU.

Table 6-42 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	19	24
A and B	Module 1, phase 2	0	20	25
A and B	Module 1, phase 3	0	19	24

[Table 6-43](#) lists the threshold values for Oracle Exadata Database Machine Half Rack using a single-phase, high-voltage PDU.

Table 6-43 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	19	21
A	Module 1, phase 2	0	12	15
A	Module 1, phase 3	0	0	0

Table 6-43 (Cont.) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	12	15
B	Module 1, phase 3	0	19	21

Table 6-44 lists the threshold values for Oracle Exadata Database Machine Half Rack using a three-phase, high-voltage PDU.

Table 6-44 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	11	14
A and B	Module 1, phase 2	0	11	14
A and B	Module 1, phase 3	0	10	13

6.6.5.3 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack

The following tables list the threshold values for Oracle Exadata Database Machine Quarter Rack:

- [Table 6-45](#)
Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack with Single-phase, Low-voltage PDU
- [Table 6-46](#)
Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack with Three-phase, Low-voltage PDU
- [Table 6-47](#)
Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack with Single-phase, High-voltage PDU
- [Table 6-48](#)
Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack with Three-phase, High-voltage PDU

Table 6-45 lists the threshold values for Oracle Exadata Database Machine Quarter Rack using a single-phase, low-voltage PDU.

Table 6-45 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	12	15
A	Module 1, phase 2	0	6	8
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	6	8
B	Module 1, phase 3	0	12	15

Table 6-46 lists the threshold values for Oracle Exadata Database Machine Quarter Rack using a three-phase, low-voltage PDU.

Table 6-46 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	13	17
A and B	Module 1, phase 2	0	12	15
A and B	Module 1, phase 3	0	5	7

Table 6-47 lists the threshold values for Oracle Exadata Database Machine Quarter Rack using a single-phase, high-voltage PDU.

Table 6-47 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	11	14
A	Module 1, phase 2	0	5	7
A	Module 1, phase 3	0	0	0

Table 6-47 (Cont.) Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	5	7
B	Module 1, phase 3	0	11	14

Table 6-48 lists the threshold values for Oracle Exadata Database Machine Quarter Rack using a three-phase, high-voltage PDU.

Table 6-48 Threshold Values for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	11	14
A and B	Module 1, phase 2	0	2	3
A and B	Module 1, phase 3	0	3	4

6.6.6 PDU Thresholds for Oracle Exadata Database Machine X4-8 and Later

This section contains the PDU thresholds for Oracle Exadata Database Machine Eight-Socket systems for X4-8 and later.

Oracle Exadata Configuration Assistant (OECA) is the only source for PDU thresholds for Oracle Exadata Database Machine X4-8 and later systems. OECA is available on Oracle Technology Network at <http://www.oracle.com/technetwork/database/exadata/oeca-download-2817713.html>

If for some reason you are unable to utilize OECA, contact Oracle Support.

6.6.7 PDU Thresholds for Oracle Exadata Database Machine X3-8 Full Rack

This section contains the PDU thresholds for Oracle Exadata Database Machine X3-8 Full Rack.

- [Table 6-49](#) Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Single-phase, Low-voltage 22 kVA PDU

- [Table 6-50](#) Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Three-phase, Low-voltage 24 kVA PDU
- [Table 6-51](#) Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Single-phase, High-voltage 22 kVA PDU
- [Table 6-52](#) Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Three-phase, High-voltage 24 kVA PDU

[Table 6-49](#) lists the threshold values for Oracle Exadata Database Machine X3-8 Full Rack using a single-phase, low-voltage 22 kVA PDU.

Table 6-49 Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Single-phase, Low-voltage 22 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	25	31
A	Module 1, phase 2	0	26	33
A	Module 1, phase 3	0	24	30
B	Module 1, phase 1	0	24	30
B	Module 1, phase 2	0	26	33
B	Module 1, phase 3	0	25	31

[Table 6-50](#) lists the threshold values for Oracle Exadata Database Machine X3-8 Full Rack using a three-phase, low-voltage 24 kVA PDU.

Table 6-50 Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Three-phase, Low-voltage 24 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	21	27
A	Module 1, phase 2	0	23	30
A	Module 1, phase 3	0	22	28
A	Module 2, phase 1	0	20	25
A	Module 2, phase 2	0	21	27
A	Module 2, phase 3	0	23	29
B	Module 1, phase 1	0	20	25

Table 6-50 (Cont.) Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Three-phase, Low-voltage 24 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
B	Module 1, phase 2	0	21	27
B	Module 1, phase 3	0	23	29
B	Module 2, phase 1	0	21	27
B	Module 2, phase 2	0	23	30
B	Module 2, phase 3	0	22	28

[Table 6-51](#) lists the threshold values for Oracle Exadata Database Machine X3-8 Full Rack using a single-phase, high-voltage 22 kVA PDU.

Table 6-51 Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Single-phase, High-voltage 22 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	22	28
A and B	Module 1, phase 2	0	24	30
A and B	Module 1, phase 3	0	22	28

[Table 6-52](#) lists the threshold values for Oracle Exadata Database Machine X3-8 Full Rack using a three-phase, high-voltage 24 kVA PDU.

Table 6-52 Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Three-phase, High-voltage 24 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	12	16
A	Module 1, phase 2	0	13	17
A	Module 1, phase 3	0	10	13
A	Module 2, phase 1	0	10	13

Table 6-52 (Cont.) Threshold Values for Oracle Exadata Database Machine X3-8 Full Rack with Three-phase, High-voltage 24 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 2, phase 2	0	13	17
A	Module 2, phase 3	0	11	15
B	Module 1, phase 1	0	10	13
B	Module 1, phase 2	0	13	17
B	Module 1, phase 3	0	11	15
B	Module 2, phase 1	0	12	16
B	Module 2, phase 2	0	13	17
B	Module 2, phase 3	0	10	13

6.6.8 PDU Thresholds for Oracle Exadata Database Machine X2-8 Full Rack

This section contains the PDU thresholds for Oracle Exadata Database Machine X2-8 Full Rack.

- [Table 6-53](#) Threshold Values for Oracle Exadata Database Machine X2-8 Full Rack with Single-phase, Low-voltage 22 kVA PDU
- [Table 6-54](#) Threshold Values for Oracle Exadata Database Machine X2-8 Full Rack with Three-phase, Low-voltage 24 kVA PDU
- [Table 6-55](#) Threshold Values for Oracle Exadata Database Machine X2-8 Full Rack with Single-phase, High-voltage 22 kVA PDU
- [Table 6-56](#) Threshold Values for Oracle Exadata Database Machine X2-8 Full Rack with Three-phase, High-voltage 24 kVA PDU

[Table 6-53](#) lists the threshold values for Oracle Exadata Database Machine X2-8 Full Rack using a single-phase, low-voltage 22 kVA PDU.

Table 6-53 Threshold Values for Oracle Exadata Database Machine X2-8 Full Rack with Single-phase, Low-voltage 22 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	36	37

Table 6-53 (Cont.) Threshold Values for Oracle Exadata Database Machine X2-8 Full Rack with Single-phase, Low-voltage 22 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 2	0	35	37
A	Module 1, phase 3	0	35	37
B	Module 1, phase 1	0	35	37
B	Module 1, phase 2	0	35	37
B	Module 1, phase 3	0	36	37

Table 6-54 lists the threshold values for Oracle Exadata Database Machine X2-8 Full Rack using a three-phase, low-voltage 24 kVA PDU.

Table 6-54 Threshold Values for Oracle Exadata Database Machine X2-8 Full Rack with Three-phase, Low-voltage 24 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	32	40
A and B	Module 1, phase 2	0	34	43
A and B	Module 1, phase 3	0	33	42

Table 6-55 lists the threshold values for Oracle Exadata Database Machine X2-8 Full Rack using a single-phase, high-voltage 22 kVA PDU.

Table 6-55 Threshold Values for Oracle Exadata Database Machine X2-8 Full Rack with Single-phase, High-voltage 22 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	26	32
A and B	Module 1, phase 2	0	25	32
A and B	Module 1, phase 3	0	26	32

Table 6-56 lists the threshold values for Oracle Exadata Database Machine X2-8 Full Rack using a three-phase, high-voltage 24 kVA PDU.

Table 6-56 Threshold Values for Oracle Exadata Database Machine X2-8 Full Rack with Three-phase, High-voltage 24 kVA PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	18	21
A and B	Module 1, phase 2	0	18	21
A and B	Module 1, phase 3	0	17	21

6.6.9 PDU Thresholds for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers

This section contains the PDU thresholds for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers.

- [PDU Thresholds for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers](#)
- [PDU Thresholds for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers](#)
- [PDU Thresholds for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers](#)

6.6.9.1 PDU Thresholds for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers

The following tables list the threshold values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers Full Rack:

- [Table 6-57](#) Threshold Values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, Low-voltage PDU
- [Table 6-58](#) Threshold Values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, Low-voltage PDU
- [Table 6-59](#) Threshold Values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, High-voltage PDU
- [Table 6-60](#) Threshold Values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, High-voltage PDU

[Table 6-57](#) lists the threshold values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a single-phase, low-voltage PDU.

Table 6-57 Threshold Values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	18	23
A	Module 1, phase 2	0	22	24
A	Module 1, phase 3	0	18	23
B	Module 1, phase 1	0	18	23
B	Module 1, phase 2	0	22	24
B	Module 1, phase 3	0	18	23

Table 6-58 lists the threshold values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a three-phase, low-voltage PDU.

Table 6-58 Threshold Values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	32	40
A and B	Module 1, phase 2	0	34	43
A and B	Module 1, phase 3	0	33	42

Table 6-59 lists the threshold values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a single-phase, high-voltage PDU.

Table 6-59 Threshold Values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	16	20

Table 6-59 (Cont.) Threshold Values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 2	0	20	21
A	Module 1, phase 3	0	16	20
B	Module 1, phase 1	0	16	20
B	Module 1, phase 2	0	20	21
B	Module 1, phase 3	0	16	20

Table 6-60 lists the threshold values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a three-phase, high-voltage PDU.

Table 6-60 Threshold Values for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	18	21
A and B	Module 1, phase 2	0	18	21
A and B	Module 1, phase 3	0	17	21

6.6.9.2 PDU Thresholds for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers

The following tables list the threshold values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers:

- [Table 6-61](#) Threshold Values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, Low-voltage PDU
- [Table 6-62](#) Threshold Values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, Low-voltage PDU
- [Table 6-63](#) Threshold Values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, High-voltage PDU

- [Table 6-64](#) Threshold Values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, High-voltage PDU

[Table 6-61](#) lists the threshold values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a single-phase, low-voltage PDU.

Table 6-61 Threshold Values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	18	23
A	Module 1, phase 2	0	13	17
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	13	17
B	Module 1, phase 3	0	18	23

[Table 6-62](#) lists the threshold values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a three-phase, low-voltage PDU.

Table 6-62 Threshold Values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	18	23
A and B	Module 1, phase 2	0	17	22
A and B	Module 1, phase 3	0	18	23

[Table 6-63](#) lists the threshold values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a single-phase, high-voltage PDU.

Table 6-63 Threshold Values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	16	20
A	Module 1, phase 2	0	12	15
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	12	15
B	Module 1, phase 3	0	16	20

Table 6-64 lists the threshold values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a three-phase, high-voltage PDU.

Table 6-64 Threshold Values for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	9	12
A and B	Module 1, phase 2	0	9	12
A and B	Module 1, phase 3	0	10	13

6.6.9.3 PDU Thresholds for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers

The following tables list the threshold values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers:

- [Table 6-65](#) Threshold Values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, Low-voltage PDU
- [Table 6-66](#) Threshold Values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, Low-voltage PDU

- [Table 6-67](#) Threshold Values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, High-voltage PDU
- [Table 6-68](#) Threshold Values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, High-voltage PDU

[Table 6-65](#) lists the threshold values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a single-phase, low-voltage PDU.

Table 6-65 Threshold Values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	10	13
A	Module 1, phase 2	0	6	8
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	6	8
B	Module 1, phase 3	0	10	13

[Table 6-66](#) lists the threshold values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a three-phase, low-voltage PDU.

Table 6-66 Threshold Values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, Low-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	12	15
A and B	Module 1, phase 2	0	11	14
A and B	Module 1, phase 3	0	5	7

[Table 6-67](#) lists the threshold values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a single-phase, high-voltage PDU.

Table 6-67 Threshold Values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Single-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	9	12
A	Module 1, phase 2	0	5	7
A	Module 1, phase 3	0	0	0
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	5	7
B	Module 1, phase 3	0	9	12

[Table 6-68](#) lists the threshold values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers using a three-phase, high-voltage PDU.

Table 6-68 Threshold Values for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers with Three-phase, High-voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	9	12
A and B	Module 1, phase 2	0	3	4
A and B	Module 1, phase 3	0	4	5

6.7 Checking Exadata Storage Servers

After powering up the rack, check the Oracle Exadata Storage Servers.

 **Note:**

- (X2-2 only) When using the KVM switch, and pressing the ESC key, the BIOS may receive two ESC characters, and prompt to exit. Select **CANCEL**.
- If you need to connect to the ILOM serial management port, then the baud rate setting on Oracle Exadata Storage Servers is changed from the default ILOM setting 9600 to 115200 baud, 8 bits, no parity, and 1 stop bit.

1. Log in to the first storage server as the `root` user.

 **Note:**

If you do not have the password for the `root` user, then contact Oracle Support Services.

2. Verify the SysSN label for the storage server matches the `product_serial_number` value using the following command:

```
#ipmitool sunoem cli "show /SYS product_serial_number"
Connected.Use ^ D to exit.
->show /SYS product_serial_number
/SYS
Properties:
product_serial_number =0937XFG03B
->Session closed
Disconnected
```

3. Check the hardware and firmware using the following command:

```
# /opt/oracle.SupportTools/CheckHWnFWProfile
```

If output does not show success, then examine the output for mismatched hardware or firmware and take steps to correct those as explained in the following steps.

4. Verify the disks on the server are visible and online using the following command. The disks are numbered slot 0 to slot 11.
 - If you are using Oracle Exadata System Software release 19c or later, then get the enclosure ID:

```
# cd /opt/MegaRAID/storcli
# ./storcli64 /c0/eall show
Controller = 0
Status = Success
Description = None

Properties :
```

```

=====
-----
EID State Slots PD PS Fans TSS Alms SIM
Port#                               ProdID VendorSpecific
-----
      8 OK          14 12 0    0  0  0  0 Port 0 - 3 & Port 4 - 7
x8 X5-2L  SPMS    v`
252 OK          8  0  0    0  0  0  1
Internal                               SGPIO
-----
-----

```

EID-Enclosure Device ID |PD-Physical drive count |PS-Power Supply count|
TSS-Temperature sensor count |Alms-Alarm count |SIM-SIM Count

Then show the state of the physical devices in that enclosure:

```

# cd /opt/MegaRAID/storcli
# ./storcli64 /c0/e8/sall show all | grep "8:"
8:0          16 Onln   0 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:1          17 Onln   1 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:2          20 Onln   2 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:3          15 Onln   3 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:4          12 Onln   4 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:5          18 Onln   5 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:6          11 Onln   6 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:7           9 Onln   7 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:8          14 Onln   8 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:9          21 Onln  11 7.151 TB SAS HDD N  N  512B
H7280B520SUN8.0T U
8:10         13 Onln   9 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U
8:11         10 Onln  10 7.151 TB SAS HDD N  N  512B
H7280A520SUN8.0T U

```

The rows displayed are from a table that displays the following information:

```

EID:Slt DID State DG      Size Intf Med SED PI SeSz
Model                               Sp
-----

```



```

-----
...
EID-Enclosure Device ID|Slt-Slot No.|DID-Device ID|DG-DriveGroup
DHS-Dedicated Hot Spare|UGood-Unconfigured Good|GHS-Global
Hotspare
UBad-Unconfigured Bad|Onln-Online|Offln-Offline|Intf-Interface
Med-Media Type|SED-Self Encryptive Drive|PI-Protection Info
SeSz-Sector Size|Sp-Spun|U-Up|D-Down|T-Transition|F-Foreign
UGUnsp-Unsupported|UGShld-UnConfigured shielded|HSPShld-Hotspare
shielded
CFShld-Configured shielded|Cpybck-CopyBack|CBSHld-Copyback
Shielded

```

- If you are using Oracle Exadata System Software release 18c or earlier, then use the following command:

```

#cd /opt/MegaRAID/MegaCli
#./MegaCli64 -Pdlist -aAll |grep "Slot \|Firmware"
Slot Number:0
Firmware state:Online, Spun Up
Slot Number:1
Firmware state:Online, Spun Up
Slot Number:2
Firmware state:Online, Spun Up
Slot Number:3
Firmware state:Online, Spun Up
Slot Number:4
Firmware state:Online, Spun Up
Slot Number:5
Firmware state:Online, Spun Up
Slot Number:6
Firmware state:Online, Spun Up
Slot Number:7
Firmware state:Online, Spun Up
Slot Number:8
Firmware state:Online, Spun Up
Slot Number:9
Firmware state:Online, Spun Up
Slot Number:10
Firmware state:Online, Spun Up
Slot Number:11
Firmware state:Online, Spun Up
[root@cell01 ~]#

```

5. If applicable, verify the PMEM modules are visible.

```
# cellcli -e list physicaldisk where disktype=PMEM
```

6. Verify the flash modules are visible. The number of cards and modules varies depending on the model.

```
# cellcli -e list physicaldisk where disktype=flashdisk
```

If any flash modules are missing, then power cycle the system, and wait 10 minutes after restarting before checking again. If this does not resolve the problem, then reseal the module and flash card. If the problem continues, then replace the card with the on-site spare.

7. Check the rear of the rack to ensure the flash card LEDs are green all the way across. If any LED is amber, then check again after the next step. If the LED is still amber after restart, then reseal the card and ESM or flash module. If the problem continues, then replace the card with the on-site spare.

 **Note:**

This step can be done when all systems are on. It does not have to be done on a per-server basis.

8. Log out of the system using the following command:

```
# logout
```

9. (X2-2 systems) Return to the User Interface on the KVM switch by pressing the Ctrl key on the left to return to the KVM switch, and select **Disconnect Session**, if using the KVM switch to access the server.
10. (X2-2 systems) Select **Target Devices** from Unit View, if using the KVM switch to access the server.
11. Proceed to the next Oracle Exadata Storage Server.

6.8 Checking Oracle Exadata Database Servers

Before installing the software, check the configuration of the Oracle Exadata Database Servers.

1. Log in as the `root` user to the first database server.

 **Note:**

If you do not have the password for the `root` user, then contact Oracle Support Services.

2. Verify the SysSN label for the server matches the `product_serial_number` value using the following command:

```
# ipmitool sunoem cli "show /SYS product_serial_number"
Connected.Use ^ D to exit.
->show /SYS product_serial_number

/SYS
  Properties:
    product_serial_number =093ABCD5001

->Session closed
Disconnected
```

3. Change to the `/opt/oracle.SupportTools/firstconf` directory.

```
cd /opt/oracle.SupportTools/firstconf
```

4. Create a group file with the list of server IP addresses to use in the `dcli` utility commands, if required.

If you are configuring Oracle Exadata Rack X8M, then you must manually create a file and add the IP addresses.

On racks that use InfiniBand Network Fabric, you can use the following command to identify the servers and copy the results into a text file.

```
# ibhosts | awk '/elastic/ {print $8}' | cut -c1-14 > group_file
```

In the preceding command, `group_file` is name of the file to create with the server IP addresses.

5. Verify the hardware profile is correct using the following command. The command checks CPU type and count, component firmware, and so on.

```
# dcli -l root -g group_file "/opt/oracle.SupportTools/  
CheckHWnFWProfile \  
-c strict" > /tmp/checkhfw.out
```

```
# more /tmp/checkhfw.out
```

The following is an example of a check that failed:

```
172.108.1.6:[WARNING ]The hardware and firmware are not supported.  
See details below  
[DiskControllerPCIEslotWidth ]  
Requires:  
x8  
Found:  
x4  
[WARNING ]The hardware and firmware are not supported.See details  
above
```

If a profile check fails and the firmware needs to be refreshed, then run the following command from the system with the problem.

```
# /opt/oracle.SupportTools/CheckHWnFWProfile -U /opt/  
oracle.cellos/iso/cellbits  
Now updating the ILOM and the BIOS ...
```

The command must be entered on a single line. After running the command, power off the system, reset the Integrated Lights Out Manager (ILOM), and wait 10 minutes before powering on the system.

6. Verify the disks on the server are visible and online.

- If you are using Oracle Exadata System Software release 19c or later, then use the following command:

```
# cd /opt/MegaRAID/storcli
# ./storcli64 /call show | grep "Physical Drives" -A 25
Physical Drives = 8
```

```
PD LIST :
=====
```

```
-----
-----
EID:SlT DID State DG          Size Intf Med SED PI SeSz
Model                Sp Type
-----
-----
39:0      31 Onln  0 278.464 GB SAS  HDD N  N  512B
ST930003SSUN300G U -
39:1      35 Onln  0 278.464 GB SAS  HDD N  N  512B
ST930003SSUN300G U -
39:2      41 Onln  0 278.464 GB SAS  HDD N  N  512B
H103030SCSUN300G U -
39:3      37 Onln  0 278.464 GB SAS  HDD N  N  512B
ST930003SSUN300G U -
39:4      38 Onln  0 278.464 GB SAS  HDD N  N  512B
ST930003SSUN300G U -
39:5      44 Onln  0 278.464 GB SAS  HDD N  N  512B
H109030SESUN300G U -
39:6      33 Onln  0 278.464 GB SAS  HDD N  N  512B
ST930003SSUN300G U -
39:7      43 Onln  0 278.464 GB SAS  HDD N  N  512B
ST930003SSUN300G U -
-----
-----
```

```
EID-Enclosure Device ID|SlT-Slot No.|DID-Device ID|DG-DriveGroup
DHS-Dedicated Hot Spare|UGood-Unconfigured Good|GHS-Global
Hotspare
UBad-Unconfigured Bad|Onln-Online|Offln-Offline|Intf-Interface
Med-Media Type|SED-Self Encryptive Drive|PI-Protection Info
SeSz-Sector Size|Sp-Spun|U-Up|D-Down|T-Transition|F-Foreign
UGUnsp-Unsupported|UGShld-UnConfigured shielded|HSPShld-Hotspare
shielded
CFShld-Configured shielded|Cpybck-CopyBack|CBShld-Copyback
Shielded
```

- If you are using Oracle Exadata System Software release 18c or earlier, then use the following command:

```
#cd /opt/MegaRAID/MegaCli
#./MegaCli64 -Pdlist -aAll |grep "Slot \|Firmware"
Slot Number:0
Firmware state:Online, Spun Up
Slot Number:1
```

```
Firmware state:Online, Spun Up
Slot Number:2
Firmware state:Online, Spun Up
Slot Number:3
Firmware state:Online, Spun Up
Slot Number:4
Firmware state:Online, Spun Up
Slot Number:5
Firmware state:Online, Spun Up
Slot Number:6
Firmware state:Online, Spun Up
Slot Number:7
Firmware state:Online, Spun Up
Slot Number:8
Firmware state:Online, Spun Up
Slot Number:9
Firmware state:Online, Spun Up
Slot Number:10
Firmware state:Online, Spun Up
Slot Number:11
Firmware state:Online, Spun Up
[root@cell01 ~]#
```

7. Examine the output of the following command:

- If you are using Oracle Exadata System Software release 19c or later, then use the following command:

```
/opt/MegaRAID/storcli/storcli64 -LdInfo -LAll -aAll
```

- If you are using Oracle Exadata System Software release 18c or earlier, then you can use the following command:

```
/opt/MegaRAID/MegaCli/MegaCli64 -LdInfo -LAll -aAll
```

8. (Oracle Exadata Database Machine based on Sun Fire X4170 Oracle Database Servers only) Check that the ILOM is configured for the proper fan cooling on each database server using the following command:

```
# ipmitool sunoem cli "show /SP/policy"
```

- If the policy is present and enabled, then no additional steps are needed.
- If the policy is present and disabled, then enable the policy using the following command:

```
# ipmitool sunoem cli "set /SP/policy FLASH_ACCELERATOR_CARD_ \
INSTALLED=enabled"
```

The command must be entered as a single line.

9. Log out of the system using the following command:

```
# logout
```

10. (X2-2 racks) Return to the User Interface on the KVM switch by pressing the Ctrl key on the left to return to the KVM switch, and select **Disconnect Session**, if using the KVM switch to access the server.
11. (X2-2 racks) Select **Target Devices** from Unit View, and proceed to the next database server, if using the KVM switch to access the server.
12. Verify that NTP servers are configured the same for all servers using the following command. The command is run on each database server individually.

```
ntpq -p
```

6.9 Performing Additional Checks and Configuration

Perform these additional checks and steps to verify the configuration of your Oracle Exadata Rack.

1. (Oracle Exadata Database Machine only) Log in to the first database server as the root user.
2. Change to the `firstconf` directory using the following command:

```
# cd /opt/oracle.SupportTools/firstconf
```

3. Create a group file with the list of IP addresses to use in the `dcli` utility commands, if you want to query multiple racks.
4. Verify the rack master serial number is set correctly using the following command:

```
# dcli -l root -g group_file_name "ipmitool sunoem cli 'show \
/SP system_identifer'" | grep "system_identifer =" > /tmp/show-
rack-msn.out
```

```
# more /tmp/show-rack-msn.out
```

If none of the serial numbers are set, then contact Oracle Support Services. If one serial number is not set correctly, then run the following command to correct it, where `rack_type` is the rack designation, such as X8-8 and the string `serial_number` represents the serial number:

```
# ipmitool sunoem cli 'set /SP system_identifer= \
"Exadata Database Machine rack_type serial_number"'
```

5. If you are running an Oracle Exadata System Software release earlier than 18.1, then ensure disk cache policy is set to `Disabled` using the following command as the root user on the first database server in the cluster:

```
# dcli -g /opt/oracle.SupportTools/onecommand/all_group -l root \
/opt/MegaRAID/MegaCli/MegaCli64 -LdPdInfo -aALL | grep -i 'Disk
Cache Policy'
```

The following is an example of the output from the command:

```
dm01db01: Disk Cache Policy   : Disabled
dm01db02: Disk Cache Policy   : Disabled
dm01db03: Disk Cache Policy   : Disabled
...
dm01cel01: Disk Cache Policy   : Disabled
dm01cel02: Disk Cache Policy   : Disabled
...
```

If any server shows different output, then identify the LUN and use the following command to reset the policy to `Disabled`, where `LUNn` is the LUN number, such as `L2`:

```
# MegaCli64 -LDSetProp -DisDskCache -LUNn -a0
```

6.10 Verifying the RoCE Network Fabric Configuration

This procedure describes how to verify the RoCE Network Fabric configuration.

1. Verify the proper `oracle-rdma-release` software versions are being used on the database servers.

```
[root@dbm01adm08 ~]# rpm -qa |grep oracle-rdma-release
oracle-rdma-release-0.11.0-1.el7ora.x86_64
```

The `oracle-rdma-release` software and adapter firmware versions are automatically maintained on the Oracle Exadata Database Machine storage servers.

2. Check the adapter firmware versions on the database servers.

Use the `CheckHWnFWProfile` script to check firmware versions for the RDMA Network Fabric adapters.

```
# /opt/oracle.SupportTools/CheckHWnFWProfile -action list
```

3. Visually check all the RDMA Network Fabric cable connections within the rack.

The port lights should be on, and the LEDs should be on. Do not press each connector to verify connectivity.

4. Complete the steps described in My Oracle Support Doc ID 2587717.1

Related Topics

- [Verify RoCE Cabling on Oracle Exadata Database Machine X8M-2 and X8M-8 Servers \(My Oracle Support Doc ID 2587717.1\)](#)

6.11 Verifying the InfiniBand Network Fabric Network

This procedure describes how to verify the InfiniBand Network Fabric network.

1. Visually check all the RDMA Network Fabric cable connections within the rack. The port lights should be on, and the LEDs should be on. Do not press each connector to verify connectivity.
2. Log in as the `root` user on any component in the rack.
3. Verify the InfiniBand Network Fabric topology using the following commands:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify-topology [-t rack_size]
```

The following example shows the output when the network components are correct.

```
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch.....
[SUCCESS ]
Are any external switches connected to each other.....
[SUCCESS ]
Are any hosts connected to spine switch.....
[SUCCESS ]
Check if all hosts have 2 CAs to different switches.....
[SUCCESS ]
Leaf switch check:cardinality and even distribution.....
[SUCCESS ]
Check if each rack has an valid internal ring.....
[SUCCESS ]
```

In the preceding command, `rack_size` is the size of the rack. The `-t rack_size` option is needed if the rack is Oracle Exadata Database Machine Half Rack or Oracle Exadata Database Machine Quarter Rack. Use the value `halfrack` or `quarterrack`, if needed.

The following example shows the output when there is a bad RDMA Network Fabric switch to cable connection:

```
#./verify-topology
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch.....
[SUCCESS ]
Are any external switches connected to each other.....
[SUCCESS ]
Are any hosts connected to spine switch.....
[SUCCESS ]
Check if all hosts have 2 CAs to different switches.....
[SUCCESS ]
Leaf switch check:cardinality and even distribution.....
[SUCCESS ]
Check if each rack has an valid internal ring.....
[ERROR ]
```

```
Switches 0x21283a87cba0a0 0x21283a87b8a0a0 have 6 connections
between them.
They should have at least 7 links between them
```


The following example shows the output when there is a bad RDMA Network Fabric cable on a database server:

```

#./verify-topology
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch.....
[SUCCESS ]
Are any external switches connected to each other.....
[SUCCESS ]
Are any hosts connected to spine switch.....
[SUCCESS ]
Check if all hosts have 2 CAs to different switches.....
[ERROR ]
Node db01 has 1 endpoints.(Should be 2)
Port 2 of this node is not connected to any switch
-----fattree End Point Cabling verification failed-----
Leaf switch check:cardinality and even distribution.....
[ERROR ]
Internal QDR Switch 0x21283a87b8a0a0 has fewer than 4 compute nodes
It has only 3 links belonging to compute nodes
[SUCCESS ]
Check if each rack has an valid internal ring.....
[SUCCESS ]

```

The following example shows the output when there is a bad connection on the switch and the system:

```

#./verify-topology
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch.....
[SUCCESS ]
Are any external switches connected to each other.....
[SUCCESS ]
Are any hosts connected to spine switch.....
[SUCCESS ]
Check if all hosts have 2 CAs to different switches.....
[ERROR ]

Node burxdb01 has 1 endpoints.(Should be 2)
Port 2 of this node is not connected to any switch
-----fattree End Point Cabling verifation failed-----
Leaf switch check:cardinality and even distribution.....
[ERROR ]
Internal QDR Switch 0x21283a87b8a0a0 has fewer than 4 compute nodes
It has only 3 links belonging to compute nodes.....
[SUCCESS ]
Check if each rack has an valid internal ring.....
[ERROR ]

Switches 0x21283a87cba0a0 0x21283a87b8a0a0 have 6 connections
between them
They should have at least 7 links between them

```

6.12 Imaging a New System

You have different options for preparing an Oracle Exadata system for deployment.

Using a `preconf.csv` file with the PXE, USB and ISO options enables an unattended first boot configuration and greatly simplifies the installation but is not mandatory. It is possible to boot from the installation media and follow the prompts to configure the system.

Starting with Oracle Exadata System Software release 19.1.0, Secure Eraser is automatically started during re-imaging if the hardware supports Secure Eraser. This significantly simplifies the re-imaging procedure while maintaining performance. Now, when re-purposing a rack, you only have to image the rack and the secure data erasure is taken care of transparently as part of the process.

In release 12.1.2.2.0 and above, the installation media has been simplified, and it is no longer necessary to build files using the ImageMaker process.

- [Using ISO Images to Image a New System](#)
You can download ISO images to use when upgrading your Oracle Exadata Database Machine.
- [Using Elastic Configuration to Image a New System](#)
Elastic configuration is now the standard methodology for all new deployments, and applies to Oracle Exadata Database Machine X4-8 and later servers.
- [Using ISO to Image a New System](#)
You can use ISO to prepare a new Exadata system for deployment.
- [Using PXE to Image a New System](#)
You can use PXE to prepare a new Exadata system for deployment.
- [Using USB to Image a New System](#)
You can use USB to prepare a new Exadata system for deployment.
- [Updates to Imaging Parameters](#)
Some of the command line options for imaging have changed with 12.1.2.1.0 and above.

6.12.1 Using ISO Images to Image a New System

You can download ISO images to use when upgrading your Oracle Exadata Database Machine.

In versions prior to 12.1.2.2.0, ISO images are available as patches that you can download for bare metal imaging an Exadata system. This is now the default method in 12.1.2.2.0 and above. You are no longer required to build the ISO files using `imagemaker`. Patches are available for both database nodes and cells. You can use the ISO images to boot the server using a virtual media device on the ILOM (from the Devices menu of the ILOM console). For example, the two ISO images available for 12.1.2.1.1 are listed in the following table.

Table 6-69 ISO Images for 12.1.2.1.1

Version	Type	Edelivery Part / Notes	ARU Bug #/ULN Channel	Description
12.1.2.1.1	Cell ISO image and md5sum: 90305b9e2c7e050d7588b8a7d5c731e7 cell_12.1.2.1.1_LINU X.X64_150316.2-1.x 86_64.iso	Refer to 12.1.2.1.1 Readme (Doc ID 1959143.1) as well as the readme provided with the download.	20757971	Oracle Database Machine Exadata Storage Cell (X5-2L, X4-2L, X4270M3, X4270M2, X4275) Image 12c Release 1 (12.1.2.1.1) for Linux x86_64
12.1.2.1.1	DB ISO image and md5sum: 39f12722cb338b2d4de5acba90adf8fe compute_12.1.2.1.1_ LINUX.X64_150316. 2-1.x86_64.iso	Refer to 12.1.2.1.1 Readme (Doc ID 1959143.1) as well as the readme provided with the download.	20757974	Oracle Database Machine Database Host (X4-8, X4800M2, X4800, X5-2, X4-2, X4170M3, X4170M2, X4170) Image 12c Release 1 (12.1.2.1.1) for Linux x86_64

To prepare a new Exadata system for deployment using ISO images, perform the following steps:

1. Obtain the latest Oracle Exadata Deployment Assistant (OEDA).
2. Run the OEDA configuration tool to generate the configuration files. You will need the configuration files when you run `applyElasticConfig.sh` later.
3. Power on all the servers.
 - a. Database nodes and cells will automatically assign to themselves unused IP addresses on `eth0` in the 172.16 range.
 - b. Host names of the nodes will be of the format `nodeN`, where *N* is a number.
4. After booting from the ISO, follow the prompts and imaging will start automatically. You will be prompted to enter IP addresses.

You can use a `preconf.csv` by following the steps in [Using ISO to Image a New System](#) to mount the image on an NFS repository. You need to place the `preconf.csv` file in the same NFS directory.

5. If you are using virtualization on the database servers (Oracle VM Server or Oracle Linux KVM), then run `switch_to_ovm.sh` on each database server. The database servers will reboot if you are using Oracle VM Server.

For example:

```
# /opt/oracle.SupportTools/switch_to_ovm.sh
2014-12-07 11:58:36 -0800 [INFO] Switch to DOM0 system
partition /dev/VGExaDb/LVDbSys3 (/dev/mapper/VGExaDb-LVDbSys3)
2014-12-07 11:58:36 -0800 [INFO] Active system device: /dev/mapper/
VGExaDb-LVDbSys1
2014-12-07 11:58:36 -0800 [INFO] Active system device in boot
area: /dev/mapper/VGExaDb-LVDbSys1
```

```
2014-12-07 11:58:36 -0800 [INFO] Set active system device to /dev/
VGExaDb/LVDbSys3 in /boot/I_am_hd_boot
2014-12-07 11:58:36 -0800 [INFO] Reboot has been initiated to
switch to the DOM0 system partition
```

6. Run `reclaimdisks.sh` on each database server (except KVM environments on X8M servers) before installing any software.

The `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim` command reclaims disk space reserved for the deployment type not selected. The command takes less than 5 minutes, approximately. Systems are imaged with disks configured with RAID5; a RAID rebuild is no longer part of the `reclaimdisks.sh` process.

 **Caution:**

Do not skip this step. Skipping this step results in unused space that can no longer be reclaimed by `reclaimdisks.sh`.

```
# /opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim
Model is ORACLE SERVER X5-2
Number of LSI controllers: 1
Physical disks found: 4 (252:0 252:1 252:2 252:3)
Logical drives found: 1
...
[INFO      ] Copying /usr/share/grub/x86_64-redhat/* to /boot/
grub ...
[INFO      ] Create filesystem on device /dev/sda1
[INFO      ] Tune filesystem on device /dev/sda1

GNU GRUB  version 0.97  (640K lower / 3072K upper memory)

[ Minimal BASH-like line editing is supported.  For the first word,
TAB
lists possible command completions.  Anywhere else TAB lists the
possible
completions of a device/filename.]
grub> root (hd0,0)
  Filesystem type is ext2fs, partition type 0x83
grub> setup (hd0)
  Checking if "/boot/grub/stage1" exists... no
  Checking if "/grub/stage1" exists... yes
  Checking if "/grub/stage2" exists... yes
  Checking if "/grub/e2fs_stage1_5" exists... yes
  Running "embed /grub/e2fs_stage1_5 (hd0)"... failed (this is not
fatal)
  Running "embed /grub/e2fs_stage1_5 (hd0,0)"... failed (this is not
fatal)
  Running "install /grub/stage1 (hd0) /grub/stage2 p /grub/grub.conf
"... succeeded
Done.
grub> quit
```

7. If you have an Oracle Exadata Rack X8 or earlier, then run `ibhosts` and verify that all nodes show the correct IP addresses and host names. There should be no nodes with `elasticNode` in the description.
8. Run the OEDA tool to deploy.

```
[root] # ./install.sh -cf ../MAA-dm01.xml -l
1. Validate Configuration File
2. Setup Capacity-on-Demand
3. Create Virtual Machine
4. Create Users
5. Setup Cell Connectivity
6. Create Cell Disks
7. Create Grid Disks
8. Configure Alerting
9. Install Cluster Software
10. Initialize Cluster Software
11. Install Database Software
12. Relink Database with RDS
13. Create ASM Diskgroups
14. Create Databases
15. Apply Security Fixes
16. Install Exachk
17. Setup ASR Alerting
18. Create Installation Summary
19. Resecure Machine
```

Related Topics

- [Exadata 12.1.2.1.1 release and patch \(20240049\) \(My Oracle Support Doc ID 1959143.1\)](#)

6.12.2 Using Elastic Configuration to Image a New System

Elastic configuration is now the standard methodology for all new deployments, and applies to Oracle Exadata Database Machine X4-8 and later servers.

You can also use elastic configuration to add additional database servers or cells to an existing configuration. The process allows initial IP addresses to be assigned to database servers and cells, regardless of the exact customer configuration ordered.

Elastic configuration is only applicable to machines shipped from the factory (or to machines imaged exactly to mimic that) and only for X4-8 or later systems. You can run elastic configuration at the first time of deployment, and when adding new factory-built components to a rack. You cannot use elastic configuration to configure or reconfigure systems. It is recommended that you use elastic configuration only with machines shipped from the factory for first-time deployment.

To prepare a new Exadata system for deployment using the elastic configuration method, perform the following steps:

1. Download the latest version of Oracle Exadata Deployment Assistant (OEDA).
2. Run the OEDA configuration tool to generate the configuration files.
3. Power on all the servers.

- a. Database nodes and cells will automatically assign to themselves unused IP addresses on `eth0` in the `172.16.*` range.
 - b. Host names of the nodes will be of the format `nodeN`, where `N` is a number.
4. If you are using Oracle VM, then, before you run the OEDA installation tool, run `switch_to_ovm.sh` on each database server. The database servers will reboot. If you are using physical deployment and not virtualized deployment, then you can skip this step.

For example:

```
# /opt/oracle.SupportTools/switch_to_ovm.sh
2014-12-07 11:58:36 -0800 [INFO] Switch to DOM0 system
partition /dev/VGExaDb/LVDbSys3 (/dev/mapper/VGExaDb-LVDbSys3)
2014-12-07 11:58:36 -0800 [INFO] Active system device: /dev/mapper/
VGExaDb-LVDbSys1
2014-12-07 11:58:36 -0800 [INFO] Active system device in boot
area: /dev/mapper/VGExaDb-LVDbSys1
2014-12-07 11:58:36 -0800 [INFO] Set active system device to /dev/
VGExaDb/LVDbSys3 in /boot/I_am_hd_boot
2014-12-07 11:58:36 -0800 [INFO] Reboot has been initiated to
switch to the DOM0 system partition
```

5. Run `reclaimdisks.sh` on each database server (except KVM environments on X8M servers) before installing any software.

The `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim` command reclaims disk space reserved for the deployment type not selected. The command takes less than 5 minutes, approximately. Systems are imaged with disks configured with RAID5; a RAID rebuild is no longer part of the `reclaimdisks.sh` process.

Caution:

Do not skip this step. Skipping this step results in unused space that can no longer be reclaimed by `reclaimdisks.sh`.

```
# /opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim
Model is ORACLE SERVER X5-2
Number of LSI controllers: 1
Physical disks found: 4 (252:0 252:1 252:2 252:3)
Logical drives found: 1
...
[INFO      ] Copying /usr/share/grub/x86_64-redhat/* to /boot/
grub ...
[INFO      ] Create filesystem on device /dev/sda1
[INFO      ] Tune filesystem on device /dev/sda1

GNU GRUB  version 0.97  (640K lower / 3072K upper memory)

[ Minimal BASH-like line editing is supported.  For the first word,
TAB
lists possible command completions.  Anywhere else TAB lists the
```

```
possible
completions of a device/filename.]
grub> root (hd0,0)
  Filesystem type is ext2fs, partition type 0x83
grub> setup (hd0)
  Checking if "/boot/grub/stage1" exists... no
  Checking if "/grub/stage1" exists... yes
  Checking if "/grub/stage2" exists... yes
  Checking if "/grub/e2fs_stage1_5" exists... yes
  Running "embed /grub/e2fs_stage1_5 (hd0)"... failed (this is not
fatal)
  Running "embed /grub/e2fs_stage1_5 (hd0,0)"... failed (this is not
fatal)
  Running "install /grub/stage1 (hd0) /grub/stage2 p /grub/grub.conf
"... succeeded
Done.
grub> quit
```

6. Apply your specific configuration to the nodes by running the `applyElasticConfig.sh` script that is supplied with OEDA.
7. Run `ibhosts` and verify that all nodes show the correct IP addresses and host names. There should be no nodes with `elasticNode` in the description.
8. Run the OEDA tool to deploy, supplying the name of the configuration file generated by OEDA.

For example, if your rack is named `exadm01`, then you would use the following command:

```
[root] # ./install.sh -cf ../exadm01.xml -l
1. Validate Configuration File
2. Setup Capacity-on-Demand
3. Create Virtual Machine
4. Create Users
5. Setup Cell Connectivity
6. Create Cell Disks
7. Create Grid Disks
8. Configure Alerting
9. Install Cluster Software
10. Initialize Cluster Software
11. Install Database Software
12. Relink Database with RDS
13. Create ASM Diskgroups
14. Create Databases
15. Apply Security Fixes
16. Install Exachk
17. Setup ASR Alerting
18. Create Installation Summary
19. Resecure Machine
```

6.12.3 Using ISO to Image a New System

You can use ISO to prepare a new Exadata system for deployment.

1. Obtain the latest Oracle Exadata Deployment Assistant (OEDA).

2. Run the OEDA configuration tool to generate the configuration files.
3. Power on all the servers.
 - a. Database nodes and cells will automatically assign to themselves unused IP addresses on eth0 in the 172.16 range.
 - b. Host names of the nodes will be of the format `nodeN`, where *N* is a number.
4. Run OEDA to generate the `preconf.csv` file that can be used to image the system.
5. Connect to the first database node. You can do this by connecting a laptop to the ILOM serial management port (9600, 8, N, 1), and then connecting to the console and log in to the database node.
6. Obtain the eth0 MAC address for each database node. This is the `fru_macaddress` field from `/SYS/MB/NET0` in the ILOM (or through `ip addr` if logged into the console).
 - a. Connect to the other database nodes using `ssh` from the console.
 - b. Run `ibhosts` to get a list of node names and IP addresses.
7. Insert the MAC addresses using uppercase characters in the seventh field of the `preconf.csv` file for each node. This field should be empty, that is, you should see two consecutive comma characters (`,,`). For example:

```
orhb42-b08-01,example.com,db,eth0,eth0,Management,00:10:E0:69:34:52,
10.196.3.106,255.255.224.0,10.196.31.250,orhb42-b08-01-
priv,Private:ac
tive-bond-ib,192.168.10.15,255.255.255.0,,America/Los_Angeles
```

8. Run `ipconf` to verify the generated `preconf.csv` file. This can be done on any Exadata server, and although not required, it is useful to ensure the `preconf.csv` is valid before proceeding.

```
# ipconf -verify -preconf /tmp/preconf.csv
[Info]: ipconf command line: /opt/oracle.cellos/ipconf.pl -verify -
preconf /tmp
/preconf -nocodes
[Done]: Pre config verification OK
```

9. If you are using versions prior to 12.1.2.2.0, proceed to step [11](#) to use ImageMaker to create the image files.

If you are using version 12.1.2.2.0 or later, it is no longer necessary to use `imagemaker` to create the image files. The relevant PXE, ISO, and USB image files have already been created in the patches themselves. Simply download the patch containing the prebuilt image file from Oracle Software Delivery Cloud and place the downloaded ISO `.iso` image file in the appropriate location. For 12.1.2.2.0 the patch numbers can be obtained from My Oracle Support note 2038073.1. There are separate ImageMaker files for database nodes and cells.

Note: For image version 12.1.2.2.0 the format of the downloadable image files has changed.

There will be an image file in the downloaded zip file. For example, the database node zip file will contain:

```
compute_12.1.2.2.0_LINUX.X64_150917-1.x86_64.iso
```

For cell image, the file will be similar to:

```
cell_12.1.2.2.0_LINUX.X64_150917-1.x86_64.iso
```

10. (For 12.1.2.2.0 and above) Follow the steps in the patch `README.txt` to copy the ISO file and optional `preconf.csv` file into place.

- a. After download, unzip the file as `root` user for the following ISO image.

```
b2d6d9ac61e2e673d05fba11ec65107d
compute_12.1.2.2.0_LINUX.X64_150917-1.x
86_64.iso
```

 **Note:**

- ISO + `preconf.csv` is supported only if the ISO image is mounted on an NFS repository in ILOM and the `preconf.csv` file is located in the same NFS directory. All other methods of mounting the ISO image in ILOM (SAMBAs URI and Java Remote Console) and burning ISO image to DVD will not work with the `preconf.csv` file. For instructions on how to mount the ISO image from a NFS central repository, see "[Configure a Host Storage Device Redirection Session](#)" in the *Oracle X5 Series Servers Administration Guide*
- Because of possible late echos of text to the console, be aware that the boot process will be waiting for input to complete the interactive `ipconf` by prompting for `Nameserver` :
- Due to a bug, `ipconf` will prompt you twice to enter the network settings.

- b. Proceed to step 15, and boot the system with the ISO image and follow installation prompts. The ISO can be made available with a `preconf.csv` following the steps in the Note box above, or as virtual media on the ILOM of the node to be re-imaged when a `preconf.csv` is not used.

Steps 11 to 14, which create the image files, can be skipped.

11. Download the `imagemaker` kit to prepare and create the files for the USB drive.

This kit is available publicly on the Oracle Software Delivery Cloud (edelivery.oracle.com) site under "Oracle Database Products -> Linux x86-64". There are separate ImageMaker files for database nodes and cells.

12. After download, unzip the first file, called something similar to `V75080-01.zip`, and then unzip the subsequent file as well. Perform the unzips as the root user. Then untar the file.

```
# unzip V75080-01.zip

# unzip cellImageMaker_12.1.2.1.1_LINUX.X64_150316.2-1.x86_64.zip

# tar -pxvf
cellImageMaker_12.1.2.1.1_LINUX.X64_150316.2-1.x86_64.tar
```

View the `README_FOR_FACTORY.txt` file in the resulting directory for additional details.

13. Run `makeImageMedia.sh` to build the ISO file for imaging.

For building images for 12.1.2.1.0 and above, the recommended `makeImageMedia.sh` options for building ISO images are:

- a. (12.1.2.1.0 up to 12.2.1.1.0) Change directory to the `dl180` or `dl360` directory.
- b. (12.2.1.1.0 and above) Change directory to the `cell` or `compute` directory.
- c. Make the ISO file.

```
./makeImageMedia.sh -factory -stit -reboot-on-success -
nodisktests [-preconf
preconf_file] filename.iso
```

For ISO imaging, note that you must run `makeImageMedia.sh` from an Oracle Exadata database server if you are using the `-preconf` option.

For example:

```
# ./makeImageMedia.sh -factory -stit -reboot-on-success -
nodisktests -preconf
preconf.csv computenode.iso
[Info]: ipconf command line: /root/imagemaker/compute/initrd/opt/
oracle.cellos
/ipconf.pl -preconf preconf.csv -verify -nocodes
[Done]: Pre config verification OK
Please wait. Calculating md5 checksums for cellbits ...
Calculating md5 checksum for exaos.tbz ...
Calculating md5 checksum for dbboot.tbz ...
Calculating md5 checksum for dbfw.tbz ...
Calculating md5 checksum for kernel.tbz ...
Calculating md5 checksum for ofed.tbz ...
Calculating md5 checksum for sunutils.tbz ...
Calculating md5 checksum for commonos.tbz ...
Calculating md5 checksum for debugos.tbz ...
Calculating md5 checksum for dbrpms.tbz ...
Calculating md5 checksum for exaosovs.tbz ...
Calculating md5 checksum for dbovsboot.tbz ...
Calculating md5 checksum for kernelovs.tbz ...
Calculating md5 checksum for ofedovs.tbz ...
Calculating md5 checksum for commonovsos.tbz ...
Calculating md5 checksum for debugosovs.tbz ...
```

```
Calculating md5 checksum for dbovsrpms.tbz ...
Calculating md5 checksum for sunovsutils.tbz ...
Calculating md5 checksum for dbvmboot.tbz ...
Calculating md5 checksum for ofedvm.tbz ...
Calculating md5 checksum for sunvmutils.tbz ...
Please wait. Making initrd ...
395864 blocks
Please wait. Calculating md5 checksums for boot ...
Size of boot image is 4 sectors -> No emulation
0.30% done, estimate finish Wed Nov 29 10:47:27 2016
0.60% done, estimate finish Wed Nov 29 10:47:27 2016
0.90% done, estimate finish Wed Nov 29 10:47:27 2016
1.21% done, estimate finish Wed Nov 29 10:47:27 2016
1.51% done, estimate finish Wed Nov 29 10:47:27 2016
1.81% done, estimate finish Wed Nov 29 10:47:27 2016
2.11% done, estimate finish Wed Nov 29 10:47:27 2016
2.41% done, estimate finish Wed Nov 29 10:47:27 2016
2.71% done, estimate finish Wed Nov 29 10:47:27 2016
3.02% done, estimate finish Wed Nov 29 10:47:27 2016
3.32% done, estimate finish Wed Nov 29 10:47:27 2016
3.62% done, estimate finish Wed Nov 29 10:47:27 2016
3.92% done, estimate finish Wed Nov 29 10:47:27 2016
4.22% done, estimate finish Wed Nov 29 10:47:27 2016
4.52% done, estimate finish Wed Nov 29 10:47:27 2016
4.83% done, estimate finish Wed Nov 29 10:47:27 2016
5.13% done, estimate finish Wed Nov 29 10:47:27 2016
5.43% done, estimate finish Wed Nov 29 10:47:27 2016
5.73% done, estimate finish Wed Nov 29 10:47:27 2016
<output truncated>
97.72% done, estimate finish Wed Nov 29 10:47:30 2016
98.02% done, estimate finish Wed Nov 29 10:47:30 2016
98.32% done, estimate finish Wed Nov 29 10:47:30 2016
98.62% done, estimate finish Wed Nov 29 10:47:30 2016
98.92% done, estimate finish Wed Nov 29 10:47:30 2016
99.22% done, estimate finish Wed Nov 29 10:47:30 2016
99.52% done, estimate finish Wed Nov 29 10:47:30 2016
99.83% done, estimate finish Wed Nov 29 10:47:30 2016
Total translation table size: 2048
Total rockridge attributes bytes: 0
Total directory bytes: 2048
Path table size(bytes): 26
Max brk space used 0
1657890 extents written (3238 MB)
Done creation of installation ISO for compute node
```

14. Attach the ISO as virtual media to the ILOM of the database server to be re-imaged, and select CDROM as the next boot device.
15. Reboot the database node or cell to start the imaging process from the ISO.
16. If you are using Oracle VM, run `switch_to_ovm.sh` on each database server. The database servers will reboot.

For example:

```
# /opt/oracle.SupportTools/switch_to_ovm.sh
2014-12-07 11:58:36 -0800 [INFO] Switch to DOM0 system
partition /dev/VGExaDb/LVDb
Sys3 (/dev/mapper/VGExaDb-LVDbSys3)
2014-12-07 11:58:36 -0800 [INFO] Active system device: /dev/mapper/
VGExaDb-LVDbSys
1
2014-12-07 11:58:36 -0800 [INFO] Active system device in boot
area: /dev/mapper/VG
ExaDb-LVDbSys1
2014-12-07 11:58:36 -0800 [INFO] Set active system device to /dev/
VGExaDb/LVDbSys3
in /boot/I_am_hd_boot
2014-12-07 11:58:36 -0800 [INFO] Reboot has been initiated to
switch to the DOM0 s
ystem partition
```

17. Run `reclaimdisks.sh` on each database server (except KVM environments on X8M servers) before installing any software.

The `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim` command reclaims disk space reserved for the deployment type not selected. The command takes less than 5 minutes, approximately. Systems are imaged with disks configured with RAID5; a RAID rebuild is no longer part of the `reclaimdisks.sh` process.

Caution:

Do not skip this step. Skipping this step results in unused space that can no longer be reclaimed by `reclaimdisks.sh`.

```
# /opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim
Model is ORACLE SERVER X5-2
Number of LSI controllers: 1
Physical disks found: 4 (252:0 252:1 252:2 252:3)
Logical drives found: 1
...
[INFO      ] Copying /usr/share/grub/x86_64-redhat/* to /boot/
grub ...
[INFO      ] Create filesystem on device /dev/sda1
[INFO      ] Tune filesystem on device /dev/sda1

GNU GRUB  version 0.97  (640K lower / 3072K upper memory)

[ Minimal BASH-like line editing is supported.  For the first word,
TAB
lists possible command completions.  Anywhere else TAB lists the
possible
completions of a device/filename.]
grub> root (hd0,0)
Filesystem type is ext2fs, partition type 0x83
```

```
grub> setup (hd0)
  Checking if "/boot/grub/stage1" exists... no
  Checking if "/grub/stage1" exists... yes
  Checking if "/grub/stage2" exists... yes
  Checking if "/grub/e2fs_stage1_5" exists... yes
  Running "embed /grub/e2fs_stage1_5 (hd0)"... failed (this is not
fatal)
  Running "embed /grub/e2fs_stage1_5 (hd0,0)"... failed (this is not
fatal)
  Running "install /grub/stage1 (hd0) /grub/stage2 p /grub/grub.conf
"... succeeded
Done.
grub> quit
```

18. Run `ibhosts` and verify that all nodes show the correct IP addresses and host names. There should be no nodes with `elasticNode` in the description.
19. Run the OEDA tool to deploy.

```
[root] # ./install.sh -cf ../MAA-dm01.xml -l
1. Validate Configuration File
2. Setup Capacity-on-Demand
3. Create Virtual Machine
4. Create Users
5. Setup Cell Connectivity
6. Create Cell Disks
7. Create Grid Disks
8. Configure Alerting
9. Install Cluster Software
10. Initialize Cluster Software
11. Install Database Software
12. Relink Database with RDS
13. Create ASM Diskgroups
14. Create Databases
15. Apply Security Fixes
16. Install Exachk
17. Setup ASR Alerting
18. Create Installation Summary
19. Resecure Machine
```

Related Topics

- [Oracle Software Delivery Cloud](#)
- [Exadata 12.1.2.2.0 release and patch \(20131726\) \(My Oracle Support Doc ID 2038073.1\)](#)
- [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID 888828.1\)](#)

6.12.4 Using PXE to Image a New System

You can use PXE to prepare a new Exadata system for deployment.

- [Using PXE to Image a New System Using Release 12.1.2.2.0 or Later](#)
You can use PXE to prepare a new Exadata system running Oracle Exadata System Software release 12.1.2.2.0 or later for deployment using a simplified method.
- [Using PXE to Image a New System for Releases Earlier than 12.1.2.2.0](#)
You can use PXE to prepare a new Exadata system running an Oracle Exadata System Software release earlier than 12.1.2.2.0 for deployment.

6.12.4.1 Using PXE to Image a New System Using Release 12.1.2.2.0 or Later

You can use PXE to prepare a new Exadata system running Oracle Exadata System Software release 12.1.2.2.0 or later for deployment using a simplified method.

1. Obtain the latest Oracle Exadata Deployment Assistant (OEDA) from <http://www.oracle.com/technetwork/database/exadata/oeda-download-2076737.html>.
2. Run the OEDA configuration tool to generate the configuration files.
3. Power on all the servers.

When the system is first started, the following conditions exist:

- a. Database nodes and cells will automatically assign to themselves unused IP addresses on eth0 in the 172.16 range.
 - b. Host names of the nodes will be of the format `nodeN`, where *N* is a number.
4. Run OEDA to generate the `preconf.csv` file that you will use to image the system.

5. Connect to the first database node.

You can do this by connecting a laptop to the ILOM serial management port (9600, 8, N, 1). Then you can connect to the console and log in to the database node.

6. Obtain the eth0 MAC address for each database node.

The eth0 MAC address is the `fru_macaddress` field from `/SYS/MB/NET0` in the ILOM (or from `ip addr` if logged into the console).

- a. Connect to the each database node through SSH from the console.
 - b. Run `ibhosts` to get a list of node names and IP addresses.
7. Insert the MAC addresses using uppercase characters in the seventh field of the `preconf.csv` file for each node.

This field should be empty, that is, you should see two consecutive comma characters (`,,`).

For example:

```
orhb42-  
b08-01,example.com,db,eth0,eth0,Management,00:10:E0:69:34:52,10.196.  
3.  
106,255.255.224.0,10.196.31.250,orhb42-b08-01-priv,Private:active-  
bond-ib,192.1  
68.10.15,255.255.255.0,,,America/Los_Angeles
```

8. Create a PXE configuration file for each node.

The file name is `hostname.xx-xx-xx-xx-xx-xx` in the `pxelinux.cfg` directory. Below are the recommended options for PXE imaging of versions 12.1.2.1.0 and above. The line starting with `append` should be entered as a single line with no line breaks.

```
kernel vmlinux-iso-file-name
append initrd=img_file pxe factory stit reboot-on-success
notests=diskgroup dhcp
preconf=n.n.n.n:/directory/preconf.csv iso-uri=http://web_server/
x86_64/iso/image_n
ame.iso console=ttyS0,115200n8
```

9. Optionally, configure key-based authentication for the `root` OS user.

Starting with Oracle Exadata System Software release 20.1.0, you can configure key-based authentication for the `root` OS user as part of the PXE imaging process.

If you want to use this option, then perform the following for each node:

- a.** Create an RSA public and private key pair using a utility such as `ssh-keygen`.

For example:

```
# ssh-keygen -t rsa -b 4096
```

- b.** Rename the public key file so that the file name conforms to the following format:

```
id_rsa.short_hostname.root.pub
```

In the public key file name, *short_hostname* is the node host name as defined in the first field of the corresponding node entry in the `preconf.csv` file.

- c.** Copy the renamed public key file to the directory that contains the `preconf.csv` file.

This is the same directory that you previously specified in the PXE configuration file `preconf` parameter, which defines the path to the `preconf.csv` file.

If you configure key-based authentication, then password-based authentication is disabled after the imaging process. Consequently, `root` access to the node is only possible using the matching private key. ILOM access using the `root` password is also disabled.

By default, if the public key file is not named or located correctly, then password-based authentication is enabled for the node. You can also make key-based authentication mandatory by specifying `root-ssh-key=yes` as a boot option in the PXE configuration file. In this case, the imaging process fails if the public key is not in the correct location.

Regardless of any public keys, you can disable the configuration of key-based authentication by specifying `root-ssh-key=no` as a boot option in the PXE configuration file.

10. Create the image files, if necessary.

If you are using version 12.1.2.2.0 or later, it is no longer necessary to use ImageMaker to create the image files. The relevant PXE, ISO, and USB image files have already been created in the patches themselves. Simply download the patch containing the pre-built PXE image files from Oracle Software Delivery Cloud and place them into the appropriate directories on the PXE server. For all releases, see My Oracle Support note 888828.1 for the patch numbers. There are separate PXE image patches for database servers and storage servers.

Note: For image version 12.1.2.2.0 the format of the downloadable image files has changed.

The downloaded zip file contains four image files:

- For 12.1.2.2.0 up to 12.2.1.1.0:

```
kernel: vmlinux-version-ImageName-{DL180|DL360}
initrd: initrd-version-ImageName-{DL180|DL360}.img
image: nfsimg-version-ImageName-{DL180|DL360}.tar
image: nfsimg-version-ImageName-{DL180|DL360}.tar.md5
```

- For 12.2.1.1.0 and above:

```
kernel: vmlinux-iso-ImageName-{cell|compute}
initrd: initrd-iso-ImageName-{cell|compute}.img
image: vmlinux-iso-ImageName-{cell|compute}.tar
```

11. Change the boot order using the ILOM to boot from PXE and then restart the server to start the imaging process.

For example, from the Oracle Linux prompt on the host:

```
# ipmitool chassis bootdev pxe

# shutdown -r now
```

12. If you are using Oracle VM, run `switch_to_ovm.sh` on all Oracle VM database servers before running the OEDA installation tool.

When you run this command, the database servers will reboot.

```
# /opt/oracle.SupportTools/switch_to_ovm.sh
2014-12-07 11:58:36 -0800 [INFO] Switch to DOM0 system
partition /dev/VGExaDb/LV
DbSys3 (/dev/mapper/VGExaDb-LVDbSys3)
2014-12-07 11:58:36 -0800 [INFO] Active system device: /dev/mapper/
VGExaDb-LVDbS
ysl
2014-12-07 11:58:36 -0800 [INFO] Active system device in boot
area: /dev/mapper/
VGExaDb-LVDbSys1
2014-12-07 11:58:36 -0800 [INFO] Set active system device to /dev/
VGExaDb/LVDbSy
s3 in /boot/I_am_hd_boot
2014-12-07 11:58:36 -0800 [INFO] Reboot has been initiated to
switch to the DOM0
system partition
```


13. Run `reclaimdisks.sh` on each database server (except KVM environments on X8M servers) before installing any software.

The `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim` command reclaims disk space reserved for the deployment type not selected. The command takes less than 5 minutes, approximately. Systems are imaged with disks configured with RAID5; a RAID rebuild is no longer part of the `reclaimdisks.sh` process.

▲ Caution:

Do not skip this step. Skipping this step results in unused space that can no longer be reclaimed by `reclaimdisks.sh`.

```
# /opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim
Model is ORACLE SERVER X5-2
Number of LSI controllers: 1
Physical disks found: 4 (252:0 252:1 252:2 252:3)
Logical drives found: 1
...
[INFO      ] Copying /usr/share/grub/x86_64-redhat/* to /boot/
grub ...
[INFO      ] Create filesystem on device /dev/sda1
[INFO      ] Tune filesystem on device /dev/sda1

GNU GRUB  version 0.97  (640K lower / 3072K upper memory)

[ Minimal BASH-like line editing is supported.  For the first word,
TAB
lists possible command completions.  Anywhere else TAB lists the
possible
completions of a device/filename.]
grub> root (hd0,0)
  Filesystem type is ext2fs, partition type 0x83
grub> setup (hd0)
  Checking if "/boot/grub/stage1" exists... no
  Checking if "/grub/stage1" exists... yes
  Checking if "/grub/stage2" exists... yes
  Checking if "/grub/e2fs_stage1_5" exists... yes
  Running "embed /grub/e2fs_stage1_5 (hd0)"... failed (this is not
fatal)
  Running "embed /grub/e2fs_stage1_5 (hd0,0)"... failed (this is not
fatal)
  Running "install /grub/stage1 (hd0) /grub/stage2 p /grub/grub.conf
"... succeeded
Done.
grub> quit
```

14. Run `ibhosts` and verify that all nodes show the correct IP addresses and host names.

There should be no nodes with `elasticNode` in the description.

15. Run the OEDA tool to deploy.

 **Note:**

If you configured key-based authentication for the `root` OS user, then you must first copy the node private keys into the OEDA WorkDir.

```
[root] # ./install.sh -cf ../machine_config_file.xml -l
1. Validate Configuration File
2. Setup Capacity-on-Demand
3. Create Virtual Machine
4. Create Users
5. Setup Cell Connectivity
6. Create Cell Disks
7. Create Grid Disks
8. Configure Alerting
9. Install Cluster Software
10. Initialize Cluster Software
11. Install Database Software
12. Relink Database with RDS
13. Create ASM Diskgroups
14. Create Databases
15. Apply Security Fixes
16. Install Exachk
17. Setup ASR Alerting
18. Create Installation Summary
19. Resecure Machine
```

Related Topics

- [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID 888828.1\)](#)

6.12.4.2 Using PXE to Image a New System for Releases Earlier than 12.1.2.2.0

You can use PXE to prepare a new Exadata system running an Oracle Exadata System Software release earlier than 12.1.2.2.0 for deployment.

1. Obtain the latest Oracle Exadata Deployment Assistant (OEDA) from <http://www.oracle.com/technetwork/database/exadata/oeda-download-2076737.html>.
2. Run the OEDA configuration tool to generate the configuration files.
3. Power on all the servers.

When the system is first started, the following conditions exist:

- a. Database nodes and cells will automatically assign to themselves unused IP addresses on `eth0` in the 172.16 range.
 - b. Host names of the nodes will be of the format `nodeN`, where *N* is a number.
4. Run OEDA to generate the `preconf.csv` file that you will use to image the system.
 5. Connect to the first database node.

You can do this by connecting a laptop to the ILOM serial management port (9600, 8, N, 1). Then you can connect to the console and log in to the database node.

6. Obtain the eth0 MAC address for each database node.

The eth0 MAC address is the `fru_macaddress` field from `/SYS/MB/NET0` in the ILOM (or from `ip addr` if logged into the console).

- a. Connect to the each database node through SSH from the console.
 - b. Run `ibhosts` to get a list of node names and IP addresses.
7. Insert the MAC addresses using uppercase characters in the seventh field of the `preconf.csv` file for each node.

This field should be empty, that is, you should see two consecutive comma characters (, ,).

For example:

```
orhb42-
b08-01,example.com,db,eth0,eth0,Management,00:10:E0:69:34:52,10.196.
3.
106,255.255.224.0,10.196.31.250,orhb42-b08-01-priv,Private:active-
bond-ib,192.1
68.10.15,255.255.255.0,,,America/Los_Angeles
```

8. Create a PXE configuration file for each node.

Below are the recommended options for PXE imaging of versions prior to release 12.1.2.2.0. The line starting with `append` should be entered as a single line with no line breaks.

```
append initrd=img_file pxe factory stit reboot-on-success
notests=diskgroup dhcp
preconf=n.n.n.n:/directory/preconf.csv sk=n.n.n.n:/directory
console=ttyS0,
115200n8
```

9. Download the ImageMaker kit to prepare and create the files for the PXE server.

The kit is available publicly on the Oracle Software Delivery Cloud (`edelivery`) site under “Oracle Database Products -> Linux x86-64”. There are separate ImageMaker files for database nodes and cells.

10. After downloading the kit, as the `root` user, unzip the first file that has a name similar to `V75080-01.zip`, and then unzip the subsequent ZIP file. Finally, untar the file.

```
# unzip V75080-01.zip

# unzip cellImageMaker_12.1.2.1.1_LINUX.X64_150316.2-1.x86_64.zip

# tar -pxvf
cellImageMaker_12.1.2.1.1_LINUX.X64_150316.2-1.x86_64.tar
```

View the `README_FOR_FACTORY.txt` file in the resulting directory for additional details.

11. Run `makeImageMedia.sh` to build the kernel, `initrd`, and image files, if necessary.

Place the image files in the relevant directories on the PXE server. These files will be sent at boot time to the node being imaged.

- Change directory to the `dl180` or `dl360` directory and make the PXE build using:

```
./makeImageMedia.sh -pxe -pxeout ImageName
```

If the `dl180` or `dl360` directory are not available, then change directory to the `cell` or `compute` directory.

The `makeImageMedia.sh` command creates four files in the `./PXE` subdirectory:

```
kernel: vmlinux-version-ImageName-{DL180|DL360}
initrd: initrd-version-ImageName-{DL180|DL360}.img
image: nfsimg-version-ImageName-{DL180|DL360}.tar
image: nfsimg-version-ImageName-{DL180|DL360}.tar.md5
```

12. Change the boot order using the ILOM to boot from PXE and then restart the server to start the imaging process.

For example, from the Oracle Linux prompt on the host:

```
# ipmitool chassis bootdev pxe
```

```
# shutdown -r now
```

13. If you are using Oracle VM run `switch_to_ovm.sh` on all Oracle VM database servers before running the OEDA installation tool.

When you run this command, the database servers will reboot.

```
# /opt/oracle.SupportTools/switch_to_ovm.sh
2014-12-07 11:58:36 -0800 [INFO] Switch to DOM0 system
partition /dev/VGExaDb/LV
DbSys3 (/dev/mapper/VGExaDb-LVDbSys3)
2014-12-07 11:58:36 -0800 [INFO] Active system device: /dev/mapper/
VGExaDb-LVDbS
ysl
2014-12-07 11:58:36 -0800 [INFO] Active system device in boot
area: /dev/mapper/
VGExaDb-LVDbSys1
2014-12-07 11:58:36 -0800 [INFO] Set active system device to /dev/
VGExaDb/LVDbSy
s3 in /boot/I_am_hd_boot
2014-12-07 11:58:36 -0800 [INFO] Reboot has been initiated to
switch to the DOM0
system partition
```

14. Run `reclaimdisks.sh` on each database server (except KVM environments on X8M servers) before installing any software.

The `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim` command reclaims disk space reserved for the deployment type not selected. The command

takes less than 5 minutes, approximately. Systems are imaged with disks configured with RAID5; a RAID rebuild is no longer part of the `reclaimdisks.sh` process.

▲ Caution:

Do not skip this step. Skipping this step results in unused space that can no longer be reclaimed by `reclaimdisks.sh`.

```
# /opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim
Model is ORACLE SERVER X5-2
Number of LSI controllers: 1
Physical disks found: 4 (252:0 252:1 252:2 252:3)
Logical drives found: 1
...
[INFO      ] Copying /usr/share/grub/x86_64-redhat/* to /boot/
grub ...
[INFO      ] Create filesystem on device /dev/sda1
[INFO      ] Tune filesystem on device /dev/sda1

GNU GRUB  version 0.97  (640K lower / 3072K upper memory)

[ Minimal BASH-like line editing is supported.  For the first word,
TAB
lists possible command completions.  Anywhere else TAB lists the
possible
completions of a device/filename.]
grub> root (hd0,0)
  Filesystem type is ext2fs, partition type 0x83
grub> setup (hd0)
  Checking if "/boot/grub/stage1" exists... no
  Checking if "/grub/stage1" exists... yes
  Checking if "/grub/stage2" exists... yes
  Checking if "/grub/e2fs_stage1_5" exists... yes
  Running "embed /grub/e2fs_stage1_5 (hd0)"... failed (this is not
fatal)
  Running "embed /grub/e2fs_stage1_5 (hd0,0)"... failed (this is not
fatal)
  Running "install /grub/stage1 (hd0) /grub/stage2 p /grub/grub.conf
"... succeeded
Done.
grub> quit
```

15. Run `ibhosts` and verify that all nodes show the correct IP addresses and host names.

There should be no nodes with `elasticNode` in the description.

16. Run the OEDA tool to deploy.

```
[root] # ./install.sh -cf ../machine_config_file.xml -l
1. Validate Configuration File
2. Setup Capacity-on-Demand
```

3. Create Virtual Machine
4. Create Users
5. Setup Cell Connectivity
6. Create Cell Disks
7. Create Grid Disks
8. Configure Alerting
9. Install Cluster Software
10. Initialize Cluster Software
11. Install Database Software
12. Relink Database with RDS
13. Create ASM Diskgroups
14. Create Databases
15. Apply Security Fixes
16. Install Exachk
17. Setup ASR Alerting
18. Create Installation Summary
19. Resecure Machine

Related Topics

- [Oracle Software Delivery Cloud](#)

6.12.5 Using USB to Image a New System

You can use USB to prepare a new Exadata system for deployment.

1. Obtain the latest Oracle Exadata Deployment Assistant (OEDA).
2. Run OEDA to generate the `preconf.csv` file that you can use to image the system.
3. Power on all the servers.
 - a. Database nodes and cells will automatically assign to themselves unused IP addresses on `eth0` in the 172.16 range.
 - b. Host names of the nodes will be of the format `nodeN`, where `N` is a number.
4. Connect to the first database node. You can do this by connecting a laptop to the ILOM serial management port (9600, 8, N, 1), and then connecting to the console and logging in to the database node.
5. Obtain the `eth0` MAC address for each database node. This is the `fru_macaddress` field from `/SYS/MB/NET0` in the ILOM, or use the command `ip addr` if logged into the console.
 - a. Connect to the other database nodes using SSH from the console.
 - b. Run `ibhosts` to get a list of node names and IP addresses.
6. Insert the MAC addresses using uppercase characters in the seventh field of the `preconf.csv` file for each node. This field should be empty, that is, you should see two consecutive comma characters (, ,). For example, if the MAC address is 00:10:E0:69:34:52:

```
orhb42-b08-01,example.com,db,eth0,eth0,Management,00:10:E0:69:34:52,  
10.196.3.106,255.255.224.0,10.196.31.250,orhb42-b08-01-  
priv,Private:active-bond-ib,  
192.168.10.15,255.255.255.0,,,America/Los_Angeles
```

7. Run `ipconf` to verify the generated `preconf.csv` file. This can be done on any Oracle Exadata server, and although not required, it is useful to ensure the `preconf.csv` file is valid before proceeding.

```
# ipconf -verify -preconf /tmp/preconf.csv
[Info]: ipconf command line: /opt/oracle.cellos/ipconf.pl -verify -
preconf /tmp/preconf -nocodes
[Done]: Pre config verification OK
```

8. For versions prior to 12.1.2.2.0 proceed to step 11 to use ImageMaker to create the image files.

If you are using version 12.1.2.2.0 or later, it is no longer necessary to use `imagemaker` to create the image files in version 12.1.2.2.0. The relevant PXE, ISO, and USB image files have already been created in the patches themselves. Simply download the patch containing the prebuilt image file from Oracle Software Delivery Cloud (edelivery) and place the downloaded USB `.img` image file onto the USB drive. For 12.1.2.2.0 the patch numbers can be obtained from My Oracle Support note 2038073.1. There are separate ImageMaker files for database nodes and cells.

Note: For image version 12.1.2.2.0 the format of the downloadable image files has changed.

There will be an image file in the downloaded zip file. For example, the database node zip file will contain `computeImageMaker_12.1.2.2.0_LINUX.X64_150917.x86_64.img`. For cell images, the file will be similar to `cellImageMaker_12.1.2.2.0_LINUX.X64_150917.x86_64.img`.

9. (For 12.1.2.2.0 and above) Follow the steps in the patch `README.txt` file to copy the file and optional `preconf.csv` file onto the USB drive.

- a. Place the `computeImageMaker_12.1.2.2.0_LINUX.X64_150917.x86_64.img` file in a directory on an Oracle Linux server.
- b. Verify that the `md5sums` match: `f3daae6a9757d4feb4a0f4262e610945`
`computeImageMaker_12.1.2.2.0_LINUX.X64_150917.x86_64.img`
- c. Prepare the USB drive using a command similar to the following, where `/dev/sdd` is the name of the inserted USB drive:

```
# dd if=/dev/zero of=/dev/sdd bs=1M count=100 oflag=direct
```

You can determine the exact name of the USB drive by looking in `/var/log/messages` after inserting the USB drive.

- d. Write the `.img` file to the USB drive. This may take 15 minutes or more and no output is shown during the operation.

```
# dd if=filename.img of=/dev/sdd bs=1M oflag=direct
```

- e. Rescan the partition table on Oracle Linux to recognize the new partition:

```
# partprobe
```

- f. Verify that the USB is mountable as long as the system supports ext4 file system. This must be an Oracle Linux 6 system for this step even if the previous `dd` command was run on an Oracle Linux 5 system.

```
mount /dev/sdd1 /mnt
```

- g. (Optional) Prepare and place the `preconf.csv` file on the USB drive. The file name must be `preconf.csv` on the USB drive. The `preconf.csv` file must contain MAC addresses for each node in the order to be used during the image. If no `preconf.csv` file is used during imaging, the node will prompt for its network configuration the first time it boots up.

```
# cp /path/preconf.csv /mnt/preconf.csv
```

```
# umount /mnt
```

The `umount` command above ensures that the file system is synchronized.

- h. Proceed to step 15. Steps 11 to 14, which create the image files, can be skipped.
10. Download the ImageMaker kit to prepare and create the files for the USB drive. This kit is available publicly on the Oracle Software Delivery Cloud (edelivery) site under "Oracle Database Products -> Linux x86-64". There are separate ImageMaker files for database nodes and cells.
11. After downloading ImageMaker, unzip the first file, which is named similar to `V75080-01.zip`, and then also unzip the subsequent file. Perform the unzip as the `root` user. Then untar the file.

```
# unzip V75080-01.zip
```

```
# unzip cellImageMaker_12.1.2.1.1_LINUX.X64_150316.2-1.x86_64.zip
```

```
# tar -pxvf  
cellImageMaker_12.1.2.1.1_LINUX.X64_150316.2-1.x86_64.tar
```

View the `README_FOR_FACTORY.txt` file in the resulting directory for additional details

12. Insert a blank USB device in the slot of the server where ImageMaker was unpacked.
13. Run `makeImageMedia.sh` to build the kernel, `initrd`, and image files for the USB device. `makeImageMedia.sh` will prompt to confirm creation of the image on the USB device. For building 12.1.2.1.0 and above images, the recommended `makeImageMedia.sh` options for building USB bootable media are:
 - a. (For 12.1.2.1.0 up to 12.2.1.1.0) Change directory to the `d1180` or `d1360` directory.
 - b. (For 12.2.1.1.0 and above) Change directory to the `cell` or `compute` directory.

c. Make files for the USB.

```
./makeImageMedia.sh -factory -stit -reboot-on-success -  
nodisktests [-preconf path_to_preconf_file]
```

For USB imaging, note the following:

- `makeImageMedia.sh` must be run from an Oracle Exadata database server if using the `-preconf` option.
 - Storage servers and database servers can hang during boot when attempting a reimage using a bootable USB created with the `makeImageMedia.sh` script provided with storage server and database server host images available on the Oracle Software Delivery Cloud (edelivery). To work around the issue, see My Oracle Support note 1919001.1
14. After the USB drive has been prepared, place it into the slot of the server to be imaged. Ensure the image type matches the server type (database node or cell).
 15. Reboot the database node or cell to start the imaging process from USB.
 16. If you are using Oracle VM, run `switch_to_ovm.sh` on each database server. The database servers will reboot.

For example:

```
# /opt/oracle.SupportTools/switch_to_ovm.sh  
2014-12-07 11:58:36 -0800 [INFO] Switch to DOM0 system  
partition /dev/VGExaDb/LVDbSys3 (/dev/mapper/VGExaDb-LVDbSys3)  
2014-12-07 11:58:36 -0800 [INFO] Active system device: /dev/mapper/  
VGExaDb-LVDbSys1  
2014-12-07 11:58:36 -0800 [INFO] Active system device in boot  
area: /dev/mapper/VGExaDb-LVDbSys1  
2014-12-07 11:58:36 -0800 [INFO] Set active system device to /dev/  
VGExaDb/LVDbSys3 in /boot/I_am_hd_boot  
2014-12-07 11:58:36 -0800 [INFO] Reboot has been initiated to  
switch to the DOM0 system partition
```

17. Run `reclaimdisks.sh` on each database server (except KVM environments on X8M servers) before installing any software.

The `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim` command reclaims disk space reserved for the deployment type not selected. The command takes less than 5 minutes, approximately. Systems are imaged with disks configured with RAID5; a RAID rebuild is no longer part of the `reclaimdisks.sh` process.

 **Caution:**

Do not skip this step. Skipping this step results in unused space that can no longer be reclaimed by `reclaimdisks.sh`.

```
# /opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim  
Model is ORACLE SERVER X5-2  
Number of LSI controllers: 1
```

```
Physical disks found: 4 (252:0 252:1 252:2 252:3)
Logical drives found: 1
...
[INFO      ] Copying /usr/share/grub/x86_64-redhat/* to /boot/
grub ...
[INFO      ] Create filesystem on device /dev/sda1
[INFO      ] Tune filesystem on device /dev/sda1

GNU GRUB  version 0.97  (640K lower / 3072K upper memory)

[ Minimal BASH-like line editing is supported.  For the first word,
TAB
lists possible command completions.  Anywhere else TAB lists the
possible
completions of a device/filename.]
grub> root (hd0,0)
  Filesystem type is ext2fs, partition type 0x83
grub> setup (hd0)
  Checking if "/boot/grub/stage1" exists... no
  Checking if "/grub/stage1" exists... yes
  Checking if "/grub/stage2" exists... yes
  Checking if "/grub/e2fs_stage1_5" exists... yes
  Running "embed /grub/e2fs_stage1_5 (hd0)"... failed (this is not
fatal)
  Running "embed /grub/e2fs_stage1_5 (hd0,0)"... failed (this is not
fatal)
  Running "install /grub/stage1 (hd0) /grub/stage2 p /grub/grub.conf
"... succeeded
Done.
grub> quit
```

18. Run `ibhosts` and verify that all nodes show the correct IP addresses and host names. There should be no nodes with "elasticNode" in the description.
19. Run the OEDA tool to deploy.

```
[root] # ./install.sh -cf ../MAA-dm01.xml -l
1. Validate Configuration File
2. Setup Capacity-on-Demand
3. Create Virtual Machine
4. Create Users
5. Setup Cell Connectivity
6. Create Cell Disks
7. Create Grid Disks
8. Configure Alerting
9. Install Cluster Software
10. Initialize Cluster Software
11. Install Database Software
12. Relink Database with RDS
13. Create ASM Diskgroups
14. Create Databases
15. Apply Security Fixes
16. Install Exachk
17. Setup ASR Alerting
```

18. Create Installation Summary
19. Resecure Machine

Related Topics

- [Oracle Software Delivery Cloud](#)
- [Exadata 12.1.2.2.0 release and patch \(20131726\) \(My Oracle Support Doc ID 2038073.1\)](#)
- [Exadata reimaging using USB fails to boot USB \(My Oracle Support Doc ID 1919001.1\)](#)
- [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID 888828.1\)](#)

6.12.6 Updates to Imaging Parameters

Some of the command line options for imaging have changed with 12.1.2.1.0 and above.

The changes are:

- `updfirm`
Do NOT use the `updfirm` option in any form. It will be removed from the imaging options in the future. Allow the firmware to be updated automatically upon initial boot instead.
- `dualboot`
`dualboot` is no longer required with 12.1.2.1.0 as there is no support for Solaris. It has no meaning in releases 12.1.2.1.0 and above.
- `factory`
The meaning of the `factory` keyword has changed as of 12.1.2.1.0. The use of "factory" now enforces additional options such as "ovs=yes". This should be used for all imaging, regardless of the ultimate system configuration (Oracle VM, Physical etc.)
- `OVS=yes|no`
This option defaults to "no". However it is **not recommended** to use this at all because it may change in the future. In its place, use the "factory" option which will override some default options.

There are some additional command line options which may be useful:

- To create a hot spare on dbnodes, use the `enable_hot_spare` option.
 - `makeImageMedia.sh` option: `-enable_hot_spare`
 - PXE option: `enable_hot_spare`
- To force active bond or no active bond for InfiniBand interfaces (mainly for expansion rack cases), use the `active-bond-ib` option.
 - `makeImageMedia.sh` option: `-active-bond-ib <yes|no>`
 - PXE option: `active-bond-ib=<yes|no>`

The `README_FOR_FACTORY.txt` file shipped with the `imagemaker` software explains each option in more detail.

Related Topics

- [Exadata 12.1.2.2.0 release and patch \(20131726\) \(My Oracle Support Doc ID 2038073.1\)](#)
- [Reimaging Exadata Cell Node Guidance \(My Oracle Support Doc ID 2151671.1\)](#)

6.13 Performing Initial Elastic Configuration of Oracle Exadata Database Machine

This topic provides background information on elastic configuration and describes how to perform the initial configuration of Oracle Exadata Database Machine.

You need the configuration file generated by Oracle Exadata Deployment Assistant (OEDA).

Elastic configuration is new in Oracle Exadata System Software release 12.1.2.1.0. It applies to all rack configurations, such as the following:

- Racks ordered with a standard number of database nodes and cells, for example, quarter or half racks
- Rack configurations that feature additional database nodes and cells
- Oracle Exadata Database Machine Two-Socket X5 generation and later storage servers
- Oracle Exadata Racks that use an InfiniBand Network Layer or a RoCE Network Layer.

Starting with the Oracle Exadata Database Machine Two-Socket X5 generation of storage servers, the same process can also be used to add additional database servers or storage servers to an existing configuration.


The process allows initial IP addresses to be assigned to database servers and storage servers, regardless of the exact configuration you ordered. Your specific configuration can then be applied to the nodes. The procedure can be run from any database node (usually the first node,) and also from a storage server in cases where expansion involves adding only storage servers.

Every Oracle Exadata Database Machine has a predefined method for the cabling of nodes to the RDMA Network Fabric switch ports. Therefore, there is a fixed mapping from each node's location in the rack to the ports of the RDMA Network Fabric switches. Assuming the rack is always populated following this map, a node's rack unit location can be identified by querying the RDMA Network Fabric to determine the RDMA Network Fabric switch port the node is connected to. Once found, that information is used to determine the rack unit location of the node. With this information, nodes can be allocated initial IP addresses based on their rack unit location, with nodes lower in the rack getting lower IP addresses.

 **Note:**

- If a RAID rebuild/reconstruction (that is, a disk replacement or expansion) is in progress, wait until it is finished before performing the steps in this section. The reason is that the `reclaimdisks.sh` script is blocked from running until the RAID rebuild/reconstruction is complete.
- Disk space for unused system partitions must be reclaimed before installing software on all database servers (except KVM environments on X8M servers). Use the following script to reclaim the disk space: `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim`.

1. Open a console session to the first database server. The first database server is the lowest database server in the rack, which is rack position U16. In the case of an expansion involving only cells, the procedure should be run from a storage server instead of a database node.

 **Note:**

The network connection and configuration for Oracle Exadata Database Machine has not occurred yet.

2. Log in as the `root` user on the first database server.

 **Note:**

If you do not have the password for the `root` user, then contact Oracle Support Services.

3. (Not for X8M racks) Verify all servers are configured with the `172.16.*` IP address for the `eth0` interface using the following command:

```
# ibhosts
```

```
Ca : 0x0010e00001486fb8 ports 2 "node10 elasticNode  
172.16.2.46,172.16.2.46 ETH0"  
Ca : 0x0010e00001491228 ports 2 "node9 elasticNode  
172.16.2.45,172.16.2.45 ETH0"  
Ca : 0x0010e000014844f8 ports 2 "node8 elasticNode  
172.16.2.44,172.16.2.44 ETH0"  
Ca : 0x0010e00001488218 ports 2 "node4 elasticNode  
172.16.2.40,172.16.2.40 ETH0"  
Ca : 0x0010e000014908b8 ports 2 "node2 elasticNode  
172.16.2.38,172.16.2.38 ETH0"  
Ca : 0x0010e0000148ca68 ports 2 "node1 elasticNode  
172.16.2.37,172.16.2.37 ETH0"  
Ca : 0x0010e00001485fd8 ports 2 "node3 elasticNode  
172.16.2.39,172.16.2.39 ETH0"
```

Each server description should be `hostname elasticNode. ipaddress ETH0`.

4. (Optional on database servers) Switch to the Oracle VM when using virtualization on the rack using the following command:

 **Note:**

For this and following steps Oracle VM refers to OVM and KVM. Starting with Oracle Exadata Database Machine X8M-2, Oracle Linux KVM is the virtualization technology for systems that use RoCE Network Fabric. Otherwise, Oracle VM is a Xen-based virtualization technology, which is used across Oracle Exadata Database Machine systems that use InfiniBand Network Fabric.

The database servers will reboot during this step.

 **Note:**

It is not possible or supported to switch between physical and Oracle VM configuration.

```
# /opt/oracle.SupportTools/switch_to_ovm.sh
```

5. Reclaim the unused disk space on the database servers.

Disk space for unused system partitions must be reclaimed before installing software on all database servers (except KVM environments on X8M servers). Use the following script to reclaim the disk space: `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim`.

If Oracle VM was chosen, this step locks in the image as Oracle VM. Going back and forth will not be possible.

`reclaimdisks.sh` performs partition reconfiguration and a few other necessary configuration changes that take typically less than a minute to complete.

6. On database servers, verify the disk configuration.

```
# /opt/oracle.SupportTools/reclaimdisks.sh -check
```

If the system is a physical (bare metal) deployment, the last line of output will be the following:

```
Valid. Booted: Linux. Layout: Linux.
```

If the system is an Oracle VM deployment, the last line of output will be the following:

```
Valid. Booted: DOM0. Layout: DOM0.
```

If the system is an Oracle VM deployment (KVM) on X8M, the last line of output will be the following:

This utility `/opt/oracle.SupportTools/reclaimdisks.sh` is not applicable for KVM based Exadata database nodes

7. Copy the OEDA configuration files to the database server using one of the following methods. If adding only storage servers, the steps can be run from a storage server.

- Using a USB flash drive:
 - a. Copy the configuration file to a USB flash drive.
 - b. Connect the USB drive to the database server or storage server.
 - c. Locate the USB drive using the following command:

```
for x in `ls -l /sys/block`; do udevadm info --attribute-
walk --path=/sys/block/$x | grep \
-iq 'DRIVERS=="usb-storage"'; if [ $? -eq 0 ] ; then
echo /dev/${x}1; \
fi ; done
```

The expected output is similar to:

```
/dev/sdb1
```

Note:

If an Oracle Exadata Storage Server X6 or X5 Extreme Flash (EF) is used, the command will also return the two internal USB drives, usually `/dev/sda1` and `/dev/sdb1`. The newly inserted USB drive will be the additional device, usually `/dev/sdc1`.

The expected output is similar to:

```
/dev/sda1
/dev/sdb1
/dev/sdc1
```

The internal USBs can be identified by viewing `mdstat`:

```
# grep md4 /proc/mdstat | awk '{print $5,$6;}'
```


- d. Create a directory on the database server or storage server using the following command:

```
# mkdir /mnt/usb
```

- e. Create a directory for the installation using the following command:

```
# mkdir /opt/oracle.SupportTools/onecommand
```

If the directory already exists, then rename the existing one to `onecommand.old`, and then re-create the empty directory.

 **Note:**

Do not use a directory under `/u01` or `/EXAVMIMAGES` because `reclaimdisks.sh` will fail if there is content in those directories.

- f. Mount the device. Use the device name given in step 7.c. The following is an example of the command.

```
# mount -t vfat /dev/sdb1 /mnt/usb
```

- g. Copy OEDA and the configuration files from the USB flash drive to the `onecommand` directory on the database server or storage server using the following command:

```
# cp /mnt/usb/path_to_file/*.* /opt/oracle.SupportTools/onecommand
```

In the preceding command, *path_to_file* is the directory path to the configuration file on the USB flash drive.

- h. Unmount the USB flash drive using the following command:

```
# umount /mnt/usb
```

- i. Remove the USB flash drive from the database server or storage server.

- Using a temporary network connection:

This option is only needed if there were IP address conflicts, and Oracle Exadata Database Machine is not already connected to the network.

- a. Connect a network cable to the Ethernet port labeled NET0 on the back of the first database server.

- b. Bring up the `eth0:1` interface as follows:

```
# ifconfig eth0:1 real_ip netmask real_netmask up
```

In the preceding command, *real_ip* is the IP address to assign to the database server, and *real_netmask* is the subnet mask to assign to the database server.

- c. Add the default route as follows:

```
# route add -net 0.0.0.0 netmask 0.0.0.0 gw real_gateway
```


In the preceding command, *real_gateway* is the default gateway to use for the database server.

- d. Copy the configuration files to the `/opt/oracle.SupportTools/onecommand` directory on the database server or storage server using a file transfer utility such as `scp`.
- e. Disconnect the network cable from Ethernet port NET0.

 **Note:**

It is not necessary to reconfigure the network because the network is reconfigured automatically during the restart step.

8. (Not for X8M racks) Verify all servers are configured with the `172.16.*` IP address for the `eth0` interface using the following command:

```
# ibhosts
```

No server description should include `elasticNode` in its description.

9. Download and unzip the OEDA archive to a host on the same network as the servers in the rack (`172.16.*`).

The `applyElasticConfig.sh` script is not configured by default on the servers in the rack.

10. Run the `applyElasticConfig.sh` script.

For example, if you unzipped the OEDA ZIP file into the `/opt/oracle.SupportTools/onecommand/linux-x64` directory, then use the following commands:

```
# cd /opt/oracle.SupportTools/onecommand/linux-x64  
  
# ./applyElasticConfig.sh -cf customer_name-configFile.xml
```

The `applyElasticConfig.sh` script performs network configuration for all database servers and storage servers. All servers restart at the end of the process.

11. Connect the corporate network cable for the management network to the Cisco Management Switch.
12. Connect the corporate network cables for the client access network to the database servers.
13. Restart all database servers and storage servers.
14. Log in to the first database server (or storage server, if only storage servers were added for expansion) to verify network connectivity using the following commands:

```
# cd /opt/oracle.SupportTools/onecommand/linux-x64  
  
# ./checkip.sh -cf configuration_file
```

In the preceding command, *configuration_file* is the name of the configuration file from OEDA, and includes the full path.

 **Note:**

- The `checkip.sh` script can be run on a server outside of Oracle Exadata Database Machine to verify network connectivity to Oracle Exadata Database Machine. Copy the `checkip.sh` file to the external system, and run the command as shown in this step.
- If the command is run from a Microsoft Windows machine, then the command is `checkip.cmd`.
- Download the version of OEDA for the relevant platform.

15. Use OEDA to install and configure the software for Oracle Exadata Database Machine.

Related Topics

- [Cabling Information for Oracle Exadata Database Machine](#)
- [Powering On and Off Oracle Exadata Rack](#)
- [Adding Additional Elastic Nodes to an Existing Rack](#)
You can add nodes to an existing rack using elastic configuration.

6.14 Adding Additional Elastic Nodes to an Existing Rack

You can add nodes to an existing rack using elastic configuration.

The elastic configuration procedure described in [Performing Initial Elastic Configuration of Oracle Exadata Database Machine](#) can also be used to add new servers to an existing rack. The new servers and existing servers must be at least X5. The procedure assigns IP addresses to the new database servers and cells, in preparation for adding them into existing setups. The addition of X5 or later model servers into previous generation racks using this methodology is not supported.

The XML configuration file generated by the Oracle Exadata Deployment Assistant (OEDA) must contain the existing nodes, plus the new nodes to be added.

 **Note:**

On database servers (excluding X8M-2 and higher), run `reclaimdisks.sh` before running the `applyElasticConfig.sh` script.

If you are adding only storage servers, then you can run the `applyElasticConfig.sh` script from a storage server rather than a database server.

The overall process for adding a new server to an existing rack is:

1. Use Oracle Exadata Deployment Assistant (OEDA) to add the new server to the rack configuration and generate the configuration files.

2. Download and unzip the OEDA archive to an existing server in the rack where you are adding the new server.
3. Install the new server into the rack.
4. Review the OEDAconfiguration file for the rack to determine the location within the rack (the `uLocation`) for the new server. For example:

```
<uLocation>16</uLocation>
```

5. Log in to the ILOM on the new server using the local serial management connection (SER MGT port), and review the system location in the SP.

For example:

```
-> show /sp/ system_location
```

```
/SP
```

```
Properties:
```

```
system_location = ru_06
```

```
->
```

6. If the location in the ILOM SP is different from the OEDA location, then modify the location in the ILOM SP to match the value from OEDAILOM.

 **Note:**

Ensure that there are no duplicate locations in the OEDA rack configuration XML file.

Set the system location in the ILOM SP with the following command:

```
-> set /SP/ system_location=RU16
```

7. Log in to the server where you unzipped the OEDA configuration files and run the `applyElasticConfig.sh` script.

For example, if you unzipped the OEDA ZIP file into the `/opt/oracle.SupportTools/onecommand/linux-x64` directory, then use the following commands:

```
# cd /opt/oracle.SupportTools/onecommand/linux-x64
```

```
# ./applyElasticConfig.sh -cf customer_name-configFile.xml
```

The `applyElasticConfig.sh` script performs network configuration for the new servers. The new servers are restarted at the end of the process.

6.15 Using the OEDA setuprootssh Utility

You can set up key-based authentication for the `root` user by using the `setuprootssh.sh` utility included with OEDA.

Starting with the April 2020 release of Oracle Exadata Deployment Assistant (OEDA), the `setuprootssh.sh` utility is located in the same directory as the `install` and `config` scripts.

- To set up the required keys for SSH key-based authentication, run the `setuprootssh.sh` utility.

```
./setuprootssh.sh -cf config.xml
```

For `config.xml`, specify the full path name to the deployment file created by OEDA.

For each physical and virtual machine (VM) host in the deployment file, the `setuprootssh.sh` utility creates an SSH key pair. The keys are saved in the OEDA `WorkDir` using the following file naming conventions:

- Private key: `id_rsa.short_hostname.root`
- Public key: `id_rsa.short_hostname.root.pub`

In the key file names, `short_hostname` is the corresponding server host name without any domain name qualification.

Where possible, the utility also installs the public key on the host and disables password-based authentication for the `root` user.

If the utility cannot install the public key on the host, then the key file remains ready for later use in the OEDA `WorkDir`. For example, this occurs to keys for VM hosts if you run the utility before you create the VMs.

- If your system has a non-default `root` password, you must specify it by using the `-p root_password` option.
- Use the `-h` option to display the usage information.

6.16 Loading the Configuration Information and Installing the Software

You can use Oracle Exadata Deployment Assistant to configure your rack, or you can do it manually.

Oracle Exadata Deployment Assistant (OEDA) loads the network settings, creates the user accounts, installs Oracle Database software, and secures the system based on the information the configuration files.

 **Note:**

- To enable a supported, non-English language, set the following environment variables before starting the configuration:

```
$ export LANG=preferred_language
$ export LC_ALL=preferred_language
```

- To copy the database image to a USB drive, use the procedure in the database maintenance chapter of *Oracle Exadata Database Machine Maintenance Guide* For releases 12.1.2.1.0 and later, the system used to create the image must be running Oracle Linux 6 or later.

- [Configuring Oracle Exadata Database Machine Using OEDA](#)
Use Oracle Exadata Deployment Assistant (OEDA) to configure your engineered system.
- [Using Oracle Exadata Deployment Assistant in Conjunction with Key-Based Authentication](#)
You can use the `-sshkeys` option to use SSH key-based authentication in conjunction with OEDA.
- [Using Oracle Exadata Deployment Assistant on Systems with Non-Default root Passwords](#)
You can use the `-p` option to prompt for a password instead of resetting the `root` password to the expected default value.
- [Configuring Oracle Database and Oracle ASM Instances for Oracle Exadata Database Machine Manually](#)

 **See Also:**

[Using Oracle Exadata Deployment Assistant](#) for descriptions of the configuration files generated by OEDA.

6.16.1 Configuring Oracle Exadata Database Machine Using OEDA

Use Oracle Exadata Deployment Assistant (OEDA) to configure your engineered system.

The deployment phase of OEDA can be performed on one of the database servers of your new Oracle Exadata Database Machine or any other host that is connected to the same network as the new engineered system rack. Do not run `install.sh` from a storage server. If you plan to run OEDA on a host outside of the new rack, then create the OEDA home directory and `workdir`, as described below, on the host machine instead.

 **Note:**

OEDA configures all Exadata components, including the database servers (with or without virtualization), storage servers, and the RDMA Network Fabric. For successful configuration, OEDA requires access to:

- The client and administration network interfaces on the database servers, including the VM server host and the VM guests where applicable.
- The administration network interfaces on the storage servers.
- The administration network interfaces on the RDMA Network Fabric switches.

Ensure that you run the OEDA deployment phase on a host with access to all of the required network interfaces.

 **Note:**

Running OEDA from a client machine outside of the rack means the software image files are transferred over the network, causing the process to be slower.

 **Note:**

For this procedure, Oracle VM refers to OVM and KVM. Starting with Oracle Exadata Database Machine X8M-2, Oracle Linux KVM is the virtualization technology for systems that use RoCE Network Fabric. Otherwise, Oracle VM is a Xen-based virtualization technology, which is used across Oracle Exadata Database Machine systems that use InfiniBand Network Fabric.

You must have completed the other tasks described in [Overview of Oracle Exadata Deployment Assistant](#) before starting this task.

1. If you are installing Oracle VMs on the database servers, then run `/opt/oracle.SupportTools/switch_to_ovm.sh` on each database server.
2. Disk space for unused system partitions must be reclaimed before installing software on all database servers (except KVM environments on X8M servers). Use the following script to reclaim the disk space: `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim`.
3. Create a directory for the extracted OEDA files on the host where you will run `install.sh`.

OEDA will not run if staged on the root file system (`/`) of any database server. Create a directory to store the latest version of OEDA and all required files on a file system that has at least 30 GB of available free space.

- For bare-metal database servers, create a directory on the `/u01` file system, for example `/u01/onecommand` or `/u01/oeda`.

- For Oracle VM servers, create a directory on the `/EXAVMIMAGES` file system.
4. Download the latest version of OEDA for the operating system of the client or server where you will run `install.sh`.

Unzip the downloaded file into the OEDA home directory you created. This will create a directory like `/u01/onecommand/linux-x64`. This directory is referred to as the OEDA home directory.

5. Download the required image files for the Oracle software.

OEDA requires patches available from My Oracle Support or Oracle Software Delivery Cloud (<https://edelivery.oracle.com>). Review the README file for OEDA to determine the file name or patch number, which is based on the version of Oracle software being installed. Place the files in the `WorkDir` sub-directory in the OEDA home directory.

- If you are installing on a bare-metal database server, then download the patches, as specified in Section 2 of the OEDA README file. After downloading the patch ZIP file, extract the contents into the `WorkDir` sub-directory.
- If you are configuring Oracle VM on the database servers, then download the gold image files as indicated in the topic "Oracle Exadata Deployment Assistant Support for Oracle Virtual Machines" in the OEDA README file. After downloading the gold image ZIP files, extract the contents into the `WorkDir` sub-directory.

A gold image is a copy of a software-only, installed Oracle home. It is used to copy an image of an Oracle home to a new host on a new file system to serve as an active, usable Oracle home.

6. Copy the required OEDA configuration files to the `WorkDir` sub-directory on the database server or Oracle VM.

When you used OEDA to specify the configuration details for your new engineered system, it created several configuration files and a ZIP file that contains all the generated files. The required files are also listed at the bottom of the `InstallationTemplate.html` file that was generated by OEDA. You can copy the ZIP file and extract it into the `WorkDir` sub-directory, or copy the required files directly into the `WorkDir` sub-directory.

7. Log in to the database server or Oracle VM and go to the OEDA home directory.
8. Run the `install.sh` script to install and deploy the software on your new engineered system.

In the following command syntax, *configuration_file* is the full path name for the master configuration file, *CustomerName-rackname.xml*

```
# ./install.sh -cf configuration_file [-s step_number | \
-r step_numberA stepnumberB]
```

You can use the following options with the `install.sh` script:

- `-cf file`: Configuration file name
- `-h`: Help output for the script
- `-l`: List the steps of OEDA
- `-r {n-N | n }`: Run steps *n* through *N* of OEDA, or run the specified step.

- `-s n`: Run step *n* of OEDA.
- `-u {n-N | n }`: Undo steps *n* through *N*, or one step.

**Note:**

Always use the latest OEDA patch listed in My Oracle Support note 888828.1. Refer to the README in the patch for the latest steps.

The following steps are performed by the OEDA deployment script. However, the actual steps for a specific deployment may differ depending on the deployment options chosen. For example, the "Update Nodes for Eighth Rack" step is performed only if you are deploying an Eighth Rack configuration.

1. Validate Configuration File
2. Update Nodes for Eighth Rack
3. Create Virtual Machine
4. Create Users
5. Setup Cell Connectivity
6. Create Cell Disks
7. Create Grid Disks
8. Configure Alerting
9. Install Cluster Software
10. Initialize Cluster Software
11. Install Database Software
12. Relink Database with RDS
13. Create ASM Diskgroups
14. Create Databases
15. Apply Security Fixes
16. Install Exachk
17. Setup ASR Alerting
18. Create Installation Summary
19. Resecure Machine

Related Topics

- [Oracle Exadata Deployment Assistant Downloads](#)
- [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID 888828.1\)](#)

6.16.2 Using Oracle Exadata Deployment Assistant in Conjunction with Key-Based Authentication

You can use the `-sshkeys` option to use SSH key-based authentication in conjunction with OEDA.

Starting with the April 2020 release of Oracle Exadata Deployment Assistant (OEDA), you can use SSH keys for `root` user authentication on Oracle Exadata Racks.

To use this feature, you must first set up the required SSH keys. The easiest way to do this is to use the `setuprootssh.sh` utility included with OEDA. For example:

```
./setuprootssh.sh -cf config.xml
```

For more information about the `setuprootssh.sh` utility, see [Using the OEDA setuprootssh Utility](#).

After you set up the required SSH keys, add the `-sshkeys` option to the regular command-line options when you start up OEDA. For example:

```
./install.sh -cf config.xml -s 1 -sshkeys -usesu
```

When you specify the `-sshkeys` option, OEDA uses SSH key-based authentication. In conjunction with this option, the OEDA `WorkDir` must contain the SSH private key for each host in the deployment file (`config.xml`). The private keys must conform to the following file naming convention:

```
id_rsa.short_hostname.root
```

In the key file names, `short_hostname` is the corresponding server host name without any domain name qualification.

As shown in the previous example, when you specify the `-sshkeys` option, you must also specify the `-usesu` option to instruct OEDA to run non-`root` commands from the `root` user using `su`.

When you first log in to a host following the `Resecure Machine` deployment step, you are prompted to reset the `root` password. This still occurs even when SSH key-based authentication is enabled, and password-based authentication is disabled.

6.16.3 Using Oracle Exadata Deployment Assistant on Systems with Non-Default root Passwords

You can use the `-p` option to prompt for a password instead of resetting the `root` password to the expected default value.

Starting with the November 2016 release, Oracle Exadata Deployment Assistant (OEDA) supports deployment on Oracle Exadata Racks that already have non-default `root` passwords. You can perform the deployment without having to change the `root` password back to the default. In earlier releases, OEDA expected the `root` password to be the default on all the servers in an Oracle Exadata Rack.

With the new non-default password feature, `root` passwords can have different values on each node in a rack. If non-default passwords are already set, add the `-p` option to the existing command line options when you run OEDA. For example:

```
./install.sh -cf <config.xml> -s 1 -p
```

The `-p` option causes OEDA to prompt for the `root` password. Make sure you enter the correct password because an incorrect password disables `root` access for 10 minutes, and after 5 failed login attempts, the server is locked down.

Related Topics

- Default Security Settings Enacted by OEDA

6.16.4 Configuring Oracle Database and Oracle ASM Instances for Oracle Exadata Database Machine Manually

Oracle Database and Oracle Automatic Storage Management (Oracle ASM) 11g Release 2 (11.2) or later must be installed on the database server. Follow the instructions for installing Oracle Database and Oracle ASM provided in the appropriate version of the *Oracle Database Installation Guide* for your platform.



Note:

To fully utilize the Hardware Assisted Resilient Data (HARD) features, ensure the `DB_BLOCK_CHECKSUM` initialization parameter is set to `TYPICAL` (default) or `FULL`.

When installing Oracle Database, you must select the Oracle ASM storage option. If you are installing Oracle Real Application Clusters (Oracle RAC), then refer to the platform-specific Oracle Clusterware and Oracle RAC installation guides, such as *Oracle Grid Infrastructure Installation and Upgrade Guide* and *Oracle Real Application Clusters Installation Guide*.

If you are consolidating databases and need to have more than one database with the same `DB_UNIQUE_NAME` on the Oracle Exadata Database Machine, then you must configure Oracle ASM clusters on the storage cells and configure ASM-scoped security.



Note:

For Oracle RAC, you must shut down the entire cluster, and then restart the cluster after the database and cell software have been installed, and the cell configuration files are in place.

The recommended procedure is to install the latest version of the Oracle software on Oracle Exadata Database Machine using Oracle Exadata Deployment Assistant (OEDA). If you choose to install the software manually, then do the following:

 **Note:**

Disk space for unused system partitions must be reclaimed before installing software on all database servers (except KVM environments on X8M servers). Use the following script to reclaim the disk space: `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim`.

If a RAID rebuild or reconstruction (that is, a disk replacement or expansion) is in progress, wait until it is finished before running the `reclaimdisks.sh` script. The reason is that the `reclaimdisks.sh` script is blocked from running until the RAID rebuild or reconstruction is complete.

1. Install Oracle Grid Infrastructure for a cluster software as described in *Oracle Grid Infrastructure Installation and Upgrade Guide for Linux*.
2. Install Oracle Database software as described in *Oracle Real Application Clusters Installation Guide for Linux and UNIX*.
3. Apply the latest patches for Oracle Database software and Oracle Exadata System Software.

After the cells are configured as described in *Oracle Exadata System Software User's Guide*, and the database and Oracle ASM instances are created, the instances must be configured to use the cells as part of an Oracle ASM disk group. The required configuration involves updating the initialization parameter file of both the database and Oracle ASM instances. Additionally, Oracle ASM disk groups must be created to access storage cells.

You can update the initialization parameter file of the Oracle ASM or the database instance while the instance is running, but the instance must be restarted if you manually update the initialization parameter file. If you update an initialization parameter dynamically using the SQL `ALTER SYSTEM` or `ALTER SESSION` commands, then there is no need to restart the instance.

In addition, for the database, you can configure a new Oracle Database Resource Manager plan, rather than using the default plans.

- [Configuring the Compatible Parameter for a Database Instance](#)
Oracle recommends using the best practices for Oracle Exadata System Software available at My Oracle Support.
- [Configuring Initialization Parameters for an Oracle ASM Instance](#)
- [Using the Same DB_UNIQUE_NAME for Multiple Database Instances](#)
You can create database instances that use the same `DB_UNIQUE_NAME` value if the databases are associated with separate Oracle ASM clusters.

Related Topics

- [Setting Up Oracle ASM-Scoped Security on Oracle Exadata Storage Servers](#)
- [Setting Up Configuration Files for a Database Server Host](#)
- [Managing I/O Resources](#)
- [Administering Oracle ASM Instances](#)
- [Configuring DBFS on Oracle Exadata Database Machine \(My Oracle Support Doc ID 1054431.1\)](#)

- [Oracle Exadata Best Practices \(My Oracle Support Doc ID 757552.1\)](#)

6.16.4.1 Configuring the Compatible Parameter for a Database Instance

Oracle recommends using the best practices for Oracle Exadata System Software available at My Oracle Support.

To configure a database instance to access cell storage, ensure that the `COMPATIBLE` parameter is set to 11.2.0.2 or later in the database initialization file.

You can view the initialization parameter file from the operating system or use the SQL*Plus administration command `SHOW PARAMETER` if the database is running. For example:

```
SQL> SHOW PARAMETER
SQL> SHOW PARAMETER compatible
```

If necessary, set the `COMPATIBLE` initialization parameter in the initialization parameter file to 11.2.0.2 or later. For example:

```
COMPATIBLE='18.0.0'
```

The `COMPATIBLE` parameter cannot be changed dynamically. If you change the value in the initialization parameter file, then the database must be shut down and restarted for the change to take effect.

Related Topics

- `COMPATIBLE`

6.16.4.2 Configuring Initialization Parameters for an Oracle ASM Instance

To enable an Oracle ASM instance to discover and access Exadata Storage Server grid disks, you must configure the `ASM_DISKSTRING` initialization parameter as follows:

- Set the `ASM_DISKSTRING` initialization parameter to the empty string (`' '`) to discover all grid disks on the cells listed in the `cellip.ora` file. This causes Oracle ASM discovery for `/dev/raw/*` and any ASMLIB disks on any non-Exadata Storage Servers.
- Set the `ASM_DISKSTRING` initialization parameter to restrict the discovery by Oracle ASM to only specific sets of disks, as shown in the examples.
 - To discover only a subset of cells, use a pattern for the `ASM_DISKSTRING` value such as the following:

```
o/cell-connect-info-pattern/griddisk-name-pattern.
```

For example:

```
ASM_DISKSTRING = 'o/*/data*'
```

- To discover non-Exadata Storage Server disks and all grid disks on all the cells specified in the `cellip.ora` file, use a wildcard pattern for Exadata Storage Server grid disks. For example:

```
ASM_DISKSTRING = '/dev/rdsk/disk*', 'o/*/*'
```

You can change the `ASM_DISKSTRING` initialization parameter when the Oracle ASM instance is running with the `SQL ALTER SYSTEM` command. If you edit the `ASM_DISKSTRING` initialization parameter in the initialization parameter file when the Oracle ASM instance is running, then the Oracle ASM instance must be shut down and restarted for the change to take effect.

See Also:

- *Oracle Automatic Storage Management Administrator's Guide* for additional information about the following:
 - Oracle ASM discovery strings
 - starting up and connecting to an Oracle ASM instance
- *Oracle Database Reference* for additional information about the `ASM_DISKSTRING` initialization parameter
- *Oracle Database SQL Language Reference* for additional information about the `ALTER SYSTEM` command

6.16.4.3 Using the Same `DB_UNIQUE_NAME` for Multiple Database Instances

You can create database instances that use the same `DB_UNIQUE_NAME` value if the databases are associated with separate Oracle ASM clusters.

Starting with Oracle Exadata System Software release 19.1.0, Oracle Multitenant database clusters sharing the same storage can contain database instances that use the same `DB_UNIQUE_NAME`.

WARNING:

If you configure databases to have the same `DB_UNIQUE_NAME`, then those databases cannot be backed up to Oracle Zero Data Loss Recovery Appliance.

1. Create Oracle ASM clusters for the database clusters that use identical `DB_UNIQUE_NAME` values for the database instances.

Each database instance that uses the same `DB_UNIQUE_NAME` must be in a separate Oracle ASM cluster. The Oracle ASM cluster name is used to qualify the `DB_UNIQUE_NAME` in I/O Resource Management (IORM), Exadata Smart Flash Cache, and Exadata Smart Scan Offload operations.

2. Configure ASM-scoped security for each cluster database.

When configuring ASM-scoped security, the Oracle ASM client must be configured as an `ASMCLUSTER` client.

If you already have ASM-scoped security configured, but the Oracle ASM client is not configured as an `ASMCLUSTER` client, then you can use the `ASSIGN KEY` command upgrade the key to an `ASMCLUSTER` key.

3. Create the databases in each cluster using the same `DB_UNIQUE_NAME`.

Related Topics

- [Setting Up Oracle ASM-Scoped Security on Oracle Exadata Storage Servers](#)
- [Upgrading ASM-Scoped Security Key for ASMCLUSTER](#)

6.17 Using the OEDA changePassword Utility

You can change the password of the `grid` and `oracle` users across all clusters by using the `changePassword.sh` utility included with OEDA.

The utility is located in the same directory as the `install` and `config` scripts.

You cannot use this utility if the servers are locked down.

- To change the password of both `oracle` and `grid` users on all database servers and clusters, run the `changePassword` command.

```
./changePassword.sh -cf config.xml -userName root
```

For `config.xml`, specify the full path name to the deployment file created by OEDA.

The utility prompts for the current `root` password and allows you to change the password of any users in the rack.

- To change the password for a specific user, such as `oracle` or `grid`, use the `-userName user_name` option with the command.
- To change the password for all users within a specific cluster, use the `-clusterName cluster_name` option with the command.
- Use the `-h` option with the command to display the usage information.

6.18 Installing Oracle Enterprise Manager Cloud Control

Oracle Exadata Database Machine should be monitored using Oracle Enterprise Manager Cloud Control. Ensure that the Enterprise Manager environment is configured for high availability and disaster recovery to ensure sufficient monitoring of the Exadata environments.

Related Topics

- [Oracle Enterprise Manager Cloud Control Basic Installation Guide](#)
- [Installation and Deployment in Oracle Enterprise Manager Exadata Management Getting Started Guide](#)
- [Configuring Enterprise Manager for High Availability in Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide](#)
- [Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide](#)
- [Discovery in Oracle Enterprise Manager Cloud Control Administrator's Guide](#)
- [How to Configure Oracle Configuration Manager for Oracle Exadata Database Machine \(My Oracle Support Doc ID 1319476.1\)](#)
- [MAA Best Practices website](#)

6.19 Adding a VM Cluster to Oracle Exadata Database Machine Using OEDA

You can use Oracle Exadata Deployment Assistant (OEDA) to add a virtual machine (VM) cluster to an existing Oracle Exadata Database Machine deployment.

The procedure has the following prerequisites:

- The system must support an existing VM deployment.
- You must have free space on your system to accommodate the new VM cluster. Specifically:
 - You must not have the maximum number of VM clusters already deployed.
 - You must have available CPU, memory, and local storage resources on the database servers.
 - You must have available storage space on the storage servers.
- You require a current copy of the OEDA system configuration XML file.

To add a VM cluster to an existing Oracle Exadata Database Machine deployment:

1. Start the OEDA Web-based configuration tool.
See [OEDA Browser-based User Interface](#).
2. Import the existing system configuration XML file.
Use the **Import** option located in the drop-down menu at the top of the page.
3. Click **Clusters** in the navigation menu on the left of the page.
The **Define Clusters** section displays the configuration information for the existing clusters.
4. Click the **Add Clusters** button.
The **Add Clusters** button contains the plus symbol (+) and is adjacent to the existing cluster tabs.
Clicking the **Add Clusters** button adds a tab in the **Define Clusters** section to contain the configuration details for the new cluster.
5. In the new cluster tab, specify the configuration details for the new cluster. Then, click **Apply**.
OEDA now displays the **Diskgroups** section.
6. In the **Diskgroups** section, click the tab for the storage that is associated with the newly defined cluster.
7. Specify the configuration details for the storage that is associated with the new cluster. Then, click **Apply**.
OEDA now displays the **Create Database Home** section.
8. Click the **Add Database Home** button, which contains the plus symbol (+) and is adjacent to the existing tabs.
A new tab appears in the **Create Database Home** section to contain the configuration details for a new database home.

9. In the new tab, specify the configuration details for the new database home. Then, click **Apply**.

Ensure that you select the newly defined cluster as the location to contain the new database home.

Afterward, OEDA displays the **Create Database** section.

10. Click the **Add Database** button, which contains the plus symbol (+) and is adjacent to the existing tabs.

A new tab appears in the **Create Database** section to contain the configuration details for a new database.

11. In the new tab, specify the configuration details for the new database. Then, click **Apply**.

Ensure that you select the newly defined database home to support the new database.

Afterward, OEDA displays the **Cluster Networks** section.

12. For each network in the **Cluster Networks** section, select the tab that is associated with the new cluster and specify the network settings. Then, click **Apply**.

As you step through the network settings, ensure that you specify valid IP address settings that are compatible with the existing system configuration.

Afterward, OEDA displays the **Alerting** section.

13. Click **Comments** in the navigation menu. Then, click **Apply and Download** to save the updated system configuration information.

14. Use the updated system configuration XML file to deploy the new cluster.

Use the OEDA deployment utility (`install.sh`) in conjunction with the updated system configuration XML file.

Perform all of the deployment steps except for the following, which do not apply. Ensure that you do not perform the following deployment steps as the unnecessary performance of these steps may impact system performance and availability.

- Update Nodes for Eighth Rack (For Eighth Rack systems only)
- Calibrate Cells
- Create Cell Disks
- Resecure Machine

Also, when you perform the `Create Installation Summary` step, ensure that you use the `-nocalibratecell` option to ensure that cell calibration is not performed during that step. For example, if the step number for the `Create Installation Summary` step is 18:

```
$ ./install.sh -cf es.xml -s 18 -nocalibratecell
```


To view a complete list of the deployment steps and the corresponding step numbers, use `install.sh` with the `-l` option. For example:

```
$ ./install.sh -cf es.xml -l

1. Validate Configuration File
2. Setup Required Files
3. Create Users
4. Setup Cell Connectivity
5. Verify Infiniband
6. Calibrate Cells
7. Create Cell Disks
8. Create Grid Disks
9. Configure Cell Alerting
10. Install Cluster Software
11. Initialize Cluster Software
12. Install Database Software
13. Relink Database with RDS
14. Create ASM Diskgroups
15. Create Databases
16. Apply Security Fixes
17. Setup ASR Alerting
18. Create Installation Summary
19. Resecure Machine
```

To run an individual step, use `install.sh` with the `-s` option and specify the step number. For example, to run step 1:

```
$ ./install.sh -cf es.xml -s 1
```

To run a range of steps, use `install.sh` with the `-r` option and specify the step range. For example, to run steps 12-16:

```
$ ./install.sh -cf es.xml -r 12-16
```

Related Topics

- [Configuring Oracle Exadata Database Machine Using OEDA](#)
Use Oracle Exadata Deployment Assistant (OEDA) to configure your engineered system.

A

Site Checklists

Complete the checklists in this appendix to ensure the site is prepared for your new Oracle Engineered System.

 **Note:**

For ease of reading, the Oracle Engineered System is sometimes referred to simply as "the rack", which indicates the entire system, not just the cabinet.

- [System Components Checklist](#)
Use this checklist to ensure that the system component considerations have been addressed.
- [Data Center Room Checklist](#)
Use this checklist to ensure that the data center room requirements have been addressed.
- [Data Center Environment Checklist](#)
Use this checklist to ensure that the data center environment requirements have been addressed.
- [Access Route Checklist](#)
Use this checklist to ensure that the access route requirements are addressed.
- [Facility Power Checklist](#)
Use this checklist to ensure that the facility power requirements have been addressed.
- [Safety Checklist](#)
Use this checklist to ensure that the safety requirements have been addressed.
- [Logistics Checklist](#)
Use this checklist to ensure that the logistics requirements have been addressed.
- [Network Configuration Checklist](#)
Use this checklist to ensure that the network configuration requirements have been addressed.
- [Auto Service Request Checklist](#)
Use this checklist if you are planning to use Auto Service Request with the Oracle Exadata Database Machine.
- [Re-racking Checklists](#)
Use these checklists prior to re-racking configurations approved for re-rack by Oracle Exadata Configuration Assistant (OECA) or your Oracle contact.
- [PDU Checklist for Re-racking](#)
Use this checklist to verify the power distribution unit (PDU) requirements when using the re-racking service.

- [Recycling the Racks](#)
Use this checklist when recycling the rack.

A.1 System Components Checklist

Use this checklist to ensure that the system component considerations have been addressed.

The cells in the second column of the following table are intentionally left blank so that the site survey team can fill in the requested information.

Table A-1 System Components Checklist

System Components Checklist Items	Ye s	No	N/ A	Comment
What type of engineered system rack are you installing? (Exadata Database Machine, Big Data Appliance, Oracle Cloud Machine, Exadata Storage Expansion, Zero Data Loss Recovery Appliance, etc.)				
Is the rack installation an extension to an existing system, rather than a new installation? If it is an extension to an existing system, refer to the multi-rack cabling information for your Engineered System.				
How many racks will be installed?				
Have you ordered the multi-rack cabling service?				
Are all Engineered System racks adjacent to each other?				
If the connecting racks are not within the specified proximity, or you are connecting more than 3 racks together, then have you completed the following tasks? <ul style="list-style-type: none"> • Purchased longer InfiniBand cables from an approved third-party provider? • Asked Oracle Support Services to provide and schedule the custom multi-rack cabling service with the installation? 				
If you are connecting the new rack to other Oracle racks, then does each rack (excluding Eighth Rack configurations) have a spine switch?				
If any of the racks you plan to interconnect using the InfiniBand network are missing a spine switch, then have you ordered the InfiniBand switch and cables?				
If your data center does not have a 10 GbE network infrastructure, then have you ordered an external switch that supports both 1 GbE and 10 GbE (or 40 GbE)?				
Have you purchased the cables and transceivers to connect the InfiniBand Gateway switches to your 10 GbE network infrastructure?				
If the Ethernet switch be replaced, then have you arranged for the replacement switch installation?				

Table A-1 (Cont.) System Components Checklist

System Components Checklist Items	Yes	No	N/A	Comment
Will additional equipment be attached to or installed in the rack? If yes, then: <ul style="list-style-type: none"> • Describe the additions • Have you verified that the additional equipment falls within Oracle guidelines? • Is sufficient power and cooling available for the additional equipment? 				

A.2 Data Center Room Checklist

Use this checklist to ensure that the data center room requirements have been addressed.

Data Center Room Checklist Items	Yes	No	N/A	Comment
Has the rack location been allocated and is vacant?				
Does the floor layout meet the equipment maintenance access requirements?				
Will the rack be positioned so that the exhaust air of one rack does not enter the air inlet of another rack?				
Have cabinet stabilization measures been considered?				

Data Center Room Checklist Items	Yes	No	N/A	Comment
If the data center has a raised floor: <ul style="list-style-type: none"> • Does the raised floor satisfy the weight requirements for the new hardware? • Is permission required to remove floor tiles for cabling and servicing below the floor? 				
Will the new hardware location require any non-standard cable lengths?				
Is the floor to ceiling height a minimum of 2914 mm (114.72 inches)?				
Is the depth of the raised floor a minimum of 46 cm (18 inches)?				

Related Topics

- [Space Requirements](#)
The space requirements for Oracle Exadata Database Machine depend on the hardware rack that is supplied with the system.
- [Flooring Requirements](#)
Oracle Exadata Rack may be installed on raised floor or solid floor environments.
- [Stabilize Oracle Exadata Rack with Leveling Feet](#)

A.3 Data Center Environment Checklist

Use this checklist to ensure that the data center environment requirements have been addressed.

Data Center Environment Considerations	Yes	No	N/A	Comment
Does the computer room air conditioning meet temperature and humidity requirements?				
Does the installation floor layout satisfy the ventilation requirements?				
If the room cooling is from a raised floor, then: <ul style="list-style-type: none"> • Are the perforated floor tiles each rated at 400 CFM or greater? • Can additional perforated floor tiles be obtained if required for additional cooling? 				
Does the data center air conditioning provide sufficient front-to-back airflow?				
Is airflow adequate to prevent hot spots?				
Can the data center continuously satisfy the environmental requirements?				

Related Topics

- [Temperature and Humidity Requirements](#)
Excessive internal temperatures may result in full or partial shut down of Oracle Exadata Database Machine.

- [Ventilation and Cooling Requirements](#)
Always provide adequate space in front and behind the rack to allow for proper ventilation.
- [General Environmental Requirements](#)

A.4 Access Route Checklist

Use this checklist to ensure that the access route requirements are addressed.

Access Route Considerations	Yes	No	N/A	Comment
<p>Has the access route been checked for clearances of the rack, in particular, the minimum width and height requirements for all doors on the route?</p> <p>Dimensions of the packaged rack:</p> <ul style="list-style-type: none"> • Height: 2159 mm (85 inches) • Width: 1219 mm (48 inches) • Depth: 1575 mm (62 inches) <p>Dimensions of the unpackaged rack:</p> <ul style="list-style-type: none"> • Height: 1998 mm (78.66 inches) • Width: 600 mm with side panels (23.62 inches) • Depth 1200 mm (47.24 inches) 				
<p>Are there any stairs, ramps, or thresholds that are of concern? If yes, then provide details.</p>				

Access Route Considerations	Yes	No	N/A	Comment
Are all access route incline angles within the permitted range (6 degrees)?				
Is the access route free of any obstacles that would expose the equipment to shock?				
Are all the surfaces acceptable for rolling the new unpacked and packed equipment?				
If a pallet jack is to be used: <ul style="list-style-type: none"> • Can the pallet jack support the weight of the rack? • Are the pallet jack tines compatible with the shipping pallet? 				
If there are stairs, is a loading elevator available for the equipment?				

Access Route Considerations	Yes	No	N/A	Comment
If an elevator is to be used: <ul style="list-style-type: none"> • Is the elevator car wide enough for the device to be carried into it? • Is the elevator car high enough for the device to be carried into it? • Does the elevator door meet the height requirement ? • Can the elevator car handle the weight of a fully-loaded rack? 				
Can the access route support the weight of the unpacked rack?				
Is the access route onto the raised floor rated for dynamic loading of the rack?				

Related Topics

- [Table: Access Route Requirements.](#)
- [Rack Weights for Oracle Exadata Database Machine](#)
Review the net and shipping weights for your Oracle Exadata Rack to ensure it can be delivered safely.
- [Rack Weights for Oracle Exadata Storage Expansion Rack](#)
- [Flooring Requirements](#)
Oracle Exadata Rack may be installed on raised floor or solid floor environments.

A.5 Facility Power Checklist

Use this checklist to ensure that the facility power requirements have been addressed.

Facility Power Considerations	Yes	No	N/A	Comment
Have the operating voltage and electric current requirements listed in "Electrical Power Requirements" been reviewed?				
Will single-phase (low-voltage or high-voltage) or 3-phase (low-voltage or high-voltage) power be used?				
Are enough power outlets provided within 2 meters for each rack?				
Do the power outlets have appropriate socket receptacles for the PDU option ordered? Options are low voltage or high voltage, single-phase or 3-phase.				
Will optional ground cables be attached to the rack?				
Are the circuit breakers for the equipment suitable in terms of voltage and current-carrying capacities?				
Does the power frequency meet the equipment specifications?				
Are power outlets available for the new equipment at the designated location?				

Facility Power Considerations	Yes	No	N/A	Comment
Will system power be delivered from two separate grids?				
Is there a UPS to power the equipment?				
Are the minimum required power sources available to support the power load for the new hardware? Use kilowatt (kW) / kilovolt (kVA) to express power load.				

A.6 Safety Checklist

Use this checklist to ensure that the safety requirements have been addressed.

Safety Checklist Considerations	Yes	No	N/A	Comment
Is there an emergency power shut off?				
Is there a fire protection system in the data center room?				
Is the computer room adequately equipped to extinguish a fire?				
Is antistatic flooring installed?				
Is the area below the raised floor free of obstacles and blockages?				

Related Topics

- [Temperature and Humidity Requirements](#)
Excessive internal temperatures may result in full or partial shut down of Oracle Exadata Database Machine.

 **See Also:**

"Powering On and Off Oracle Exadata Rack" in the *Oracle Exadata Database Machine Maintenance Guide*

A.7 Logistics Checklist

Use this checklist to ensure that the logistics requirements have been addressed.

Logistics Checklist Considerations	Yes	No	N/A	Comment
Is contact information for the data center personnel available?				
Is there security or access control for the data center?				
Are there any security background checks or security clearances required for Oracle personnel to access the data center? If yes, then provide the process for Oracle to follow.				
How many days in advance must background checks be completed?				
Are there any additional security access issues?				
Is computer room access available for installation personnel?				
Are laptops allowed in the data center?				

Logistics Checklist Considerations	Yes	No	N/A	Comment
Are cell phones allowed in the data center?				
Are cameras allowed in the data center?				
Does the building have a delivery dock?				
Is there a delivery/ unpacking/ staging area?				
Is inside delivery planned (direct to the final rack location in the data center room)?				
If the delivery is not inside, then is the site prepared for uncrating?				
Is the delivery/ unpacking/ staging area protected from the elements?				
Does the building have adequate receiving space?				
Is the unpacking area air-conditioned to avoid thermal shock for various hardware components?				
Will sufficient moving personnel be available to transport the rack?				

Logistics Checklist Considerations	Yes	No	N/A	Comment
(If applicable to your country) Is union labor required for any part of the delivery or installation?				
Is the site prepared for uncrating and packaging removal? Package removal should take place <i>outside</i> the data center room.				
Is uncrating of cabinet and cabinet packaging removal required?				
Are there any restrictions on delivery truck length, width or height?				
Is there storage space (cabinet) for the ride along spares? If not, does the customer allow cardboard boxes and other packing material in the computer room, since the spares are packed in cardboard boxes?				
Is there a time constraint on dock access? If yes, please provide time constraints.				

Logistics Checklist Considerations	Yes	No	N/A	Comment
Is a tail or side lift required on the delivery carrier to unload the equipment at the delivery dock?				
Will any special equipment be required to place the rack in the data center room? Examples: <ul style="list-style-type: none"> • Stair walkers • Lifters • Ramps • Steel plates • Floor covers 				
Does the delivery carrier require any special equipment, such as non-floor damaging rollers, transport dollies, pallet jacks or fork lifts?				

Related Topics

- [Space Requirements](#)
The space requirements for Oracle Exadata Database Machine depend on the hardware rack that is supplied with the system.

A.8 Network Configuration Checklist

Use this checklist to ensure that the network configuration requirements have been addressed.

Network Configuration Considerations	Yes	No	N/A	Comment
Were any conflicts found while running the <code>checkip.sh</code> script before connecting the network? See "Default IP Addresses" for information on the script.				
Will the required network cables be laid from the network equipment to the location where the Oracle Exadata Rack will be installed?				
Will the network cables that will connect to the Oracle Exadata Rack be labeled?				
Will the 10 GbE or 25 GbE interfaces be used for the client access network?				
Have the optional 10 GbE or 25 GbE SFP+ transceivers and appropriate fiber cables been ordered?				
Will the Cisco Ethernet switch have IP routing disabled (recommended)?				

Related Topics

- [Understanding the Network Requirements for Oracle Exadata Database Machine](#)
Review the network requirements for Oracle Exadata Database Machine before installing or configuring the hardware.

A.9 Auto Service Request Checklist

Use this checklist if you are planning to use Auto Service Request with the Oracle Exadata Database Machine.

Auto Service Request Considerations	Yes	No	N/A	Comment
Does the customer have a My Oracle Support account to register Auto Service Request?				
Is the Oracle Support Customer Support Identifier (CSI) number available?				
Are the host name and IP address for the server where ASR Manager will be installed available?				
Will the system need a proxy server? If so, what is the host name and IP address for the proxy server?				
Is the Customer Technical Contact information for Auto Service Request available? This information should include the first name, last name, and e-mail address of the contact.				

 **See Also:**

"Understanding Auto Service Request" in the *Oracle Exadata Database Machine Maintenance Guide*

A.10 Re-racking Checklists

Use these checklists prior to re-racking configurations approved for re-rack by Oracle Exadata Configuration Assistant (OECA) or your Oracle contact.

 **Note:**

- Customer must purchase both the Oracle Re-racking Service and Oracle Installation Service.
- Oracle does not provide support for customer-supplied equipment.

Re-racking Considerations	Yes	No	N/A	Comment
Has the customer purchased the Oracle Re-racking Service?				
Is there a cart capable of carrying the weight of the servers to move the components and associated cabling from the supplied rack to the customer supplied rack?				
Is the target rack empty?				

Re-racking Considerations	Yes	No	N/A	Comment
<p>Do the target rack dimensions meet the following requirements?</p> <ul style="list-style-type: none"> • Height: 42 RU • Width: 600 mm (23.62 inches) • Depth: 1112 mm (43.78 inches) without front and rear doors <p>If the rack is less than 42 RU, then the customer must provide compatible PDUs to install in the target rack.</p>				
<p>Is the distance between the front and rear mounting planes between the minimum of 610 mm and the maximum 915 mm (24 inches to 36 inches)?</p>				
<p>Is the clearance depth in the front of the front mounting plane (distance to the front cabinet door) at least 25.4 mm (1 inch)?</p>				
<p>Does the target rack meet the following minimum load capacity?</p> <ul style="list-style-type: none"> • 19 kg (41.89 lb) / RU 				

Re-racking Considerations	Yes	No	N/A	Comment
Is the rack a four-post rack (mounting at both front and rear)? Note: Two-post racks are not compatible.				
Does the target rack's horizontal opening and unit vertical pitch conform to ANSI/EIA 310-D-1992 or IEC 60297 standards?				
Does the target rack have RETMA rail support? Note: Oracle Engineered System racks require 19 inches (483 mm) for RETMA rail spacing width. The minimum rack width of 600 mm (23.63 inches) is recommended to accommodate the PDU and cable harnesses on the side. If the rack is less than 600 mm wide, then it must have additional depth to accommodate mounting behind the server CMAs.				
Does the target rack support Oracle cable management arms (CMAs)?				
Does the target rack support installation of Oracle vented and solid filler panels?				

Re-racking Considerations	Yes	No	N/A	Comment
Can the target rack provide tie-downs along the left rear side of the rack to support the RDMA Network Fabric cables? The side is left rear side when viewed from the front of the rack.				
Can the target rack provide tie-downs for the Ethernet wiring harness?				
Is there sufficient space for the cable harnesses and the PDUs in the target rack?				
Can a label with the Oracle Engineered System Rack serial number be printed and attached to the target rack?				

A.11 PDU Checklist for Re-racking

Use this checklist to verify the power distribution unit (PDU) requirements when using the re-racking service.

PDU Re-racking Considerations	Yes	No	N/A	Comment
Does the target rack support installation of standard Oracle PDUs? If not, then complete this checklist.				
Can the customer provide an equivalent pair of PDUs?				

PDU Re-racking Considerations	Yes	No	N/A	Comment
Can the customer provide two PDUs with capacity of 10 kVA per PDU?				
Can the customer provide at least 17 x 10A C13 plugs per PDU?				
Can the customer provide a single PDU and its circuits to support the Oracle Engineered System Rack power requirements in case one PDU fails?				
Can the customer ensure power loads are evenly distributed across all circuits of a single PDU?				
Can the customer provide appropriate power drops for the PDUs?				

A.12 Recycling the Racks

Use this checklist when recycling the rack.

Recycling Considerations	Yes	No	N/A	Comment
Does the customer want Oracle to take back and recycle the empty rack? If yes, then the Oracle service engineer can pack the empty rack for recycling.				

Recycling Considerations	Yes	No	N/A	Comment
If Oracle is recycling the rack, then has the customer completed the appropriate application at the Oracle Hardware Take Back Program website?				

Related Topics

- [Product Take Back and Recycling website](#)

B

Database High Availability Checklist

Use this checklist to ensure the requirements for high availability are met for Oracle Exadata Database Machine.

Database High Availability Considerations	Yes	No	N/A	Comment
Have high redundancy disk groups been configured for best protection against different storage failures and corruptions?				
Do you have a bonded network for client access?				
Will the pre-optimized and configured database be used after deployment?				
Have you reviewed the Oracle Exadata System Software patching overview and patch testing guidelines in My Oracle Support note 1262380.1?				
Have you reviewed the Oracle Exadata System Software Maximum Availability Architecture (MAA) Web site for post-deployment steps to achieve MAA?				

Database High Availability Considerations	Yes	No	N/A	Comment
Will you use Oracle Exadata Database Machine health check after deployment? The health check is available from My Oracle Support note 1070954.1.				
Will you follow Oracle Exadata Database Machine monitoring best practices available in My Oracle Support note 1920255.1?				
Have you reviewed the recommended software for Oracle Exadata Database Machine in My Oracle Support note 888828.1?				
Have you reviewed Oracle Exadata Database Machine best practices in My Oracle Support note 757552.1?				

Related Topics

- [Oracle Database High Availability Overview](#)
- [Exadata Patching Overview and Patch Testing Guidelines \(My Oracle Support Doc ID 1262380.1\)](#)
- [Oracle Maximum Availability Architecture website](#)
- [Oracle Exadata Database Machine exachk or HealthCheck \(My Oracle Support Doc ID 1070954.1\)](#)
- [Exadata Health and Resource Usage Monitoring Technical Paper \(My Oracle Support Doc ID 1920255.1\)](#)
- [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID 888828.1\)](#)
- [Oracle Exadata Best Practices \(My Oracle Support Doc ID 757552.1\)](#)