Oracle® Exadata

Installation and Configuration Guide



26.1
F29249-49
July 2026

Oracle Exadata Installation and Configuration Guide, 26.1

F29249-49

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

Primary Author: Peter Fusek

Contributing Authors: Craig Brown, Glenn Maxey

Contributors: Doug Archambault, Leo Agranonik, Andrew Babb, Nilesh Choudhury, Henry Chow, Ravindra Dani, Boris Erlikhman, Jaime Figueroa, Roger Hansen, Leslie Keller, Frank Kobylanski, René Kundersma, Holger Leister, Yang Liu, Juan Loaiza, Barb Lundhild, Catherine Luu, Philip Newlan, 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, Alex Blyth, Gautam Bhatt

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

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

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

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

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

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

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

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

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.

Oracle customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

Contents

Preface

Audience i
Related Documentation i

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

1.1 Ensuring That the Site is Ready 1
1.2 General Environmental Requirements 2
1.2.1 General Environmental Requirements for Oracle Exadata Rack X6 and Later 3
1.2.2 General Environmental Requirements for Oracle Exadata X5-2 3
1.2.3 General Environmental Requirements for Oracle Exadata Database Machine X5-8 5
1.2.4 General Environmental Requirements for Oracle Exadata Database Machine X4-2 6
1.2.5 General Environmental Requirements for Oracle Exadata Database Machine X4-8 with Exadata Storage Server X5-2 Servers 9
1.2.6 General Environmental Requirements for Oracle Exadata Database Machine X4-8 Full Rack 10
1.2.7 General Environmental Requirements for Oracle Exadata Database Machine X3-2 11
1.2.8 General Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X4-2L Servers 14
1.2.9 General Environmental Requirements for Oracle Exadata Database Machine X3-8 Full Rack with Exadata Storage Server X3-2 Servers 15
1.2.10 General Environmental Requirements for Oracle Exadata Database Machine X2-2 16
1.2.11 General Environmental Requirements for Oracle Exadata Database Machine X2-8 Full Rack 17
1.2.12 General Environmental Requirements for Oracle Exadata Storage Expansion Rack X5-2 18
1.2.13 General Environmental Requirements for Oracle Exadata Storage Expansion Rack X4-2 19
1.2.14 General Environmental Requirements for Oracle Exadata Storage Expansion Rack X3-2 21
1.2.15 General Environmental Requirements for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers 22

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

1.2.16 General Environmental Requirements for Single Servers 23
1.3 Space Requirements 29
1.3.1 Space Requirements for Racks up to Oracle Exadata X6 30
1.4 Receiving, Unpacking, and Access Route Requirements 30
1.4.1 Rack Weights for Oracle Exadata 31
1.4.2 Rack Weights for Oracle Exadata Storage Expansion Rack 32
1.5 Maintenance Access Requirements 33
1.6 Flooring Requirements 34
1.7 Electrical Power Requirements 34
1.7.1 PDU Power Requirements 35
1.7.1.1 Low-Voltage 15 kVA Single Phase PDU for North America, South America, Japan and Taiwan 36
1.7.1.2 Low-Voltage 22 kVA Single Phase PDU for North America, South America, Japan and Taiwan 36
1.7.1.3 Low-Voltage 15 kVA Three Phase PDU for North America, South America, Japan and Taiwan 37
1.7.1.4 Low-Voltage 24 kVA Three Phase PDU for North America, South America, Japan and Taiwan 37
1.7.1.5 High-Voltage 15 kVA Single Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan 38
1.7.1.6 High-Voltage 22 kVA Single Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan 39
1.7.1.7 High-Voltage 15 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan 39
1.7.1.8 High-Voltage 24 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan 40
1.7.2 Facility Power Requirements 40
1.7.3 Circuit Breaker Requirements 41
1.7.4 Electrical Grounding Guidelines 41
1.8 Temperature and Humidity Requirements 41
1.9 Ventilation and Cooling Requirements 43
1.10 Network Connection and IP Address Requirements for Oracle Exadata Rack 44
1.10.1 Network Connection Requirements for Oracle Exadata Rack 44
1.10.2 DNS Configuration for Oracle Exadata Rack 45
1.10.3 IP Address Requirements for Oracle Exadata Rack 46
1.10.4 Secure Fabric Requirements for Oracle Exadata 48

2 Understanding the Network Requirements for Oracle Exadata

2.1 Overview of Network Requirements 1
2.2 Network Channel Bonding Support 3
2.3 Network Partitioning on Oracle Exadata 4
2.3.1 VLAN Support on Customer-Facing Networks 4
2.3.2 Access VLAN Support with RoCE Network Fabric 5

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

2.3.3 Using Exadata Secure RDMA Fabric Isolation 6
2.3.4 Using InfiniBand Partitioning for Network Isolation with InfiniBand Network Fabric 8
2.4 IPv6 Support 9
2.5 Configuring a Separate Network for ILOM 10
2.6 Default IP Addresses 11
2.7 Default Port Assignments 11

3 Using Oracle Exadata Deployment Assistant

3.1 Overview of Oracle Exadata Deployment Assistant 1
3.1.1 Considerations and Requirement for Using OEDA 3
3.2 Getting Started with the OEDA Browser-Based User Interface 4
3.3 Using the Browser-based Version of Oracle Exadata Deployment Assistant 6

4 OEDA Command Line Interface

4.1 About the OEDA Command Line Interface 1
4.2 Running the OEDACLI Utility 4
4.2.1 Displaying Help Output for OEDACLI 5
4.2.2 Displaying Help Output for an OEDACLI Command 6
4.3 OEDACLI Command Reference 7
4.3.1 ACFSVOLUME 9
4.3.1.1 ADD ACFSVOLUME 9
4.3.1.2 ALTER ACFSVOLUME 10
4.3.1.3 DELETE ACFSVOLUME 10
4.3.1.4 LIST ACFSVOLUMES 11
4.3.2 ACTION 11
4.3.2.1 SAVE ACTION 12
4.3.2.2 MERGE ACTIONS 12
4.3.2.3 DEPLOY ACTIONS 12
4.3.3 ALERT 13
4.3.3.1 ALTER ALERT 13
4.3.3.2 DELETE ALERT 15
4.3.3.3 LIST ALERTS 15
4.3.4 CELL 16
4.3.4.1 CLONE CELL 16
4.3.4.2 CLONE NEWCELL 18
4.3.4.3 DELETE CELL 19
4.3.4.4 DELETE NEWCELL 19
4.3.4.5 LIST CELLS 20
4.3.5 CLUSTER 20
4.3.5.1 ALTER CLUSTER 21

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

4.3.5.2 CREATE CLUSTER 23
4.3.5.3 DESTROY CLUSTER 24
4.3.5.4 DOWNGRADE CLUSTER 25
4.3.5.5 LIST CLUSTER 26
4.3.5.6 LIST CLUSTERS 27
4.3.5.7 UPDATE CLUSTER 27
4.3.5.8 UPGRADE CLUSTER 28
4.3.6 COMPUTE 31
4.3.6.1 CLONE COMPUTE 31
4.3.6.2 DELETE COMPUTE 34
4.3.7 CURRENTACTION 35
4.3.7.1 CLEAR CURRENTACTION 35
4.3.7.2 LIST CURRENTACTION 36
4.3.8 DATABASE 36
4.3.8.1 ADD DATABASE 36
4.3.8.2 ALTER DATABASE 39
4.3.8.3 DELETE DATABASE 40
4.3.8.4 LIST DATABASES 42
4.3.9 DATABASEHOME 43
4.3.9.1 ADD DATABASEHOME 43
4.3.9.2 ALTER DATABASEHOME 45
4.3.9.3 DELETE DATABASEHOME 47
4.3.9.4 LIST DATABASEHOMES 48
4.3.10 DISKGROUP 48
4.3.10.1 ADD DISKGROUP 48
4.3.10.2 ALTER DISKGROUP 50
4.3.10.3 ALTER DISKGROUPS 53
4.3.10.4 DELETE DISKGROUP 54
4.3.10.5 LIST DISKGROUPS 55
4.3.11 DOM0 56
4.3.11.1 LIST DOM0S 56
4.3.12 DOMAIN 56
4.3.12.1 GET DOMAIN 56
4.3.13 ES 57
4.3.13.1 ALTER ES 57
4.3.13.2 DISCOVER ES 58
4.3.13.3 VALIDATE ES 58
4.3.14 EXASCALE 58
4.3.14.1 DEPLOY EXASCALE 59
4.3.15 EXASCALECLUSTER 59
4.3.15.1 ADD EXASCALECLUSTER 60
4.3.15.2 ALTER EXASCALECLUSTER 60

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

4.3.15.3 DELETE EXASCALECLUSTER 61
4.3.15.4 LIST EXASCALECLUSTER 62
4.3.15.5 LIST EXASCALECLUSTERS 62
4.3.16 EXITONERROR 62
4.3.16.1 SET EXITONERROR 62
4.3.17 FILE 63
4.3.17.1 LOAD FILE 63
4.3.17.2 MERGE FILES 63
4.3.17.3 SAVE FILE 63
4.3.17.4 SAVE FILES 64
4.3.18 GUEST 64
4.3.18.1 CLONE GUEST 65
4.3.18.2 CREATE GUEST 68
4.3.18.3 DELETE GUEST 70
4.3.18.4 MIGRATE GUEST 71
4.3.19 ILOM 75
4.3.19.1 ALTER ILOM 75
4.3.19.2 LIST ILOMS 76
4.3.20 LOGIN 76
4.3.20.1 CHECK LOGIN 76
4.3.21 MACHINE 77
4.3.21.1 ALTER MACHINE 77
4.3.21.2 ALTER MACHINES 80
4.3.21.3 LIST MACHINES 81
4.3.22 NETWORK 81
4.3.22.1 ADD NETWORK 82
4.3.22.2 ALTER NETWORK 83
4.3.22.3 ALTER NETWORKS 85
4.3.22.4 DELETE NETWORK 86
4.3.22.5 LIST NETWORKS 86
4.3.23 PASSWORD 87
4.3.23.1 ALTER PASSWORD 87
4.3.24 PROPERTY 88
4.3.24.1 ALTER PROPERTY 88
4.3.24.2 LIST PROPERTIES 88
4.3.25 RACK 89
4.3.25.1 ADD RACK 89
4.3.25.2 ALTER RACK 90
4.3.25.3 LIST RACKS 90
4.3.25.4 LIST RACKDESCRIPTION 91
4.3.26 REQUIREDFILES 91
4.3.26.1 LIST REQUIREDFILES 91

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

4.3.27 REMOTE 92
4.3.27.1 EXEC REMOTE 92
4.3.28 SCAN 92
4.3.28.1 ADD SCAN 92
4.3.28.2 ALTER SCAN 93
4.3.28.3 DELETE SCAN 93
4.3.28.4 LIST SCANS 94
4.3.29 SECURITY 94
4.3.29.1 SET SECURITY 95
4.3.30 SOFTWAREVERSIONS 95
4.3.30.1 LIST SOFTWAREVERSIONS 95
4.3.31 SSHKEYS 96
4.3.31.1 DEPLOY SSHKEYS 96
4.3.31.2 GENERATE SSHKEYS 97
4.3.31.3 SET SSHKEYS 98
4.3.32 STEPS 98
4.3.32.1 LIST STEPS 98
4.3.33 STORAGEPOOL 100
4.3.33.1 ADD STORAGEPOOL 100
4.3.33.2 ALTER STORAGEPOOL 101
4.3.33.3 DELETE STORAGEPOOL 102
4.3.33.4 LIST STORAGEPOOL 102
4.3.33.5 LIST STORAGEPOOLS 103
4.3.34 STORAGETEMPLATES 103
4.3.34.1 LIST STORAGETEMPLATES 103
4.3.35 SWITCH 103
4.3.35.1 ADD SWITCH 103
4.3.35.2 ALTER SWITCH 104
4.3.35.3 LIST SWITCHES 105
4.3.36 SU 105
4.3.36.1 SET SU 105
4.3.37 ULOC 106
4.3.37.1 LIST ULOC 106
4.3.38 VAULT 106
4.3.38.1 ADD VAULT 107
4.3.38.2 ALTER VAULT 107
4.3.38.3 DELETE VAULT 108
4.3.38.4 LIST VAULT 109
4.3.38.5 LIST VAULTS 109
4.3.39 VIP 109
4.3.39.1 ADD VIP 109
4.3.39.2 ALTER VIP 110

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

4.3.39.3 DELETE VIP 111
4.3.39.4 LIST VIPS 112
4.3.40 VOLUME 112
4.3.40.1 ADD VOLUME 112
4.3.40.2 DELETE VOLUME 113
4.3.40.3 LIST VOLUMES 113
4.3.41 XMLACTION 113
4.3.41.1 ALTER XMLACTION 114
4.3.41.2 DELETE XMLACTION 114
4.3.41.3 LIST XMLACTION 114
4.3.41.4 LIST XMLACTIONS 114

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

5.1 Reviewing Safety Guidelines 1
5.2 Unpacking and Installing Oracle Exadata Rack 2
5.3 Acclimating Oracle Exadata 3
5.4 Powering on the System the First Time 3
5.4.1 Inspecting the Machine After it is in Place 3
5.4.2 Connecting Power Cords 3
5.4.3 Powering on Oracle Exadata Rack 4

6 Configuring Oracle Exadata Database Machine

6.1 Verifying the Network Configuration Prior to Configuring the Rack 2
6.2 Configuring the RDMA Network Fabric Switch 3
6.2.1 Configuring the Cisco Nexus 9336C-FX2 Switch 3
6.2.2 Configuring the RoCE Network Fabric Switches to Enable Exadata Secure RDMA Fabric Isolation 9
6.2.3 Configuring Sun Datacenter InfiniBand Switch 36 Switch 9
6.3 Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack 12
6.4 Configuring the Management Network Switch 14
6.4.1 Configuring the Cisco Catalyst 4948 Ethernet Switch 14
6.4.2 Configuring the Cisco Nexus 93108-1G or 9348 Ethernet Switch 21
6.4.2.1 Performing the Initial Switch Configuration for the Cisco Nexus 93108-1G or 9348 Ethernet Switch 21
6.4.2.2 Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch 25
6.5 Configuring the Power Distribution Units 26
6.6 Configuring the Threshold Settings for the Power Distribution Units 28
6.6.1 PDU Thresholds for Oracle Exadata Database Machine X5-2 and Later 29
6.6.2 PDU Thresholds for Oracle Exadata Database Machine X4-2 29

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

6.6.2.1 PDU Thresholds for Oracle Exadata Database Machine X4-2 Full Rack 29
6.6.2.2 PDU Thresholds for Oracle Exadata Database Machine X4-2 Half Rack 31
6.6.2.3 PDU Thresholds for Oracle Exadata Database Machine X4-2 Quarter Rack 32
6.6.3 PDU Thresholds for Oracle Exadata Database Machine X3-2 34
6.6.3.1 PDU Thresholds for Oracle Exadata Database Machine X3-2 Full Rack 34
6.6.3.2 PDU Thresholds for Oracle Exadata Database Machine X3-2 Half Rack 36
6.6.3.3 PDU Thresholds for Oracle Exadata Database Machine X3-2 Quarter Rack 37
6.6.4 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) 39
6.6.4.1 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Full Rack 39
6.6.4.2 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Half Rack 41
6.6.4.3 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 M2 and X4270 M2 servers) Quarter Rack 42
6.6.5 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) 44
6.6.5.1 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Full Rack 44
6.6.5.2 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Half Rack 46
6.6.5.3 PDU Thresholds for Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) Quarter Rack 47
6.6.6 PDU Thresholds for Oracle Exadata Database Machine X4-8 and Later 49
6.6.7 PDU Thresholds for Oracle Exadata Database Machine X3-8 Full Rack 49
6.6.8 PDU Thresholds for Oracle Exadata Database Machine X2-8 Full Rack 51
6.6.9 PDU Thresholds for Oracle Exadata Storage Expansion Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers 53
6.6.9.1 PDU Thresholds for Oracle Exadata Storage Expansion Full Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers 53
6.6.9.2 PDU Thresholds for Oracle Exadata Storage Expansion Half Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers 55
6.6.9.3 PDU Thresholds for Oracle Exadata Storage Expansion Quarter Rack with Exadata Storage Server with Sun Fire X4270 M2 Servers 56
6.7 Checking Exadata Storage Servers 58
6.8 Checking Oracle Exadata Database Servers 61
6.9 Performing Additional Checks and Configuration 65
6.10 Verifying the RoCE Network Fabric Configuration 66
6.11 Verifying the InfiniBand Network Fabric Network 67
6.12 Customizing the Database Server File System Layout 68
6.12.1 Creating the Custom File System Layout 68
6.12.2 Applying the Custom File System Layout 69
6.13 Customizing the Guest Image 70

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

6.13.1 Customizing the Guest Image to Expand an Existing File System 71
6.13.2 Customizing the Guest Image to Add a Storage Volume and File System 83
6.14 Imaging a New System 96
6.14.1 Using Elastic Configuration to Image a New System 97
6.14.2 Using ISO to Image a New System 99
6.14.3 Using PXE to Image a New System 103
6.14.4 Updates to Imaging Parameters 107
6.15 Performing Initial Elastic Configuration of Oracle Exadata 108
6.16 Adding Additional Elastic Nodes to an Existing Rack 113
6.17 Using the OEDA genPasswordHash Utility 115
6.18 Using the OEDA setuprootssh Utility 115
6.19 Loading the Configuration Information and Installing the Software 116
6.19.1 Configuring Oracle Exadata Database Machine Using OEDA 117
6.19.2 Using Oracle Exadata Deployment Assistant in Conjunction with Key-Based Authentication 120
6.19.3 Using Oracle Exadata Deployment Assistant on Systems with Non-Default root Passwords 120
6.19.4 Deploying Oracle Exadata Using Enhanced Storage Server Security 121
6.20 Using the OEDA changePassword Utility 122
6.21 Installing Oracle Enterprise Manager Cloud Control 123
6.22 Adding a VM Cluster to Oracle Exadata Database Machine Using OEDA 123
6.23 Using IPv4 and IPv6 in an Existing Deployment 126
6.23.1 Adding Network Interfaces to an Existing Guest 126
6.23.2 Removing Network Interfaces from an Existing Guest 127

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-9
A.7 Logistics Checklist A-10
A.8 Network Configuration Checklist A-13
A.9 Auto Service Request Checklist A-14
A.10 Re-racking Checklists A-15
A.11 PDU Checklist for Re-racking A-18
A.12 Recycling the Racks A-19

B Database High Availability Checklist

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

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
- Related Documentation

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.

Related Documentation

The following guides contain additional information for Oracle Exadata:

Oracle Exadata System Guides

- *Oracle Exadata System Overview*
- *Oracle Exadata System Software User's Guide*
- *Oracle Exadata Maintenance Guide*
- *Oracle Exadata Security Guide*
- *Oracle Exadata Extending and Multi-Rack Cabling Guide*
- *Oracle Engineered System Safety and Compliance Guide*

Oracle Exadata Hardware Guides

- *Exadata Server X11M Service Manual* at https://docs.oracle.com/en/engineered-systems/servers/exadata/x11m/dbmhw/index.html
- *Exadata Server X10M Service Manual* at https://docs.oracle.com/en/engineered-systems/servers/exadata/x10m/dbmhw/index.html
- *Oracle Exadata X9-2 Database Server Service Manual* at https://docs.oracle.com/en/servers/x86/x9-2/exa-database/index.html
- *Oracle Exadata Storage Server X9-2 EF, HC, and XT Service Manual* at https://docs.oracle.com/en/servers/x86/x9-2l/exa-storage/index.html
- *Oracle Exadata X8-2 Database Server Service Manual* at https://docs.oracle.com/en/servers/x86/x8-2l/exadata-db-service-manual/
- *Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual* at https://docs.oracle.com/en/servers/x86/x8-2l/exadata-storage-service-manual/

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

- *Oracle Exadata Database Server X8-8 Service Manual* at https://docs.oracle.com/en/servers/x86/x8-8/exadata-storage-service-manual/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

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

- *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"

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

1

Site Requirements for Oracle Exadata and Oracle Exadata Storage Expansion Rack

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

- Ensuring That the Site is Ready
- General Environmental Requirements
- Space Requirements
- 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.
- 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 Rack

Related Topics

- Site Checklists

1.1 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.

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

Task 2 Understand the Networking Requirements and Options
Review Understanding the Network Requirements for Oracle Exadata to discern the networking requirements and configuration decisions that you must make.

Task 3 Run Oracle Exadata Deployment Assistant (OEDA)
Run OEDA to create the necessary configuration files.
See Using Oracle Exadata Deployment Assistant.

Task 4 Configure the Network
Configure the network using the files from OEDA. This includes registering the networks in the Domain Name System (DNS), assigning IP addresses, and configuring data center switches and firewalls.

Task 5 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.
Review the procedures in Installing Oracle Exadata or Oracle Exadata Storage Expansion Rack at the Site.

1.2 General Environmental Requirements

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

- General Environmental Requirements for Oracle Exadata Rack X6 and Later
- General Environmental Requirements for Oracle Exadata 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 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

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

1.2.1 General Environmental Requirements for Oracle Exadata Rack X6 and Later

Starting with X6 models of Oracle Exadata and Oracle Exadata Storage Expansion Rack, use Oracle Exadata Configuration Assistant (OECA) to determine environmental requirements, such as size, weight, acoustic level, power, cooling, and airflow. This includes all X6, X7, X8, X8M, X9M, and later models.

Related Topics

- Oracle Exadata Configuration Assistant (OECA) Downloads

1.2.2 General Environmental Requirements for Oracle Exadata X5-2

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

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

Environmental Component	Oracle Exadata X5-2 Full Rack	Oracle Exadata X5-2 Half Rack	Oracle Exadata X5-2 Quarter Rack	Oracle Exadata X5-2 Eighth Rack
Weight **See Also:** Flooring Requirements	• Extreme Flash drives: 804.9 kg (1774 lbs) • High capacity drives: 874.8 kg (1928 lbs))	• Extreme Flash drives: 525.9 kg (1160 lbs) • High capacity drives: 560.9 kg (1236 lbs)	• Extreme Flash drives: 366.3 kg (808 lbs) • High capacity drives: 381.2 kg (841 lbs)	• 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	• 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)	• 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)	• 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)	• 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)

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

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

Environmental Component	Oracle Exadata X5-2 Full Rack	Oracle Exadata X5-2 Half Rack	Oracle Exadata X5-2 Quarter Rack	Oracle Exadata X5-2 Eighth Rack
Cooling **See Also:** Temperature and Humidity Requirements, and Ventilation and Cooling Requirements	• 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)	• 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)	• 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)	• 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)
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 Maximum: Approximately 2,434 CFM Typical: Approximately 1,704 CFM • High capacity drives Maximum: Approximately 2,410 CFM Typical: Approximately 1,687 CFM	• Extreme Flash drives Maximum: Approximately 1,315 CFM Typical: Approximately 920 CFM • High capacity drives Maximum: Approximately 1,302 CFM Typical: Approximately 912 CFM	• Extreme Flash drives Maximum: Approximately 657 CFM Typical: Approximately 460 CFM • High capacity drives Maximum: Approximately 652 CFM Typical: Approximately 456 CFM	• Extreme Flash drives Maximum: Approximately 540 CFM Typical: Approximately 378 CFM • High capacity drives Maximum: Approximately 534 CFM Typical: Approximately 374 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 RDMA Network Fabric	38 for Ethernet network, assuming single cluster Up to 18 for RDMA Network Fabric	22 for Ethernet network, assuming single cluster Up to 8 for RDMA Network Fabric	22 for Ethernet network, assuming single cluster Up to 8 for RDMA Network Fabric

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

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

Environmental Component	Oracle Exadata X5-2 Full Rack	Oracle Exadata X5-2 Half Rack	Oracle Exadata X5-2 Quarter Rack	Oracle Exadata 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

1.2.3 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: • Max: 16.4 kW (16.7 kVA) • Typical: 11.5 kW (11.7 kVA) High capacity drives: • Max: 16.2 kW (16.6 kVA) • Typical: 11.4 kW (11.6 kVA)	Extreme Flash drives: • Max: 9.5 kW (9.7 kVA) • Typical: 6.7 kW (6.8 kVA) High capacity drives: • Max: 9.5 kW (9.7 kVA) • Typical: 6.6 kW (6.8 kVA)

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

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: • Max: 55,935 BTU/hour (59,012 kJ/hour) • Typical: 39,155 BTU/hour (41,308 kJ/hour) High capacity drives: • Max: 55,410 BTU/hour (58,457 kJ/hour) • Typical: 38,787 BTU/hour (40,920 kJ/hour)	Extreme Flash drives: • Max: 32,436 BTU/hour (34,220 kJ/hour) • Typical: 22,705 BTU/hour (23,954 kJ/hour) High capacity drives: • 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: • Max: 2590 CFM • Typical: 1813 CFM High capacity drives: • Max: 2565 CFM • Typical: 1796 CFM	Extreme Flash drives: • Max: 1502 CFM • Typical: 1051 CFM High capacity drives: • Max: 1496 CFM • Typical: 1048 CFM

1.2.4 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)

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

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"	• 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)	• 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)	• 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)	• 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"	• 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)	• 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)	• 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)	• 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)

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

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"	• 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	• High performance drives Maximum: Approximately 1,050 CFM Typical: Approximately 750 CFM • High capacity drives Maximum: Approximately 980 CFM Typical: Approximately 680 CFM	• High performance drives Maximum: Approximately 520 CFM Typical: Approximately 365 CFM • High capacity drives Maximum: Approximately 490 CFM Typical: Approximately 350 CFM	• 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

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

1.2.5 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"	• Extreme Flash drives: 828.1 kg (1826 lbs) • High capacity drives: 898.3 kg (1980 lbs))	• 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"	• 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)	• 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"	• 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)	• 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"	• 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	• 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

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

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

1.2.6 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"	• 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)

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

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"	• 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"	• 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

1.2.7 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)

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

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
Power **See Also:** "Electrical Power Requirements"	• 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)	• 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)	• 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)	• 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"	• 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)	• 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)	• 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)	• 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)

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

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"	• 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	• High performance drives Maximum: Approximately 1,050 CFM Typical: Approximately 750 CFM • High capacity drives Maximum: Approximately 950 CFM Typical: Approximately 670 CFM	• High performance drives Maximum: Approximately 500 CFM Typical: Approximately 375 CFM • High capacity drives Maximum: Approximately 470 CFM Typical: Approximately 330 CFM	• 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
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

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

1.2.8 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"	• 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)
Cooling **See Also:** "Temperature and Humidity Requirements", and "Ventilation and Cooling Requirements"	• 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"	• 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

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

1.2.9 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"	• 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"	• 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"	• 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

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

1.2.10 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
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

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

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
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

1.2.11 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)
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

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

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)
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

1.2.12 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	• Extreme Flash drives: 821 kg (1810 lbs) • High capacity drives: 915.8 kg (2019 lbs)	• Extreme Flash drives: 546.6 kg (1205 lbs) • High capacity drives: 591.5 kg (1304 lbs)	• 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	• 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)	• 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)	• 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)

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

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
Cooling	• 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)	• 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	• 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)
Air flow front-to-back (subject to actual data center environment)	• 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	• Extreme Flash drives Maximum: Approximately 920 CFM Typical: Approximately 645 CFM • High capacity drives Maximum: Approximately 1,065 CFM Typical: Approximately 745 CFM	• 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

1.2.13 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.

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

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)
Power	• 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)	• 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)	• 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	• 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)	• 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)	• 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)	• 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	• High performance drives Maximum: Approximately 900 CFM Typical: Approximately 650 CFM • High capacity drives Maximum: Approximately 810 CFM Typical: Approximately 570 CFM	• High performance drives Maximum: Approximately 440 CFM Typical: Approximately 310 CFM • High capacity drives Maximum: Approximately 410 CFM Typical: Approximately 300 CFM
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

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

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
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

1.2.14 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	• 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)	• 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)	• 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)
Cooling	• 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)	• 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)	• 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)

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

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
Air flow front-to-back (subject to actual data center environment)	• 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	• High performance drives Maximum: Approximately 900 CFM Typical: Approximately 650 CFM • High capacity drives Maximum: Approximately 790 CFM Typical: Approximately 550 CFM	• 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

1.2.15 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)

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

Table 1-14 (Cont.) 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
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

1.2.16 General Environmental Requirements for Single Servers

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

Table 1-15 Environmental Requirements for Oracle Exadata X11M Single Servers

Environmental Component	Exadata Database Server X11M	Exadata Database Server X11M-Z	Exadata Database Server X11M - 3 TB RAM Configuration	Exadata Storage Server X11M - High Capacity (X11M HC)	Exadata Storage Server X11M - Extreme Flash (X11M EF)	Exadata Storage Server X11M-Z - High Capacity (X11M HC-Z)	Exadata Storage Server X11 - High Capacity (X11 HC)	Exadata Storage Server X11 - Extreme Flash (X11 EF)	Exadata Storage Server X11-Z - High Capacity (X11 HC-Z)
Weight	24 kg (53 lbs)	24 kg (53 lbs)	24 kg (53 lbs)	33.6 kg (74 lbs)	27.2 kg (60 lbs)	28.6 kg (63 lbs)	33.6 kg (74 lbs)	27.2 kg (60 lbs)	28.6 kg (63 lbs)
Acoustic Noise	86 dB	85 dB	90 dB	86 dB	80 dB	86 dB	86 dB	80 dB	86 dB

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

Table 1-15 (Cont.) Environmental Requirements for Oracle Exadata X11M Single Servers

Environmental Component	Exadata Database Server X11M	Exadata Database Server X11M-Z	Exadata Database Server X11M - 3 TB RAM Configuration	Exadata Storage Server X11M - High Capacity (X11M HC)	Exadata Storage Server X11M - Extreme Flash (X11M EF)	Exadata Storage Server X11M-Z - High Capacity (X11M HC-Z)	Exadata Storage Server X11 - High Capacity (X11 HC)	Exadata Storage Server X11 - Extreme Flash (X11 EF)	Exadata Storage Server X11-Z - High Capacity (X11 HC-Z)
Power	Maximum: 1.2 kW (1.2 kVA) Typical: 0.8 kW (0.8 kVA) (varies based on application load)	Maximum: 1.1 kW (1.1 kVA) Typical: 0.8 kW (0.8 kVA) (varies based on application load)	Maximum: 1.4 kW (1.4 kVA) Typical: 1 kW (1 kVA) (varies based on application load)	Maximum: 1 kW (1 kVA) Typical: 0.7 kW (0.7 kVA) (varies based on application load)	Maximum: 0.9 kW (0.9 kVA) Typical: 0.6 kW (0.6 kVA) (varies based on application load)	Maximum: 0.8 kW (0.8 kVA) Typical: 0.5 kW (0.5 kVA) (varies based on application load)	Maximum: 0.9 kW (0.9 kVA) Typical: 0.6 kW (0.7 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: 0.7 kW (0.7 kVA) Typical: 0.5 kW (0.5 kVA) (varies based on application load)
Cooling	At max usage: 4,054 BTU/hour (4,277 kJ/hour) At typical usage: 2,838 BTU/hour (2,994 kJ/hour)	At max usage: 3,783 BTU/hour (3,991 kJ/hour) At typical usage: 2,648 BTU/hour (2,794 kJ/hour)	At max usage: 4,705 BTU/hour (4,964 kJ/hour) At typical usage: 3,294 BTU/hour (3,475 kJ/hour)	At max usage: 3,474 BTU/hour (3,665 kJ/hour) At typical usage: 2,431 BTU/hour (2,565 kJ/hour)	At max usage: 3,037 BTU/hour (3,204 kJ/hour) At typical usage: 2,126 BTU/hour (2,243 kJ/hour)	At max usage: 2,570 BTU/hour (2,712 kJ/hour) At typical usage: 1,799 BTU/hour (1,898 kJ/hour)	At max usage: 3,115 BTU/hour (3,287 kJ/hour) At typical usage: 2,181 BTU/hour (2,301 kJ/hour)	At max usage: 2,723 BTU/hour (2,873 kJ/hour) At typical usage: 1,906 BTU/hour (2,011 kJ/hour)	At max usage: 2,305 BTU/hour (2,432 kJ/hour) At typical usage: 1,614 BTU/hour (1,702 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	At max usage: Approximately 188 CFM At typical usage: Approximately 131 CFM	At max usage: Approximately 175 CFM At typical usage: Approximately 123 CFM	At max usage: Approximately 218 CFM At typical usage: Approximately 152 CFM	At max usage: Approximately 161 CFM At typical usage: Approximately 113 CFM	At max usage: Approximately 141 CFM At typical usage: Approximately 98 CFM	At max usage: Approximately 119 CFM At typical usage: Approximately 83 CFM	At max usage: Approximately 144 CFM At typical usage: Approximately 101 CFM	At max usage: Approximately 126 CFM At typical usage: Approximately 88 CFM	At max usage: Approximately 107 CFM At typical usage: Approximately 75 CFM

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

Table 1-16 Environmental Requirements for Oracle Exadata X10M Single Servers

Environmental Component	Database Server X10M	Database Server X10M - 3 TB RAM Configuration	Exadata Storage Server X10M - High Capacity	Exadata Storage Server X10M Eighth Rack - High Capacity	Exadata Storage Server X10M - Extreme Flash	Exadata Storage Server X10M XT Server
Weight	23.6 kg (52 lbs)	24 kg (53 lbs)	33.6 kg (74 lbs)	28.6 kg (63 lbs)	27.2 kg (60 lbs)	30.8 kg (68 lbs)
Acoustic Noise	84 dB	86 dB	84 dB	84 dB	82 dB	81 dB
Power	Maximum: 1.3 kW (1.3 kVA) Typical: 0.9 kW (0.9 kVA) (varies based on application load)	Maximum: 1.5 kW (1.5 kVA) Typical: 1 kW (1.1 kVA) (varies based on application load)	Maximum: 1 kW (1.1 kVA) Typical: 0.7 kW (0.7 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: 1 kW (1 kVA) Typical: 0.7 kW (0.7 kVA) (varies based on application load)	Maximum: 0.4 kW (0.5 kVA) Typical: 0.3 kW (0.3 kVA) (varies based on application load)
Cooling	At max usage: 4,299 BTU/hour (4,536 kJ/hour) At typical usage: 3,010 BTU/hour (3,175 kJ/hour)	At max usage: 5,050 BTU/hour (5,328 kJ/hour) At typical usage: 3,535 BTU/hour (3,729 kJ/hour)	At max usage: 3,559 BTU/hour (3,755 kJ/hour) At typical usage: 2,491 BTU/hour (2,628 kJ/hour)	At max usage: 2,634 BTU/hour (2,778 kJ/hour) At typical usage: 1,843 BTU/hour (1,945 kJ/hour)	At max usage: 3,337 BTU/hour (3,521 kJ/hour) At typical usage: 2,336 BTU/hour (2,464 kJ/hour)	At max usage: 1,512 BTU/hour (1,595 kJ/hour) At typical usage: 1,058 BTU/hour (1,116 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	At max usage: Approximately 199 CFM At typical usage: Approximately 139 CFM	At max usage: Approximately 234 CFM At typical usage: Approximately 164 CFM	At max usage: Approximately 165 CFM At typical usage: Approximately 115 CFM	At max usage: Approximately 122 CFM At typical usage: Approximately 85 CFM	At max usage: Approximately 154 CFM At typical usage: Approximately 108 CFM	At max usage: Approximately 70 CFM At typical usage: Approximately 49 CFM

Table 1-17 Environmental Requirements for X9M-2 Single Servers

Environmental Component	Database Server X9M-2	Exadata Storage Server X9M-2 - High Capacity	Exadata Storage Server X9M-2 Eighth Rack - High Capacity	Exadata Storage Server X9M-2 - Extreme Flash	Exadata Storage Server X9M-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 Noise	79 dB	80 dB	80 dB	80 dB	80 dB

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

Table 1-17 (Cont.) Environmental Requirements for X9M-2 Single Servers

Environmental Component	Database Server X9M-2	Exadata Storage Server X9M-2 - High Capacity	Exadata Storage Server X9M-2 Eighth Rack - High Capacity	Exadata Storage Server X9M-2 - Extreme Flash	Exadata Storage Server X9M-2 XT Server
Power	Maximum: 0.9 kW (0.9 kVA) Typical: 0.7 kW (0.7 kVA) (varies based on application load)	Maximum: 0.9 kW (0.9 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: 0.9 kW (0.9 kVA) Typical: 0.6 kW (0.6 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)
Cooling	At max usage: 3,153 BTU/hour (3,326 kJ/hour) At typical usage: 2,207 BTU/hour (2,328 kJ/hour	At max usage: 2,883 BTU/hour (3,042 kJ/hour) At typical usage: 2,018 BTU/hour (2,129 kJ/hour)	At max usage: 2,134 BTU/hour (2,251 kJ/hour) At typical usage: 1,494 BTU/hour (1,576 kJ/hour)	At max usage: 2,907 BTU/hour (3,067 kJ/hour) At typical usage: 2,035 BTU/hour (2,147 kJ/hour)	At max usage: 1,720 BTU/hour (1,814 kJ/hour) At typical usage: 1,204 BTU/hour (1,270 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	At max usage: Approximately 146 CFM At typical usage: Approximately 102 CFM	At max usage: Approximately 133 CFM At typical usage: Approximately 93 CFM	At max usage: Approximately 99 CFM At typical usage: Approximately 69 CFM	At max usage: Approximately 135 CFM At typical usage: Approximately 94 CFM	At max usage: Approximately 80 CFM At typical usage: Approximately 56 CFM

Table 1-18 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 Noise	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)

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

Table 1-18 (Cont.) 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
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 1-19 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 Noise	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: 0.7 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

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

Table 1-20 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 Noise	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)
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 1-21 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 Noise	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

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

Table 1-22 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 Noise	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: 0.7 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 1-23 Environmental Requirements for 8-Socket Servers, Model X4-8 and Later

Environmental Component	Database Server X8M-8 and X9M-8	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)	90.2 kg (198.8 lbs)	93.4 kg (206 lbs)	93.4 kg (206 lbs)
Acoustic Noise	87 dB	87 dB	81 dB(A)	81 dB(A)
Power	Maximum: 3.6 kVA Typical: 2.5 kVA (varies based on application load)	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: 11,953 BTU/hour (12,610 kJ/hour) Typical: 8,367 BTU/hour (8,827 kJ/hour)	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 553 CFM Typical: Approximately 387 CFM	Maximum: Approximately 537 CFM Typical: Approximately 376 CFM	Maximum: Approximately 569 CFM Typical: Approximately 398 CFM

1.3 Space Requirements

For all Oracle Exadata X7, X8, X8M, X9M, and later models, use Oracle Exadata Configuration Assistant (OECA) to determine the size of each rack.

In addition to the rack height, 914 mm (36 inches) of space is required above the rack height for maintenance access, as described in Table 1-25.

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

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.

- Space Requirements for Racks up to Oracle Exadata X6
  All racks up to Oracle Exadata X6-2 and X6-8 use the same hardware rack, and have the same space requirements.

Related Topics

- Oracle Exadata Configuration Assistant (OECA) Downloads

1.3.1 Space Requirements for Racks up to Oracle Exadata X6

All racks up to Oracle Exadata 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 1-25. 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.

1.4 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.

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

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

Table 1-24 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
  Review the weight of your Oracle Exadata Rack to ensure that it can be delivered and located safely.
- Rack Weights for Oracle Exadata Storage Expansion Rack

1.4.1 Rack Weights for Oracle Exadata

Review the weight of your Oracle Exadata Rack to ensure that it can be delivered and located safely.

For all Oracle Exadata X6, X7, X8, X8M, X9M, and later models, use Oracle Exadata Configuration Assistant (OECA) to determine the net rack weight. In addition to the net rack weight, allow approximately 180 kg (400 lbs) for packaging materials.

For earlier models, use the following table to determine the rack weight:

Type of Rack	Shipping Weight	Net Weight
Oracle Exadata X5-2 Full Rack	984.8 kg (2171 lbs)	901.3 kg (1987 lbs)
Oracle Exadata X5-2 Half Rack	639.6 kg (1410 lbs)	587.4 kg (1295 lbs)
Oracle Exadata X5-2 Quarter Rack	453.1 kg (999 lbs)	418.2 kg (922 lbs)
Oracle Exadata X5-2 Eighth Rack	453.1 kg (999 lbs)	418.2 kg (922 lbs)
Oracle Exadata X5-8 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 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 X4-2 Full Rack	936 kg (2064 lbs)	852.8 kg (1880 lbs)
Oracle Exadata X4-2 Half Rack	577 kg (1273 lbs)	525.3 kg (1158 lbs)

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

Type of Rack	Shipping Weight	Net Weight
Oracle Exadata X4-2 Quarter Rack	401 kg (883 lbs)	365.6 kg (806 lbs)
Oracle Exadata X4-2 Eighth Rack	401 kg (883 lbs)	365.6 kg (806 lbs)
Oracle Exadata 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 X3-2 Full Rack	954.8 kg (2105 lbs)	871.4 kg (1921 lbs)
Oracle Exadata X3-2 Half Rack	596.0 kg (1314 lbs)	543.9 kg (1199 lbs)
Oracle Exadata X3-2 Quarter Rack	403.7 kg (890 lbs)	368.8 kg (813 lbs)
Oracle Exadata X3-2 Eighth Rack	403.7 kg (890 lbs)	368.8 kg (813 lbs)
Oracle Exadata X3-8 Full Racks with Exadata Storage Server X4-2L Servers	994.7 kg (2193 lbs)	911.3 kg (2009 lbs)
Oracle Exadata X3-8 Full Racks	1013.3 kg (2234 lbs)	929.9 kg (2050 lbs)
Oracle Exadata X2-2 Full Rack	1050.1 kg (2315 lbs)	966.61 kg (2131 lbs)
Oracle Exadata X2-2 Half Rack	684.1 kg (1508 lbs)	602.8 kg (1329 lbs)
Oracle Exadata X2-2 Quarter Rack	487.6 kg (1075 lbs)	409.1 kg (902 lbs)
Oracle Exadata X2-8 Full Rack with Sun Fire X4800 Oracle Database Servers	1026.9 kg (2264 lbs)	943.8 kg (2080 lbs)
Oracle Exadata X2-8 Full Rack with Sun Server X2-8 Oracle Database Servers	1036.0 kg (2284 lbs)	980.7 kg (2162 lbs)

Related Topics

- Oracle Exadata Configuration Assistant (OECA) Downloads

1.4.2 Rack Weights for Oracle Exadata Storage Expansion Rack

For all Oracle Exadata Storage Expansion Rack X6-2, X7-2, X8-2, X8M-2, and later models, use Oracle Exadata Configuration Assistant (OECA) to determine the net rack weight. In addition to the net rack weight, allow approximately 180 kg (400 lbs) for packaging materials.

For earlier models, use the following table to determine the rack weight:

Type of Rack	Shipping Weight	Net Weight
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)

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

Type of Rack	Shipping Weight	Net Weight
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)
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)

Related Topics

- Oracle Exadata Configuration Assistant (OECA) Downloads

1.5 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.

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

Table 1-25 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.

1.6 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 1-26 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

If you want to route data or power cables through the bottom of the rack, you can cut out a rectangular hole in the floor below the rear portion of the rack. However, avoid creating a hole that impacts the rack casters or leveling feet. For details, see Floor Cutout Dimensions in *Oracle Rack Cabinet 1242 User's Guide*.

1.7 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

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

- 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 Oracle Exadata Rack configuration contains two power distribution units (PDUs). Various PDUs are available to cater for different circumstances. Select the correct PDU for your data center and Oracle Exadata Rack configuration.
- 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.

1.7.1 PDU Power Requirements

Each Oracle Exadata Rack configuration contains two power distribution units (PDUs). Various PDUs are available to cater for different circumstances. Select the correct PDU for your data center and Oracle Exadata Rack configuration.

The following are the PDUs used in Oracle Exadata:

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

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

1.7.1.1 Low-Voltage 15 kVA Single Phase PDU for North America, South America, Japan and Taiwan

> **Note**
>
> The following specifications are for a single PDU. Every Oracle Exadata Rack configuration contains two PDUs.

Table 1-27 Specifications for each Low-Voltage 15 kVA Single Phase PDU for North America, South America, Japan and Taiwan

Specification	Value
Number of Inputs	3
Source Voltage	200-240V
Frequency	50/60 Hz
Max Line Current	24A
Phase Current	N/A
Nominal Power Rating	15kVA
Recommended Circuit Breaker	30A
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Max Group Current	20A
Output Voltage	200-240V
Data Center Receptacle	1ph 2W + ground 30A NEMA L6-30

1.7.1.2 Low-Voltage 22 kVA Single Phase PDU for North America, South America, Japan and Taiwan

> **Note**
>
> The following specifications are for a single PDU. Every Oracle Exadata Rack configuration contains two PDUs.

Table 1-28 Specifications for each Low-Voltage 22 kVA Single Phase PDU for North America, South America, Japan and Taiwan

Specification	Value
Number of Inputs	3
Source Voltage	200-240V
Frequency	50/60 Hz
Max Line Current	36.8A
Phase Current	N/A
Nominal Power Rating	22kVA
Recommended Circuit Breaker	50A

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

Table 1-28 (Cont.) Specifications for each Low-Voltage 22 kVA Single Phase PDU for North America, South America, Japan and Taiwan

Specification	Value
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Max Group Current	20A
Output Voltage	200-240V
Data Center Receptacles	Hubbell CS8269 or CS8264

1.7.1.3 Low-Voltage 15 kVA Three Phase PDU for North America, South America, Japan and Taiwan

> **Note**
>
> The following specifications are for a single PDU. Every Oracle Exadata Rack configuration contains two PDUs.

Table 1-29 Specifications for each Low-Voltage 15 kVA Three Phase PDU for North America, South America, Japan and Taiwan

Specification	Value
Number of Inputs	1
Source Voltage	200-220V
Frequency	50/60 Hz
Max Line Current	40A
Phase Current	23A
Nominal Power Rating	15kVA
Recommended Circuit Breaker	50A
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Max Group Current	20A
Output Voltage	200-240V
Data Center Receptacle	3ph 3W + ground 60A IEC60309

1.7.1.4 Low-Voltage 24 kVA Three Phase PDU for North America, South America, Japan and Taiwan

> **Note**
>
> The following specifications are for a single PDU. Every Oracle Exadata Rack configuration contains two PDUs.

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

Table 1-30 Specifications for each Low-Voltage 24 kVA Three Phase PDU for North America, South America, Japan and Taiwan

Specification	Value
Number of Inputs	2
Source Voltage	200-220V
Frequency	50/60 Hz
Max Line Current	34.6A
Phase Current	20A
Nominal Power Rating	24kVA
Recommended Circuit Breaker	50A
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Max Group Current	20A
Output Voltage	200-240V
Data Center Receptacle	3ph 3W + ground 60A IEC60309

1.7.1.5 High-Voltage 15 kVA Single Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

> **Note**
>
> The following specifications are for a single PDU. Every Oracle Exadata Rack configuration contains two PDUs.

Table 1-31 Specifications for each High-Voltage 15 kVA Single Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Value
Number of Inputs	3 x 25A 1ph
Voltage	220 - 240 VAC
Frequency	50/60 Hz
Current	25A maximum (per input)
Power Rating	15 kVA
Output Current	72A (3 x 24A)
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Group Protection (UL489 1-pole circuit breaker)	20A
Data Center Receptacle	IEC309 32A 3 PIN 250 VAC IP44

> **Note**
>
> The high-voltage 15 kVA single phase PDU is no longer available. The high-voltage 22kVA single phase PDU is the recommended replacement.

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

1.7.1.6 High-Voltage 22 kVA Single Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

> **Note**
>
> The following specifications are for a single PDU. Every Oracle Exadata Rack configuration contains two PDUs.

Table 1-32 Specifications for each High-Voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Value
Number of Inputs	3
Source Voltage	200-240V
Frequency	50/60 Hz
Max Line Current	32A
Phase Current	N/A
Nominal Power Rating	22kVA
Recommended Circuit Breaker	30A
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Max Group Current	20A
Output Voltage	200-240V
Data Center Receptacle	1ph 2W + ground 32A IEC60309 3P 6H 230V

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

> **Note**
>
> The following specifications are for a single PDU. Every Oracle Exadata Rack configuration contains two PDUs.

Table 1-33 Specifications for each High-Voltage 15 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Value
Number of Inputs	1
Source Voltage	220/380 - 240/415VAC 3Ø
Frequency	50/60 Hz
Max Line Current	21A
Phase Current	21A

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

Table 1-33 (Cont.) Specifications for each High-Voltage 15 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Value
Nominal Power Rating	15kVA
Recommended Circuit Breaker	25A
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Max Group Current	20A
Output Voltage	220-240V
Data Center Receptacle	3ph 4W + ground 32A IEC60309 5P 6h 400V

1.7.1.8 High-Voltage 24 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

> **Note**
>
> The following specifications are for a single PDU. Every Oracle Exadata Rack configuration contains two PDUs.

Table 1-34 Specifications for each High-Voltage 24 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

Specification	Value
Number of Inputs	2
Source Voltage	220/380 - 240/415VAC 3Ø
Frequency	50/60 Hz
Max Line Current	16A
Phase Current	16A
Nominal Power Rating	24kVA
Recommended Circuit Breaker	25A
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Max Group Current	18A
Output Voltage	220-240V
Data Center Receptacle	3ph 4W + ground 32A IEC60309 5P 6h 400V

1.7.2 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

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

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.

1.7.3 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.

1.7.4 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.

1.8 Temperature and Humidity Requirements

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

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

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

> **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 1-35 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
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

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

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.

1.9 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.

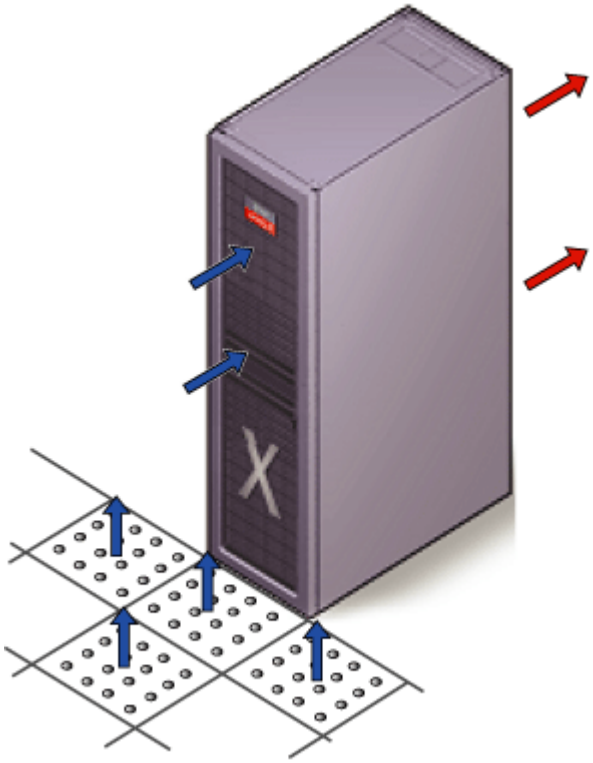
Use perforated tiles, approximately 400 CFM/tile, in front of the rack for cold air intake. The tiles can be arranged in any order in front of the rack, as long as cold air from the tiles can flow into the rack. Inadequate cold air flow could result in a higher inlet temperature in the servers due to 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 1-1 shows a typical installation of the floor tiles for Oracle Exadata Database Machine Full Rack in a typical data center.

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

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



1.10 Network Connection and IP Address Requirements for Oracle Exadata Rack

This section describes the requirements to connect Oracle Exadata Database Machine or Oracle Exadata Storage Expansion Rack to your existing network infrastructure.

- Network Connection Requirements for Oracle Exadata Rack
- DNS Configuration for Oracle Exadata Rack
- IP Address Requirements for Oracle Exadata Rack
- Secure Fabric Requirements for Oracle Exadata

1.10.1 Network Connection Requirements for Oracle Exadata Rack

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

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

Network	Connection Requirements
Administration Network	• Mandatory: 1 uplink for the Management Network Switch. • Recommended: 2 uplinks for remote monitoring of the power distribution units (PDUs). • Optional: 1 uplink for ILOM on each database server or storage server if using a separate ILOM network. • For X2 models only: 1 uplink for the KVM switch.
Client Network	• For bonded network configurations: 2 uplinks for each database server. • For non-bonded network configurations: 1 uplink for each database server. **Note:** Non-bonded network configurations are not supported on Oracle Exadata X7 and later systems.
Additional Network(s)	• For bonded network configurations: 2 uplinks for each database server and additional network. • For non-bonded network configurations: 1 uplink for each database server and additional network. **Note:** Non-bonded network configurations are not supported on Oracle Exadata X7 and later systems.
RDMA Network Fabric (Private Network)	No uplinks are required. The RDMA Network Fabric is completely self-contained within Oracle Exadata.

The same requirements apply to Oracle Exadata Storage Expansion Rack. However, since Oracle Exadata Storage Expansion Rack contains no database servers, there are no corresponding database server uplink requirements.

Related Topics

- Understanding the Network Requirements for Oracle Exadata
  Review the network requirements for Oracle Exadata before installing or configuring the hardware.

1.10.2 DNS Configuration for Oracle Exadata Rack

Prior to receiving your Oracle Exadata Rack, you must use Oracle Exadata Deployment Assistant (OEDA) to generate files that drive the configuration of the system. The system configuration files contain network host names and IP addresses for the management network, client network, and additional public networks, which should be registered in Domain Name System (DNS) prior to initial configuration. In particular, 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.

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

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

During configuration using OEDA, Grid Naming Service (GNS) is not configured on Oracle Exadata Database Machine.

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

1.10.3 IP Address Requirements for Oracle Exadata Rack

Oracle Exadata Database Machine and Oracle Exadata Storage Expansion Rack require a dedicated allocation of host names and IP addresses. The number of required IP addresses primarily depends on the size of the system and whether it is configured to use virtual machines (VMs).

The detailed network configuration, including host names and IP addresses, is generated from information supplied to Oracle Exadata Deployment Assistant (OEDA).

You should configure the IP addresses in your existing networks after you run OEDA. All IP addresses in the configuration 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 and the internal network.

You can use the information in the following sections to calculate the IP address requirement for Oracle Exadata. However, ensure that you use the configuration information generated in OEDA as your authoritative guide.

Rack-Level IP Address Requirements

The following table outlines the rack-level IP address requirements for Oracle Exadata Database Machine. These addresses are required regardless of the configuration options that are selected for the system. For Oracle Exadata Storage Expansion Rack, no further IP addresses are required beyond these rack-level IP addresses.

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

Network	IP Address Requirements
Administration Network	• 1 IP address for the administration network interface on each database server or storage server. • 1 IP address for the ILOM network interface on each database server or storage server. **Note:** By default, the database server and storage server ILOM interfaces reside in the administration network. However, starting with Oracle Exadata System Software release 19.1.0, the ILOM interfaces can reside in a separate management network. Either way, the number of required IP addresses remains the same. • 1 IP address for the management interface on each RDMA Network Fabric switch. • 1 IP address for the management interface on the Management Network Switch switch. • 1 IP address for the management interface in each power distribution unit (PDU). • For X2 models only: 1 IP address for the KVM switch.
RDMA Network Fabric (Private Network)	2 IP addresses for each database server or storage server. **Note:** Older configurations (typically X3 and prior) may be configured with Linux Ethernet bonding, which uses 1 IP address for each database server or storage server.

Bare-Metal IP Address Requirements

When Oracle Exadata Database Machine is configured without VMs (otherwise known as a bare-metal configuration), one additional set of cluster-level IP addresses is required in addition to the rack-level addresses outlined in the previous table. Use the following table to determine the additional IP addresses required for a bare-metal configuration.

Network	IP Address Requirements
Client Network	• 1 IP address for the client network interface on each database server. • 1 IP address for the virtual (VIP) network interface on each database server. • 3 IP addresses for SCAN network services.
Additional Network(s)	1 IP address for each additional network interface on each database server.

VM Cluster IP Address Requirements

When Oracle Exadata Database Machine is configured with virtual machines (VMs), each VM cluster requires a separate set of IP addresses, which are in addition to the system-level addresses outlined previously. Use the following table to determine the additional IP addresses required for each VM cluster.

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

Network	IP Address Requirements
Administration Network	1 IP address for the administration network interface on each database server VM.
Client Network	• 1 IP address for the client network interface on each database server VM. • 1 IP address for the virtual (VIP) network interface on each database server VM. • 3 IP addresses for SCAN network services.
Additional Network(s)	1 IP address for each additional network interface on each database server VM.
RDMA Network Fabric (Private Network)	2 IP addresses for each database server VM.

Related Topics

- Single Client Access Name (SCAN)
- Using Oracle Exadata Deployment Assistant

1.10.4 Secure Fabric Requirements for Oracle Exadata

Starting in October 2024, Exadata Secure RDMA Fabric Isolation is recommended by default on all new Oracle Exadata configurations.

Exadata Secure RDMA Fabric Isolation enables strict network isolation for virtual machine (VM) clusters on Oracle Exadata systems that use RDMA over Converged Ethernet (RoCE).

Secure Fabric provides critical infrastructure for secure consolidation of multiple tenants on Oracle Exadata, where each tenant resides in a dedicated 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.

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 OEDA.

   See Configuring the RoCE Network Fabric Switches to Enable Exadata Secure RDMA Fabric Isolation.

2. As part of initial system deployment using OEDA, select the option to enable Secure Fabric and specify VLAN IDs for the cluster and storage network partitions associated with each VM cluster.

   In the OEDA Web user interface, the option to enable Secure Fabric is one of the advanced options associated with the **Cluster Networks** page. When the option to enable Secure Fabric is selected, the **Cluster Networks** page automatically displays additional fields to specify the VLAN IDs required to configure Secure Fabric.

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

Commencing with the October 2024 Oracle Exadata System Software maintenance releases (24.1.5, 23.1.19, and 22.1.28), the option to enable Secure Fabric is selected by default for all new configurations using VM clusters.

See Using the Browser-based Version of Oracle Exadata Deployment Assistant.

Related Topics

- Using Exadata Secure RDMA Fabric Isolation

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

2

Understanding the Network Requirements for Oracle Exadata

Review the network requirements for Oracle Exadata 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 and Oracle Exadata Storage Expansion Rack.

- Overview of Network Requirements
  In addition to the database and storage servers, Oracle Exadata 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
- IPv6 Support
- 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 default administration network IP addresses are assigned dynamically by the elastic configuration procedure during the first start of the system.
- Default Port Assignments

2.1 Overview of Network Requirements

In addition to the database and storage servers, Oracle Exadata 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 environment.

To deploy Oracle Exadata ensure that you meet the minimum network requirements. Oracle Exadata requires a minimum of three networks, and there are interfaces available for additional networks. Each network must be on a separate and distinct subnet. The network descriptions are as follows:

- **Administration Network**: Also known as the management network, this required network connects to your existing management network infrastructure, and is used for administrative work on all components of Oracle Exadata. By default, the administration

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

network connects the database servers, storage servers, server Integrated Lights Out Manager (ILOM) interfaces, and RDMA Network Fabric switches to the Management Network Switch in the rack. One uplink is required from the Management Network Switch to your management network.

Each database server and storage server has two network interfaces for administration. One interface provides management access to the operating system through a dedicated Ethernet port. The other network interface is dedicated to ILOM. By default, Oracle Exadata is delivered with both interfaces connected to the Management Network Switch. Cabling or configuration changes to these interfaces is not permitted, except that starting with Oracle Exadata System Software release 19.1.0, the ILOM interfaces can be connected to a dedicated ILOM network, which is separate from the administration network. The administration network interfaces on the database servers should not be used for client or application network traffic.

Notes

- Separate uplinks to your 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 Management Network Switch.
- A properly secured configuration requires full isolation of the administration network from all other networks.

- **Client 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 databases 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. Non-bonded network configurations are not supported on Oracle Exadata X7 and later systems.
- **Private Network**: Also known as the RDMA Network Fabric, storage network, or interconnect. This network connects the database servers and storage servers. Oracle AI Database uses this network for Oracle RAC cluster interconnect traffic and for accessing data on the Oracle Exadata Storage Servers. The private network is automatically configured during installation. It is non-routable, fully contained in Oracle Exadata, and does not connect to your existing networks.

  Starting with Oracle Exadata X8M, the private network uses RDMA over Converged Ethernet (RoCE).

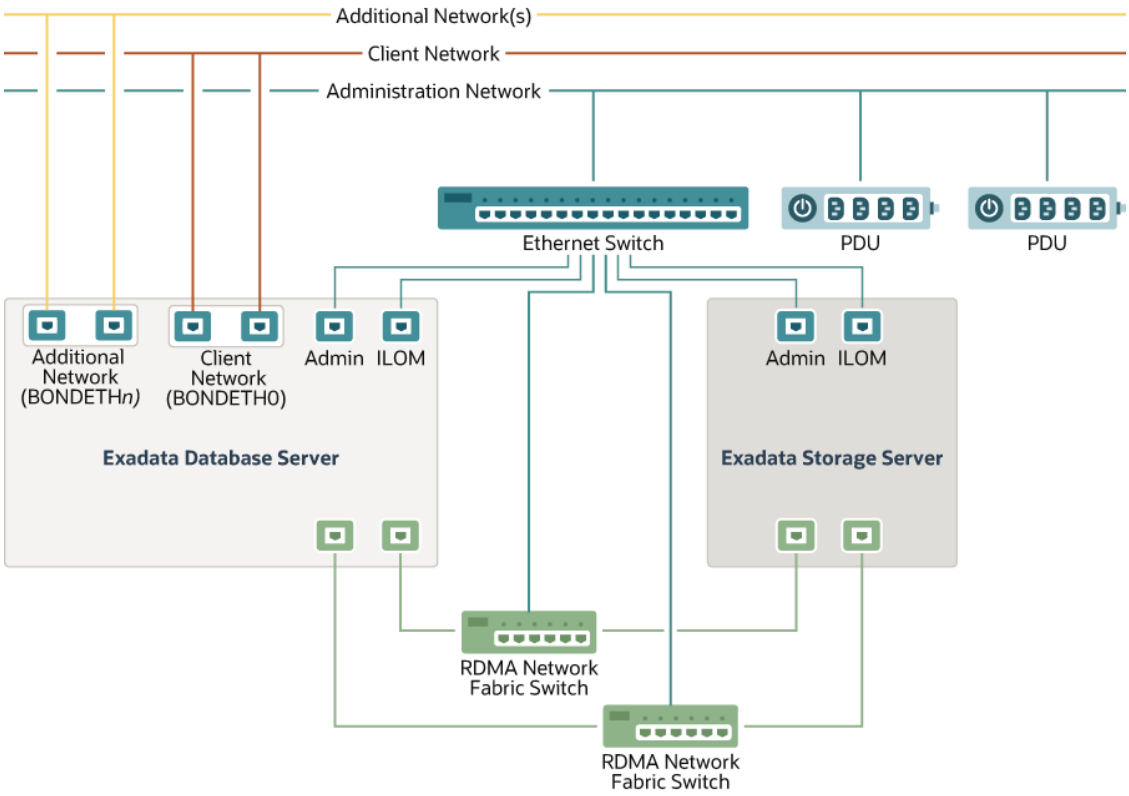
  Previously, the private network was built using InfiniBand technology. RoCE Network Fabric uses different switches and cables from those used by InfiniBand Network Fabric.
- **Additional Networks**: Database servers can optionally connect to additional networks using the available open ports not used by the administration network and the client network.

  By using the OEDA Web user interface, you can create up to two additional networks. In OEDA, the first additional network is known as the **Backup Network**, and the second additional network is known as the **Other Network**. You can create more additional networks by using the OEDA command-line interface (OEDACLI).

  Like the client network, the additional networks support channel bonding to maximize bandwidth and availability. Non-bonded network configurations are not supported on Oracle Exadata X7 and later systems.

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

The following diagram displays how the various Oracle Exadata components connect to the different networks.



Related Topics

- Using Oracle Exadata Deployment Assistant
- Hardware Components of Oracle Exadata Database Machine

2.2 Network Channel Bonding Support

A pair of database server network ports can be bonded to provide higher network availability or bandwidth for the client network and additional networks.

Non-bonded network configurations are not supported on Oracle Exadata X7 and later systems.

In a bonded network configuration:

- Use Oracle Exadata Deployment Assistant (OEDA) to specify the physical network interfaces that you want to bond for the client network and the additional networks, if configured. OEDA generates bonded network interfaces that amalgamate two physical network interfaces.
- Manual changes to the OEDA-generated bonding configuration are allowed but are discouraged. Oracle will not support questions or issues associated with non-standard bonding configurations. In any event, bonding configurations having fewer than two networking interfaces are not permitted.
- The bonded client network interface name is `bondeth0`. The bonded interface name for the first additional network, also known in OEDA as the Backup Network, is `bondeth1`. The

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

bonded interface name for the second additional network, also known in OEDA as the Other Network, is `bondeth2`, and so on.

- The XML configuration file generated by OEDA includes detailed information that maps the bonded network interfaces to the underlying Ethernet ports.
- During the initial configuration using OEDA, the Linux bonding module is configured to use active-backup mode (`mode=active-backup`) by default. Additional configuration of other bonding parameters is allowed but is outside the scope of installation services and must be performed by customer network engineers. Reconfiguration to enable a different bonding policy is permitted but is discouraged.

  For further details, refer to the "Linux Ethernet Bonding Driver HOWTO" at https://www.kernel.org/doc/Documentation/networking/bonding.txt.
- You must provide network infrastructure (switches) capable of supporting the chosen bonding mode. For example, if Link Aggregation Control Protocol (LACP) is enabled (`mode=802.3ad`), then you must supply and configure the network switches accordingly.

  Requirements for specific bonding policies are documented in the "Linux Ethernet Bonding Driver HOWTO" at https://www.kernel.org/doc/Documentation/networking/bonding.txt

2.3 Network Partitioning on Oracle Exadata

Oracle Exadata supports network partitioning using a variety of mechanisms.

- VLAN Support on Customer-Facing Networks
  Oracle Exadata 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 can use Access VLAN settings to implement server-level isolation across the RoCE Network Fabric.
- Using 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 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.

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

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 KB317502)
- Implementing Tagged VLAN Interfaces in Xen based Oracle VM Environments on Exadata (My Oracle Support Doc ID KB311742)
- Implementing Tagged VLAN Interfaces in KVM based Oracle VM Environments on Exadata (My Oracle Support Doc ID KB641497)

2.3.2 Access VLAN Support with RoCE Network Fabric

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

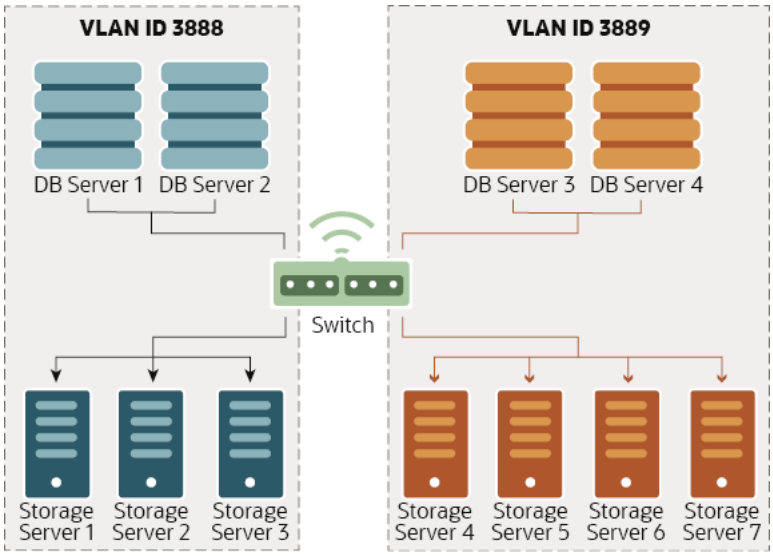
By default, Oracle Exadata uses Access VLAN ID 3888 for all RoCE Network Fabric private network traffic, on the server `re0` and `re1` interfaces. This setting enables all database servers and storage servers to communicate freely with each other, and is suitable for many system configurations. However, 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 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

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

- 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. Furthermore, from October 2024, Secure Fabric is recommended by default on all new configurations.

Exadata Secure RDMA Fabric Isolation enables strict network isolation for virtual machine (VM) clusters on Oracle Exadata systems that use RDMA over Converged Ethernet (RoCE).

Secure Fabric provides critical infrastructure for secure consolidation of multiple tenants on Oracle Exadata, where each tenant resides in a dedicated 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.

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

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.

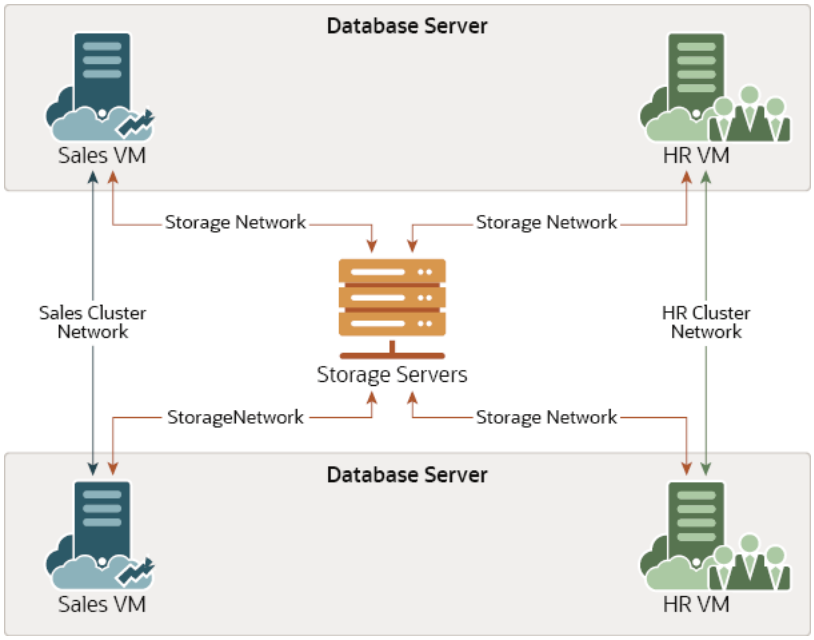
With Secure Fabric, each database cluster uses a dedicated network partition and VLAN ID for cluster networking between the database servers, which supports Oracle Real Application Clusters (Oracle RAC) inter-node messaging. In this partition, all of the database servers are full members. They can communicate 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.

The following diagram illustrates the network partitions that support Exadata Secure RDMA Fabric Isolation. In the diagram, the line connecting the Sales VMs illustrates the Sales cluster network. The Sales cluster network is the dedicated network partition that supports cluster communication between the Sales VMs. The line connecting the HR VMs illustrates the HR cluster network. The HR cluster network is another dedicated network partition that supports cluster communication between the HR VMs. The lines connecting the database server VMs (Sales and HR) to the storage servers illustrate the storage network. The storage network is the shared network partition that supports communications between the database server VMs and the storage servers. But, it does not allow communication between the Sales and HR clusters.

Figure 2-1 Secure Fabric Network Partitions



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

As illustrated in the diagram, each database server (KVM host) can support multiple VMs in separate database clusters. However, Secure Fabric does not support configurations where one database server contains multiple VMs belonging to the same database cluster. In other words, using the preceding example, one database server cannot support multiple Sales VMs or multiple HR VMs.

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.

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 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 OEDA.

   See Configuring the RoCE Network Fabric Switches to Enable Exadata Secure RDMA Fabric Isolation.

2. As part of initial system deployment using OEDA, select the option to enable Secure Fabric and specify VLAN IDs for the cluster and storage network partitions associated with each VM cluster.

   In the OEDA Web user interface, the option to enable Secure Fabric is one of the advanced options associated with the **Cluster Networks** page. When the option to enable Secure Fabric is selected, the **Cluster Networks** page automatically displays additional fields to specify the VLAN IDs required to configure Secure Fabric.

   Commencing with the October 2024 Oracle Exadata System Software maintenance releases (24.1.5, 23.1.19, and 22.1.28), the option to enable Secure Fabric is selected by default for all new configurations using VM clusters.

   See Using the Browser-based Version of 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.

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

With Oracle Exadata, 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 IPv6 Support

Oracle Exadata database and storage servers are enabled to use IPv6 for the admin network, ILOM, and the client access network. This works for both bare metal and virtualized deployments. Furthermore, IPv6 support is available on various other Exadata components.

The following table describes how various Exadata components support IPv6:

Component	Description of IPv6 Support
Oracle Exadata Deployment Assistant (OEDA)	OEDA allows users to specify IPv6 addresses for various network interfaces. When IPv6 addresses are used for the administration network, then the DNS servers, NTP servers, SMTP servers, and SNMP servers also need to be on an IPv6 network.
Cisco switch	The Cisco 4948E-F switch supports IPv6 management addresses with firmware version 15.2(3)E2 or later. Refer to My Oracle Support document KB384251 for upgrade instructions.
Auto Service Request (ASR)	Oracle ASR Manager supports IPv6 starting with Oracle ASR Manager version 5.4. This support requires dual stack IPv6/IPv4 mode on the Oracle ASR Manager server.
Enterprise Manager	Oracle Enterprise Manager supports IPv6 with Exadata starting in Enterprise Manager 13.4. Enterprise Manager needs to be on a bridged network such that it can monitor both the RDMA Network Fabric (on an IPv4 network) and the Exadata database and storage servers (on an IPv6 network).
patchmgr	If the target hosts use only IPv6 addresses, patchmgr requires remote repositories that are hosted on a machine with an IPv6 address (or a name that resolves to an IPv6 IP address) unless the customer network has a bridge or gateway to route between IPv4 and IPv6.
RDMA Network Fabric	The RDMA Network Fabric is a private network that uses IPv4 only. Note that only private addresses are used on the RDMA Network Fabric, so there is negligible benefit from using IPv6.

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

Component	Description of IPv6 Support
SMTP and SNMP	If the target host uses IPv6 addresses, SMTP and SNMP servers should usually be IPv6 (or a name that resolves to an IPv6 IP address) unless the customer network has a bridge or gateway to route between IPv4 and IPv6.
Platinum Support	Platinum Support is available for IPv6 deployments.

Related Topics

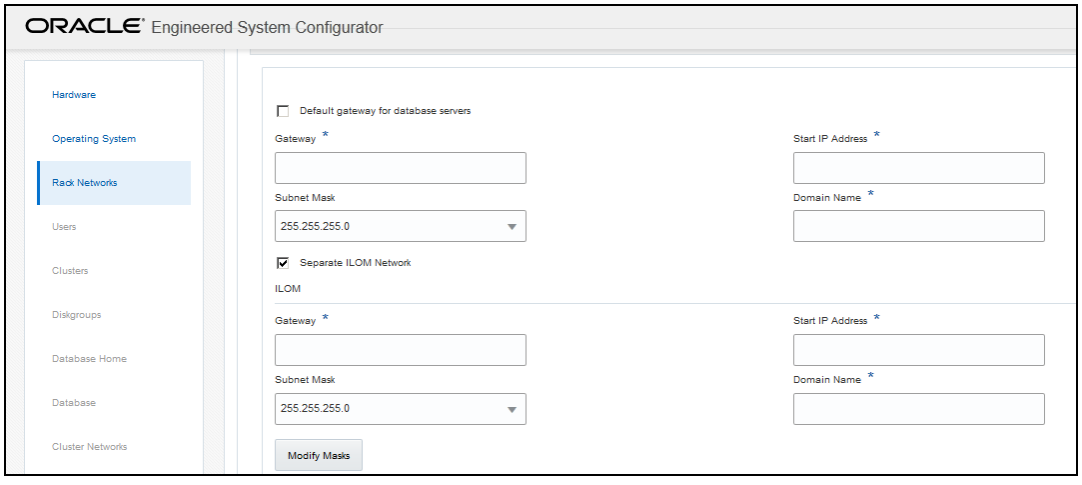
- IPv6 support status on Exadata Database Machine (My Oracle Support Doc ID KB608858)
- Upgrading firmware / Configuring SSH on Cisco Catalyst 4948 Ethernet Switch (My Oracle Support Doc ID KB384251)

2.5 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.



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

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

- 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.

2.6 Default IP Addresses

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

The default administration network IP addresses are in the 172.16.2.1 to 172.16.7.254 range. In earlier releases, Oracle Exadata 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 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 have been created, but before the Oracle Exadata 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 run OEDA on a Microsoft Windows system, then the generated script is `checkip.bat`.

2.7 Default Port Assignments

The following table lists the ports used by services on Oracle Exadata. Review the list and open the required ports to enable network communication through a firewall.

> **Note**
>
> - The table shows default port assignments. Port assignments may vary from system to system because of implementation-specific customizations.
> - Private network ports and interfaces are not listed. Private network communication should not be controlled by using a firewall.
> - Unless otherwise specified, references to Exadata database servers apply to bare-metal servers, virtual machine (VM) hosts, and VM guests. References to database server ILOM interfaces apply only to physical ILOM server interfaces.

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

Table 2-1 Default Port Assignments

Source	Target	Protocol	Port	Network	Application
Any	Database servers, Exadata Storage Servers, network switches, and Integrated Lights Out Manager (ILOM) interfaces	SSH over TCP	22	Administration	SSH
Database servers, Exadata Storage Servers, and network switches	SMTP e-mail server	SMTP	25	Administration	SMTP (Simple Mail Transfer Protocol) for Exadata alerts
Database servers	DNS servers	UDP or TCP	53	Client	DNS (Domain Name System)
Database servers, Exadata Storage Servers, network switches, and ILOM interfaces	DNS servers	UDP or TCP	53	Administration	DNS
Any	ILOM interfaces on database servers, Exadata Storage Servers, and ILOM-enabled network switches	HTTP	80	Administration	ILOM Web interface (user configurable, Default: redirection to port 443)
Any	rpcbind	TCP	111	Administration	rpcbind
Database servers	NTP servers	NTP over UDP	123	Client	Outgoing Network Time Protocol (NTP)
Database servers, Exadata Storage Servers, network switches, and ILOM interfaces	NTP servers	NTP over UDP	123	Administration	Outgoing NTP

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

Table 2-1 (Cont.) Default Port Assignments

Source	Target	Protocol	Port	Network	Application
Any	ILOM interfaces on database servers, Exadata Storage Servers, and ILOM-enabled network switches	SNMP over UDP	161	Administration	SNMP (Simple Network Management Protocol) (user configurable)
Any	PDU	SNMP over UDP	161	Administration	SNMP (user configurable)
Exadata Storage Servers	SNMP subscriber such as Oracle Enterprise Manager Cloud Control or an SNMP manager	SNMP	162	Administration	SNMP version 1 (SNMPv1) outgoing traps (user-configurable)
Database servers, Exadata Storage Servers, network switches, and ILOM interfaces	ASR Manager	SNMP	162	Administration	Telemetry messages sent to ASR Manager
ILOM interfaces on database servers, Exadata Storage Servers, and ILOM-enabled network switches	Any	IPMI over UDP	162	Administration	Outgoing Intelligent Platform Management Interface (IPMI) Platform Event Trap (PET)
Exadata Storage Server ILOMs	Management Server (MS)	SNMPv3	162	Administration	Exadata Storage Server ILOM SNMP notification rules
PDU	SNMP trap receivers	SNMP over UDP	162	Administration	Outgoing SNMPv2 traps
Any	Management Server (MS) on Exadata Storage Servers	HTTPS	443	Administration	Requests from ExaCLI and/or RESTful API calls

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

Table 2-1 (Cont.) Default Port Assignments

Source	Target	Protocol	Port	Network	Application
Any	ILOM interfaces on database servers, Exadata Storage Servers, and ILOM-enabled network switches	HTTPS	443	Administration	ILOM Web interface (user configurable)
Any	PDU	HTTPS	443	Administration	PDU Web interface
Exadata Storage Servers	SMTPS client	SMTPS	465	Administration	Simple Mail Transfer Protocol, Secure (if configured)
Database servers, Exadata Storage Servers, network switches, and ILOM interfaces	Syslog server	Syslog over UDP	514	Administration	Outgoing Syslog
PDU	Syslog server	Syslog over UDP	514	Administration	Outgoing Syslog
Any	ILOM interfaces on database servers, Exadata Storage Servers, and ILOM-enabled network switches	IPMI over UDP	623	Administration	IPMI
Any	plathwsvcd	TCP	723	Administration	plathwsvcd
Any	evnd	TCP	791	Administration	evnd
Any	partitiond	TCP	867	Administration	partitiond
Any	Database servers	TCP	1521	Administration	Oracle Net listener
Any	Database servers	TCP	1521	Client	Oracle Net listener
Any	tgtd	TCP	3260	Administration	SCSI target daemon
Any	Database servers	TCP	3872	Administration	Java EM agent
Any	Database servers	TCP	5000	Administration	Oracle Trace File Analyzer Collector

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

Table 2-1 (Cont.) Default Port Assignments

Source	Target	Protocol	Port	Network	Application
Any	Exadata Storage Servers	HTTPS	5052	Administration	Exascale control services, also known as Exascale RESTful Services (ERS)
Any	Exadata Storage Servers	TCP	5053	Administration	Fast node death detection (FNDD) on RDMA over Converged Ethernet (RoCE) systems
Any	ILOM interfaces on database servers and Exadata Storage Servers	TCP	5120	Administration	ILOM remote console: CD
Any	ILOM interfaces on database servers and Exadata Storage Servers	TCP	5121	Administration	ILOM remote console: keyboard and mouse
Any	ILOM interfaces on database servers and Exadata Storage Servers	TCP	5123	Administration	ILOM remote console: diskette
Any	ILOM interfaces on database servers and Exadata Storage Servers	TCP	5555	Administration	ILOM remote console: encryption
Any	ILOM interfaces on database servers and Exadata Storage Servers	TCP	5556	Administration	ILOM remote console: authentication
Any	Database servers	TCP	6200	Administration	Oracle Notification Service (ONS)
Any	Database servers	TCP	6200	Client	Oracle Notification Service (ONS)

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

Table 2-1 (Cont.) Default Port Assignments

Source	Target	Protocol	Port	Network	Application
ASR Manager	ILOM interfaces on database servers and Exadata Storage Servers	HTTP	6481	Administration	Service tag listener for asset activation
Any	ILOM interfaces on database servers and Exadata Storage Servers	TCP	6481	Administration	ILOM remote console: Servicetag daemon
Any	ILOM interfaces on database servers and Exadata Storage Servers	TCP	7578	Administration	ILOM remote console: video
Any	ILOM interfaces on database servers and Exadata Storage Servers	TCP	7579	Administration	ILOM remote console: serial
Any	Database servers and Exadata Storage Servers	TCP	7777	Administration	Oracle Enterprise Manager Grid Control HTTP console port
Any	Database servers and Exadata Storage Servers	TCP	7799	Administration	Oracle Enterprise Manager Grid Control HTTPS console port
Any	Database servers and Exadata Storage Servers	TCP	7777	Client	Oracle Enterprise Manager Grid Control HTTP console port
Any	Database servers and Exadata Storage Servers	TCP	7799	Client	Oracle Enterprise Manager Grid Control HTTPS console port

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

Table 2-1 (Cont.) Default Port Assignments

Source	Target	Protocol	Port	Network	Application
Any	Management Server (MS) on database servers and Exadata Storage Servers	TCP	7878 8888	Administration	MS access through Oracle WebLogic **Note:** Applies only to Oracle Exadata System Software before release 20.1.0.
Any	Management Server (MS) on database servers	HTTPS	7879	Administration	Requests from ExaCLI and/or RESTful API calls
Database servers and Exadata Storage Servers	ASR Manager	HTTPS	8100 16161	Administration	Diagpack uploads
Database server ILOM interfaces	Management Server (MS)	SNMPv3	8162	Administration	Database Server ILOM SNMP notification rules
Any	rpc.statd	TCP	21408 40801 41460 47431	Administration	rpc.statd

See Also

Managing Oracle Database Port Numbers in *Oracle Real Application Clusters Installation Guide for Linux and UNIX*.

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

3
Using Oracle Exadata Deployment Assistant

This chapter explains how to use the Oracle Exadata Deployment Assistant, a tool for configuring Oracle Exadata, Oracle Zero Data Loss Recovery Appliance, and other Oracle Engineered Systems. The information you enter in the configuration interface is used to generate various configuration files. These files then serve as input for the installation programs, allowing you to customize the automated deployment process.

- Overview of Oracle Exadata Deployment Assistant
  Use Oracle Exadata Deployment Assistant (OEDA) to specify the system configuration details and drive the system configuration processes.
- Getting Started with the OEDA Browser-Based User Interface
  The OEDA application server is available for Linux, OSX and Windows.
- Using the Browser-based Version of Oracle Exadata Deployment Assistant

3.1 Overview of Oracle Exadata Deployment Assistant

Use Oracle Exadata Deployment Assistant (OEDA) to specify the system configuration details and drive the system configuration processes.

OEDA provides a graphical user interface to gather your configuration details and create the Oracle Exadata Rack configuration file. The configuration file drives the automated installation and configuration processes for Oracle Exadata Rack.

> **Note**
>
> For ease of reading, Oracle Exadata Rack is used when information refers to both Oracle Exadata and Oracle Exadata Storage Expansion 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 (My Oracle Support Doc ID KB153930). OEDA is also available on Oracle Technology Network.

In addition to Oracle Exadata Rack, OEDA is also used for Oracle Zero Data Loss Recovery Appliance and Oracle SuperCluster.

Oracle Exadata System Software release 19.1.0 introduced the Web-based interface for OEDA, which replaces the previous Java-based user interface as the graphical user interface for configuring Oracle Exadata Rack.

The following outlines how OEDA is used during the implementation of Oracle Exadata Rack:

- 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. OEDA supports IPv6 addresses.

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

2. Define the settings for the rack, such as network configuration and backup method.
3. Download and install (unzip) the latest version of OEDA from Oracle Technology Network.
4. If you are using a version of OEDA released in conjunction with Oracle Exadata System Software release 26.1.0 or later, ensure that any system running OEDA programs contains Java Development Kit (JDK) version 17 or later. You can download a supported JDK from https://www.oracle.com/java/technologies/downloads/.

   Before running any OEDA program, check the following to ensure that the required JDK is installed and available:

   - Ensure that the `JAVA_HOME` environment variable references the intended JDK.
   - Ensure that `JAVA_HOME/bin/java` is executable.
   - Run `java --version` and confirm that the output reports Java version 17 or higher.
5. Run the configuration interface on a supported platform.

   See Getting Started with the OEDA Browser-Based User Interface.
6. 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 various naming details, networking details (including DNS and NTP servers), and other configuration details.
7. 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.
8. 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 Oracle Exadata Deployment Assistant.
- Considerations and Requirement for Using OEDA
  Review this information before using Oracle Exadata Deployment Assistant (OEDA).

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

3.1.1 Considerations and Requirement for Using OEDA

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

- Oracle Exadata ships with the Oracle Linux operating system installed on the servers.
- The rack prefix is used to generate host names for all components. For example, if you specify `dbm0` for the rack prefix:
  - The database server host names are like `dbm0db01`
  - The storage server host names are like `dbm0cel01`
  - The InfiniBand Network Fabric switch names are like `dbm0sw-iba1`
  - The RoCE Network Fabric switch names are like `dbm0sw-rocea1`

  > **Note**
  >
  > The rack prefix can contain a maximum of 20 letters and numbers. Spaces and symbols are not permitted; for example, exclamation (!), hyphen (-), and so on.

  In you have multiple racks, then, each rack prefix should be a unique value that identifies the rack. Oracle recommends using `dbm01` for the first rack, `dbm02` for the second, `dbm03` for the third, and so on.
- If you configure Oracle Exadata with storage management using Oracle Automatic Storage Management (Oracle ASM), the backup method and redundancy protection levels you use determine how to size the 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. 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.
- If you configure Oracle Exadata with storage management using Oracle Automatic Storage Management (Oracle ASM), 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:

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

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
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. Time zone values provided for Oracle Exadata 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 Getting Started with the OEDA Browser-Based User Interface

The OEDA application server is available for Linux, OSX and Windows.

The Oracle Exadata Deployment Assistant (OEDA) Web browser-based user interface is available starting with the October 2018 OEDA release. The OEDA browser-based interface can import previous OEDA XML configuration files created with the Java-based version of OEDA.

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

> **Note**
>
> The OEDA browser-based interface is supported only on Chrome and Firefox browsers.

To set up the OEDA application server and start the browser-based user interface, perform the following steps.

1. Download the latest OEDA release from Exadata Database Machine and Exadata Storage Server Supported Versions (My Oracle Support Doc ID KB153930) and copy the archive to the system where you want to run the OEDA application server.
2. On the system where you want to run the OEDA application server, extract the contents of the downloaded compressed file. The extracted contents is written to a directory based on the operating system, such as `linux-x64`, `macosx-x64`, or `windows-i586`. This is referred to as the OEDA_HOME directory.
3. If you are using a version of OEDA released in conjunction with Oracle Exadata System Software release 26.1.0 or later, ensure that any system running OEDA programs contains Java Development Kit (JDK) version 17 or later. You can download a supported JDK from https://www.oracle.com/java/technologies/downloads/.

   Before running any OEDA program, check the following to ensure that the required JDK is installed and available:

   - Ensure that the `JAVA_HOME` environment variable references the intended JDK.
   - Ensure that `JAVA_HOME/bin/java` is executable.
   - Run `java --version` and confirm that the output reports Java version 17 or higher.
4. Launch the OEDA application server (located in the OEDA_HOME directory).

   - On Linux, OSX, and UNIX run:

     ```
     installOedaServer.sh [-p port-number] [-g] [-https | -https-keystore
     keystore-path [-https-password]]
     ```

   - On Microsoft Windows run:

     ```
     installOedaServer.cmd [-p port-number] [-g] [-https | -https-keystore
     keystore-path [-https-password]]
     ```

   In the command:

   - `-p` *`port-number`* — optionally specifies the port number for the OEDA application server.

     If not specified, the default port number is `7072`. However, you can specify any unused port number. It is not recommended to use port numbers less than `1024`.
   - `-g` — optionally specifies that the OEDA application server accepts connections on all available network interfaces.

     If not specified, the OEDA application server only listens on the localhost interface (`127.0.0.1`).

   Additionally, in versions of OEDA released in conjunction with Oracle Exadata System Software release 26.1.0 or later, the following arguments enable optional support for HTTPS in the OEDA application server.

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

- `-https` — enables HTTPS in the OEDA application server using the self-signed certificate bundled with OEDA.

  The bundled self-signed certificate only supports HTTPS on the localhost interface (`127.0.0.1`). Therefore, this option cannot be used in conjunction with the `-g` option.

- `-https-keystore` *`keystore-path`* — enables HTTPS in the OEDA application server using a user-provided certificate, which is contained in a PKCS12 keystore located at *keystore-path*.

  A user-provided certificate is required to enable HTTPS support using a specific hostname.

- `-https-password` — securely prompts the user to provide the keystore password, if required.

You can run the `installOedaServer` program using a non-privileged OS user account. You do not need to be logged in as an OS administrator.

When you run the `installOedaServer` program, previous OEDA application server instances are stopped and removed before starting the latest OEDA application server version.

5. Access the OEDA browser-based user interface.

   After you start the OEDA application server, you can access the browser-based interface by opening a browser and entering the following URL:

   `http[s]://`*`host-name`*`:`*`port-number`*`/oeda`

   In the URL:

   - *host-name* is the application server host name.
   - *port-number* is the network port number that you specified when starting the `installOedaServer` program. For example, `7072`.

   If you did not specify the `-g` option when starting the OEDA application server, then you must run the browser on the same host as the OEDA application server and specify `localhost` as the application server host name.

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

Use the Oracle Exadata Deployment Assistant (OEDA) for initial configuration, and when modifying or adding to an existing deployment. You can import an existing configuration when adding new components or modifying an existing deployment.

Have the following available before running OEDA:

- IP addresses
- Machine names
- DNS and NTP information

The following options are available in the **Options** drop-down menu:

- **About**: View version information for OEDA.
- **New**: Start a new configuration.

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

- **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 diagnose issues with the OEDA Web interface.
- **Preferences**: Specify your preferences for how the Web interface behaves.
- **Validate (OEDA Step 1)**: Check the current configuration for completeness and correctness.
- **CheckIp**: Checks the IP addresses.
- **Installation Template**: Generates the Installation Template report, which summarizes the current configuration.
- **OEDA Keyboard Shortcut Keys**: Describes keyboard shortcuts for navigating the OEDA Web interface.

The following list describes the configuration flow in the OEDA Web interface.

1. Specify the hardware for the deployment on the **Select Hardware** page. This page also captures customer information and various rack-level configuration options.
2. Use the **Choose Operating System** page to choose whether to configure virtualization on the database servers.

   > **Note**
   >
   > In a virtualized configuration, each Exadata database server may only contain a limited number of VMs up to the supported maximum.
   >
   > For Exadata X8M-2 and later systems, see Oracle Linux KVM Deployment Specifications and Limits. Otherwise, see Oracle VM Deployment Specifications and Limits.

3. Enter the administration network and private network information on the **Rack Networks** page.
4. Specify the configuration of database server operating system users and groups on the **Users** page.

   You can choose between a Default, Role Separated, or Custom configuration.
5. Configure the Exascale cluster and storage pool on the **Exascale** page.
6. Define the Oracle Grid Infrastructure (GI) clusters on the **Clusters** page.

   Each cluster definition specifies the database and storage servers that are associated with the cluster.

   In a virtualized configuration, each cluster definition also specifies the CPU and memory resources allocated to each VM. You can also customize the configuration of the guest user filesystem (`/u01`) and optionally add others (`/u02`, and so on).
7. Define the Oracle ASM disk group layout on the **Diskgroups** page.
8. Specify the Oracle Database home information for each Oracle home on the **Database Home** page.
9. Specify databases and their attributes on the **Database** page.

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

> **Note**
>
> To create databases on Exadata, Oracle recommends using Oracle Exadata Deployment Assistant (OEDA) or the OEDA Command Line Utility (OEDACLI), as both utilities implement the latest Exadata best-practice settings for your system configuration.
>
> The Oracle Database Configuration Assistant (DBCA) is also available in the Oracle Database software bundle. However, DBCA is not recommended for new database creation on Exadata because it does not integrate with the OEDA XML configuration file and may not implement the recommended Exadata-specific configuration settings.
>
> Although OEDA and OEDACLI support a range of configuration options for many common use cases, they are not intended to configure every advanced Oracle Database feature, specialized security configuration, or optional database component. If your environment requires additional Oracle Database features, advanced security configurations, or optional database components, you can configure and manage those features outside the scope of OEDA and OEDACLI by using standard Oracle Database tools and procedures. In that case, you are responsible for configuring, maintaining, and validating any customizations that fall outside the database configurations supported by OEDA and OEDACLI.

10. Enter the network information on the **Cluster Networks** page.

> **Note**
>
> Commencing with the October 2024 Oracle Exadata System Software maintenance releases (24.1.5, 23.1.19, and 22.1.28), the option to enable Secure Fabric is selected by default for all new configurations using VM clusters. To configure Secure Fabric, you must also specify VLAN IDs for the private network partitions on the **Cluster Networks** page.

11. Configure the email alerts, SNMP alerts, Auto Service Request, Oracle Config Manager, and Enterprise Manager alerting on the **Alerting** page.
12. 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.
13. 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.
14. 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.

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

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.
- Running 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

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

> **Note**
>
> To create databases on Exadata, Oracle recommends using Oracle Exadata Deployment Assistant (OEDA) or the OEDA Command Line Utility (OEDACLI), as both utilities implement the latest Exadata best-practice settings for your system configuration.
>
> The Oracle Database Configuration Assistant (DBCA) is also available in the Oracle Database software bundle. However, DBCA is not recommended for new database creation on Exadata because it does not integrate with the OEDA XML configuration file (`es.xml`) and may not implement the recommended Exadata-specific configuration settings.
>
> Although OEDA and OEDACLI support a range of configuration options for many common use cases, they are not intended to configure every advanced Oracle Database feature, specialized security configuration, or optional database component. If your environment requires additional Oracle Database features, advanced security configurations, or optional database components, you can configure and manage those features outside the scope of OEDA and OEDACLI by using standard Oracle Database tools and procedures. In that case, you are responsible for configuring, maintaining, and validating any customizations that fall outside the database configurations supported by OEDA and OEDACLI.
>
> See ADD DATABASE.

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

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

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
   ```


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

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 Running the OEDACLI Utility

You can perform various Exadata deployment and management operations using the Oracle Exadata Deployment Assistant (OEDA) Command Line Utility (OEDACLI).

Prerequisites

The following recommendations and requirements apply when using OEDACLI:

- Always download and use the latest OEDA software bundle.
- If you are using a version of OEDA released in conjunction with Oracle Exadata System Software release 26.1.0 or later, ensure that any system running OEDA programs contains Java Development Kit (JDK) version 17 or later. You can download a supported JDK from https://www.oracle.com/java/technologies/downloads/.

  Before running any OEDA program, check the following to ensure that the required JDK is installed and available:

  – Ensure that the `JAVA_HOME` environment variable references the intended JDK.
  – Ensure that `JAVA_HOME/bin/java` is executable.
  – Run `java --version` and confirm that the output reports Java version 17 or higher.
- Many OEDACLI commands and operations require a current Exadata XML configuration file. If you cannot supply an up-to-date file, you can generate one using the OEDACLI `DISCOVER ES` command.

  To generate an up-to-date Exadata XML configuration file, run the `DISCOVER ES` command and specify:

  – The complete list of servers (KVM hosts and storage servers) in your Exadata system.
  – The directory location to store the Exadata XML configuration information.

  For example:

  ```
  oedacli> DISCOVER ES
  HOSTNAMES='exa01adm01,exa01adm02,exa01celadm01,exa01celadm02,exa01celadm03'
   LOCATION=/u01/log/exa01_discovered
  Discovering nodes [exa01adm01, exa01adm02, exa01celadm01,exa01celadm02,
  exa01celadm03]...
  Getting Guest info on exa01adm02.example.com
  Getting Guest info on exa01adm01.example.com
  ...
  ```

- OEDACLI performs actions on Exadata system servers using various OS user accounts. To do this, OEDACLI can use default user passwords or you must provide the user

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

passwords to OEDACLI by using the `genPasswordHash.sh` command located in the OEDA software bundle. See also Using the OEDA genPasswordHash Utility.

Syntax

```
./oedacli [ -h ] [ -l ] [ -j ] [ -q ] [ -f commandfile ] [ -c configfile [ -e
immediatecommand ]]
```

Options

- `-h`, `--help` - Display help for command-line options.
- `-l`, `--enhanced-logging` - Enable verbose logging.
- `-j`, `--json-output` - Produce `LIST` command output in JSON format.
- `-q`, `--quiet-mode` - For `LIST` commands, return only data (and no on-screen status).
- `-f`, `--command-file` *`commandfile`* - Specifies a file containing commands that you want to run. You can use this option to specify a non-interactive scripted series of commands that you want to run.
- `-c`, `--config-file` *`configfile`* - Specifies the name of the Exadata XML configuration file used by OEDACLI. If not specified, you can load the file in the command line interface using the `LOAD FILE` command.
- `-e`, `--immediate-command` *`immediatecommand`* - Specifies an immediate command to run, typically a single `LIST` command. If specified, this must be the last command option.

Usage Notes

When run without any other options, OEDACLI starts an interactive command session.

- 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.

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
```


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

```
        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>
```


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

```
 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.

- ACFSVOLUME
  You can add, alter, delete, or list Oracle ACFS volumes.
- ACTION
  You can save, merge, and deploy OEDACLI actions.
- ALERT
  You can create, alter, delete, or list Oracle alerts.
- 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.
- CURRENTACTION
  You can clear or list the current action in OEDACLI.
- DATABASE
  You can add, alter, delete, or list the Oracle databases.
- DATABASEHOME
  You can add, alter, delete, or list the Oracle Database homes.
- DISKGROUP
  You can add, alter, delete, or list the Oracle ASM disk groups.
- DOM0
  You can list Oracle VM management domains (dom0s).
- DOMAIN
  You can get live information from a domain.

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

- ES
  You can discover an engineered system or alter a limited number of attributes of an engineered system.
- EXASCALE
- EXASCALECLUSTER
- EXITONERROR
  You can set the exit status for OEDACLI.
- FILE
  You can manage the Engineered System XML file (`es.xml`).
- GUEST
  You can clone or delete an Oracle VM guest.
- ILOM
  You can alter or list the ILOMS in the loaded `es.xml` file.
- LOGIN
  You can check remote login to a machine or cluster.
- 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.
- PASSWORD
  You can specify the default passwords.
- PROPERTY
  You can set OEDA properties.
- RACK
  You can manage an Engineered System rack.
- REQUIREDFILES
  You can display information about required files.
- REMOTE
  You can run a command on a remote host.
- SCAN
  You can add, alter, delete, or list the SCANs for a cluster.
- SECURITY
  You can manage security settings.
- SOFTWAREVERSIONS
  You can check supported software versions.
- SSHKEYS
  You can generate and deploy SSH keys for `root` user authentication on Oracle Exadata Rack.
- STEPS
  You can display the steps for a virtual `CREATE CLUSTER` deployment.
- STORAGEPOOL
- STORAGETEMPLATES
- SWITCH
  You can display and manage the switches.

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

- SU
  Controls use of the su command.
- ULOC
  You can display the location of rack items.
- VAULT
- VIP
  You can add, alter, delete, or list the Oracle Clusterware VIPs.
- VOLUME
  You can display and manage storage volumes.
- XMLACTION
  You can manage actions relating to the Engineered System XML file.

4.3.1 ACFSVOLUME

You can add, alter, delete, or list Oracle ACFS volumes.

- ADD ACFSVOLUME
- ALTER ACFSVOLUME
- DELETE ACFSVOLUME
- LIST ACFSVOLUMES

4.3.1.1 ADD ACFSVOLUME

This command adds Oracle Advanced Cluster File System (Oracle ACFS) volumes.

Syntax

```
ADD ACFSVOLUME
   VOLUMENAME=volume_name
   VOLUMESIZE=volume_size
   MOUNTPATH=mount_path
WHERE
   { DISKGROUPID=diskgroup_id |
     CLUSTERNUMBER=cluster_number DISKGROUPNAME=diskgroup_name |
     CLUSTERNAME=cluster_name DISKGROUPNAME=diskgroup_name |
     CLUSTERID=cluster_id DISKGROUPNAME=diskgroup_name }
```

Arguments

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

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

- CLUSTERID : Specifies the es.xml ID of the cluster
- DISKGROUPNAME : Specifies the name of the container disk group

4.3.1.2 ALTER ACFSVOLUME

This command alters an existing Oracle Advanced 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

4.3.1.3 DELETE ACFSVOLUME

This command deletes an existing Oracle Advanced Cluster File System (Oracle ACFS) volume.

Syntax

```
DELETE ACFSVOLUME
WHERE
   { ID=acfs_volume_id |
```


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

```
     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

4.3.1.4 LIST ACFSVOLUMES

This command lists the Oracle Advanced Cluster File System (Oracle ACFS) volumes.

Syntax

```
LIST ACFSVOLUMES
[ WHERE
     { ID=acfs_volume_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

4.3.2 ACTION

You can save, merge, and deploy OEDACLI actions.

- SAVE ACTION

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

- MERGE ACTIONS
- DEPLOY ACTIONS

4.3.2.1 SAVE ACTION

Use this command to save the current action to the in-memory copy of the Engineered System XML configuration.

Syntax

```
SAVE ACTION [FORCE]
```

Arguments

`FORCE`: Also merges the current action into the in-memory copy of the Engineered System XML configuration. Using the option provides a shortcut that is equivalent to running `SAVE ACTION` immediately followed by `MERGE ACTIONS`.

> **Note**
>
> The Engineered System configuration should be saved as a file using a subsequent `SAVE FILE` command.

4.3.2.2 MERGE ACTIONS

Use this command to merge the current actions into the in-memory copy of the Engineered System XML configuration.

Syntax

```
MERGE ACTIONS [FORCE]
```

Arguments

`FORCE`: Also marks the actions as deployed in the Engineered System XML configuration. Use this option to alter the Engineered System XML configuration without running `DEPLOY ACTIONS`. This option is typically used with OEDACLI in edit mode to make modifications to the Engineered System XML configuration file before using the OEDA deployment utility (`install.sh`).

> **Note**
>
> The Engineered System configuration should be saved as a file using a subsequent `SAVE FILE` command.

4.3.2.3 DEPLOY ACTIONS

Use this command to run all actions from the Engineered System XML file.

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

Syntax

```
DEPLOY ACTIONS
```

Example 4-3 Deploying the Actions to Add a Database

```
oedacli> DEPLOY ACTIONS
Deploying Action ID : 4 add database dbname='testdb' DATADG='DATAC1'
RECODG='RECOC1' where DBHOMEID='c0_databaseHome'
Deploying ADD DATABASE
Running datapatch on database 'testdb'
Done...
Done
```

4.3.3 ALERT

You can create, alter, delete, or list Oracle alerts.

- ALTER ALERT
- DELETE ALERT
- LIST ALERTS

4.3.3.1 ALTER ALERT

This command creates or alters an Oracle alert.

Syntax

```
ALTER ALERT
   { RECIPIENTS='recipients'
     FROMNAME='from_name'
     FROMEMAIL=from_email
     SMTPHOST=smtp_host
     [ SSL=uses_ssl ]
     [ PORT=smtp_port ]
   | SNMPHOST=snmp_host
     [ PORT=snmp_port ]
     [ COMMUNITY=snmp_community ]
   | ASRHOST=asr_host
     CONTACT=asr_contact
     EMAIL=asr_email
     MOSACCOUNT=asr_mos_account
     [ PROXYHOST=asr_proxy_host ]
     [ PROXYPORT=asr_proxy_port ]
     [ PROXYUSER=asr_proxy_user ]
   | [ EMAIL=ocm_email ]
     [ HUBHOST=ocm_hub_host ]
     [ HUBUSER=ocm_hub_user ]
     [ PROXYHOST=ocm_proxy_host ]
     [ PROXYPORT=ocm_proxy_port ]
     [ PROXYUSER=ocm_proxy_user ]
   | EMBASE=em_base
```


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

```
     [ OMSHOST=oms_host ]
     PORT=oms_port }
WHERE
   { ID=alert_id |
     TYPE=alert_type |
     HOSTNAME=host_name TYPE=alert_type |
     CLUSTERNUMBER=cluster_number TYPE=alert_type |
     CLUSTERNAME=cluster_name TYPE=alert_type |
     CLUSTERID=cluster_id TYPE=alert_type }
```

Arguments

You can specify various attributes according to the alert type.

- For SMTP alerts:
  - `RECIPIENTS` : Specifies a list of email addresses. The list entries are separated by a space or a comma and the list is surrounded by single quotation marks
  - `FROMNAME` : Specifies the name of the email sending user, surrounded by single quotation marks
  - `FROMEMAIL` : Specifies the email address of the email sending user
  - `SMTPHOST` : Specifies the hostname of the SMTP email server
  - `SSL` : Optionally, set to `true` to use SSL. Default value is `false`.
  - `PORT` : Optionally, specifies the SMTP email server port. Default value is `25`.
- For SNMP alerts:
  - `SNMPHOST` : Specifies the hostname of the SNMP server
  - `PORT` : Optionally, specifies the SNMP server port. Default value is `162`.
  - `COMMUNITY` : Optionally, specifies the SNMP community value. Default value is `public`.
- For Oracle Auto Service Request (ASR) alerts:
  - `ASRHOST` : Specifies the hostname of the ASR server
  - `CONTACT` : Specifies the ASR contact name
  - `EMAIL` : Specifies the ASR email address
  - `MOSACCOUNT` : Specifies the ASR My Oracle Support (MOS) email address
  - `PROXYHOST` : Optionally, specifies the proxy hostname for connectivity to the ASR server
  - `PROXYPORT` : Optionally, specifies the proxy port for connectivity to the ASR server
  - `PROXYUSER` : Optionally, specifies the proxy username for connectivity to the ASR server
- For Oracle Configuration Manager (OCM) alerts:
  - `EMAIL` : Specifies the OCM email address
  - `HUBHOST` : Specifies the hostname of the OCM hub
  - `HUBUSER` : Specifies the OCM hub user name
  - `PROXYHOST` : Optionally, specifies the proxy hostname for connectivity to the OCM hub
  - `PROXYPORT` : Optionally, specifies the proxy port for connectivity to the OCM hub

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

  - PROXYUSER : Optionally, specifies the proxy username for connectivity to the OCM hub
- For Oracle Enterprise Manager (EM) alerts:
  - EMBASE : Optionally, specifies the EM base directory. Default value is /u01/app/oracle/product/EMbase.
  - OMSHOST : Specifies the hostname of the Oracle Management Server (OMS)
  - PORT : Optionally, specifies the OMS server port. Default value is 4900.

The following arguments are available in the WHERE clause:

- ID : Specifies the es.xml ID for an existing alert
- TYPE : Specifies the alert type: SNMP, SMTP, ASR, OCM, or EM
- HOSTNAME : Specifies the target server hostname
- 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

4.3.3.2 DELETE ALERT

This command deletes a specific alert.

Syntax

```
DELETE ALERT
WHERE
   { ID=alert_id |
     HOSTNAME=host_name [ TYPE=alert_type ] |
     CLUSTERNUMBER=cluster_number [ TYPE=alert_type ] |
     CLUSTERNAME=cluster_name [ TYPE=alert_type ] |
     CLUSTERID=cluster_id [ TYPE=alert_type ] }
```

Arguments

The following arguments are available in the WHERE clause:

- ID : Specifies the es.xml ID for an existing alert
- HOSTNAME : Specifies the target server hostname
- TYPE : Specifies the alert type: SNMP, SMTP, ASR, OCM, or EM
- 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

4.3.3.3 LIST ALERTS

This command lists alerts.

Syntax

```
LIST ALERTS
[ WHERE
```


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

```
{ ID=alert_id |
  TYPE=alert_type |
  HOSTNAME=host_name [ TYPE=alert_type ] } ]
```

Arguments

The following arguments are available in the `WHERE` clause:

- `ID` : Specifies the `es.xml` ID for an existing alert
- `TYPE` : Specifies the alert type: `SNMP`, `SMTP`, `ASR`, `OCM`, or `EM`
- `HOSTNAME` : Specifies the target server hostname

4.3.4 CELL

You can clone, delete, or list Oracle Exadata storage cells.

- CLONE CELL
- CLONE NEWCELL
- DELETE CELL
- DELETE NEWCELL
- LIST CELLS

4.3.4.1 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 ]
   [ WAIT=wait ]
   [ POWER=power ]
[ 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.

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

- TGTNAME: Specifies the name of the new cell being added.
- WAIT: Default is true, in which case the command waits for completion of the ASM rebalance operation prior to completion. Optionally, set to false to return control without waiting for the ASM rebalance operation.
- POWER : Optionally, specifies the ASM rebalance power limit. The range of valid values is 1 - 1024.
- 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
```


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

```
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

# Rebalance ASM Diskgroups
CLONE CELL SRCNAME=dbm01celadm09 TGTNAME=dbm01celadm14 where
STEPNAME=REBALANCE

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS
```

4.3.4.2 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:
  - `X8XT`: Specifies that the target is an XT storage server with InfiniBand Network Fabric.
  - `X8MXT`: 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.

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

- `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
```

4.3.4.3 DELETE CELL

This command deletes a cell from the OEDA XML configuration file (`es.xml`).

Syntax

```
DELETE CELL
WHERE
   SRCNAME=cell_hostname
```

Arguments

- `SRCNAME` : Specifies the hostname of the cell being deleted.

Usage Notes

- The command performs no deployment actions, and no deployed cells are affected.
- Use `SAVE ACTION FORCE` to save changes to the XML configuration file.
- To remove a deployed cell, use the `ALTER CLUSTER` command instead of `DELETE CELL`.

4.3.4.4 DELETE NEWCELL

This command deletes a cell from an OEDA configuration file (`es.xml`).

Syntax

```
DELETE NEWCELL
WHERE
```


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

```
{ SRCNAME=cell_host_name |
  SRCNAMES=cell_host_names }
```

Arguments

`WHERE` clause arguments:

- `SRCNAME` : Specifies the host name of the cell being deleted from `es.xml`.
- `SRCNAMES` : Specifies a comma-separated list of host names for the cells being deleted from `es.xml`.

Usage Notes

Use `SAVE ACTION FORCE` to save changes only to `es.xml`. No deployment actions are performed.

4.3.4.5 LIST CELLS

This command lists the storage cells.

Syntax

```
LIST CELLS
```

Usage Notes

This command is deprecated. Instead, use:

```
LIST MACHINES WHERE TYPE=CELL
```

4.3.5 CLUSTER

You can alter or list the Oracle RAC Clusters on the Oracle Exadata Database Machine.

- ALTER CLUSTER
- CREATE CLUSTER
  Use this command to create virtual guests for a virtual machine (VM) cluster, including the steps to configure and deploy an Oracle RAC database.
- DESTROY CLUSTER
  This command is used to cleanly remove a running virtual machine (VM) cluster in one operation.
- DOWNGRADE CLUSTER
  Use this command downgrade a previously upgraded cluster.
- LIST CLUSTER
  This command lists a specific Exadata cluster.
- LIST CLUSTERS
  This command lists the existing Exadata clusters.
- UPDATE CLUSTER
  Use this command to update an Oracle Grid Infrastructure cluster, but only within the same major release.

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

- UPGRADE CLUSTER
  Use this command to upgrade an Oracle Grid Infrastructure cluster, either within or across major release boundaries.

4.3.5.1 ALTER CLUSTER

This command alters Oracle Grid Infrastructure (GI) cluster attributes.

Syntax

```
ALTER CLUSTER
   { CLUSTERNAME=cluster_name |
     GIVERSION=grid_version |
     GIHOMELOC=grid_home_location |
     INVLOC=inventory_location |
     BASEDIR=oracle_base_dir |
     PATCHLIST=patch_list |
     VCPU=vcpu |
     VMEM=vmem |
     ASMSCOPEDSECURITY=asm_scoped_security |
     HOMESIZE=home_size |
     INSTALLTYPE={ CLONE | ZIP }
   } ...
WHERE
   { CLUSTERNUMBER=cluster_number |
     CLUSTERNAME=cluster_name |
     ID=cluster_id }


ALTER CLUSTER
   { ADDCELLS='add_cells' | DROPCELLS='drop_cells' }
      [ DISKGROUPS='asm_diskgroups' ] [ POWER=power ] [ WAIT={ TRUE |
FALSE }]
WHERE
   { CLUSTERNUMBER=cluster_number |
     CLUSTERNAME=cluster_name |
     ID=cluster_id }
   [ STEPNAME=step_name ]
```

Arguments

You can modify the following attributes of the GI cluster:

- `CLUSTERNAME`: Specifies a new, unique name for the cluster
- `GIVERSION`: Specifies the clusterware version, including the release update (RU). For example: `12.1.0.2.170418`
- `GIHOMELOC`: Specifies a new path for the Grid home location
- `INVLOC`: Specifies a new path for the software inventory
- `BASEDIR`: Specifies a new path for the Oracle base directory
- `PATCHLIST`: Specifies a comma-separated list of patches. The list must be enclosed in single quotes. By default, the patch list is derived from the RU included in the `GIVERSION`.
- `VCPU`: Specifies the number of virtual CPUs for each virtual machine (VM) in a VM cluster. This option only applies to a cluster in a VM deployment.

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

- VMEM: Specifies the amount of memory (in GB) for each virtual machine (VM) in a VM cluster. This option only applies to a cluster in a VM deployment.
- ASMSCOPEDSECURITY: Specify true to enable ASM scoped security. The default value is false.
- HOMESIZE: Specifies the cluster home size (in GB).
- 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.
- ADDCELLS: Specifies a comma-separated list of cells to add to the disk groups. The list must be enclosed in single quotes. If multiple cells are specified, they are all added in one operation, and the data is rebalanced once at the end of the process.
- DROPCELLS: Specifies a comma-separated list of cells to drop from the disk groups. The list must be enclosed in single quotes. If multiple cells are specified, they are all dropped in one operation, and the data is rebalanced once at the end of the process.
- DISKGROUPS: Specifies a comma-separated list of disk groups from which you want to add or remove cells specified in ADDCELLS or DROPCELLS. The list must be enclosed in single quotes. If not specified, then the specified cells are added to or removed from all of the disk groups.
- POWER: Specifies the ASM rebalance power limit. The range of valid values is 1 - 1024.
- WAIT: Specifies whether or not to wait for the completion of the ASM rebalance operation that occurs during the ADD_DISKS_TO_ASM step of an ADDCELLS or DROPCELLS operation. This option is only applicable when the WHERE clause contains STEPNAME=ADD_DISKS_TO_ASM.

  Specify false to proceed without waiting for the ASM rebalance operation to complete. Specify true to wait for the ASM rebalance operation to complete. The default value is true.

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
- ID: Specifies the es.xml ID of the cluster
- STEPNAME: Specifies one of the following steps for an ADDCELLS or DROPCELLS operation:
  - CONFIG_CELL
  - CREATE_GRIDDISKS
  - ADD_DISKS_TO_ASM

  For an ADDCELLS operation, the steps are performed in the following order:

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

1. CONFIG_CELL: Configures the storage server, including cell connectivity and cell disk creation.
2. CREATE_GRIDDISKS: Creates grid disks for use by the specified cluster.
3. ADD_DISKS_TO_ASM: Adds the newly created grid disks to the ASM disk groups associated with the specified cluster.

For a DROPCELLS operation, the steps are performed in reverse order:

1. ADD_DISKS_TO_ASM: Removes the grid disks from the ASM disk groups associated with the specified cluster.
2. CREATE_GRIDDISKS: Deletes grid disks associated with the specified cluster.
3. CONFIG_CELL: Deletes the cell configuration on the storage server.

Example 4-4 Deploying a Cluster Change

The following example shows an OEDACLI session that uses the ALTER CLUSTER command along with typical output at each step.

```
oedacli> alter cluster clustername=cluster-c1 where id=Cluster-c6e838fe4-e317-
ee1a-dd89-53f0b49a689b_id
oedacli> save action
oedacli> merge actions
 processMerge
 processMergeActions
 Merging Action : alter cluster clustername=cluster-c1 where id=Cluster-
c6e838fe4-e317-ee1a-dd89-53f0b49a689b_id
 Merging ALTER CLUSTER
 Action Validated and Merged OK
oedacli> deploy actions
 Deploying Action ID : 1 alter cluster clustername=cluster-c1 where
id=Cluster-c6e838fe4-e317-ee1a-dd89-53f0b49a689b_id
 Deploying ALTER CLUSTER
 Done...
 Done
```

4.3.5.2 CREATE CLUSTER

Use this command to create virtual guests for a virtual machine (VM) cluster, including the steps to configure and deploy an Oracle RAC database.

Syntax

```
CREATE CLUSTER
   [ STEPS='step_short_names' [ DELETEDB={ true | false } ] |
     SKIPSTEPS='step_short_names' [ DELETEDB={ true | false } ] |
     DELETEDB={ true | false } ]
[ WHERE
     { CLUSTERNAME=cluster_name |
       CLUSTERNUMBER=cluster_number |
       ID=cluster_id } ]
```

Arguments

- STEPS : Optionally, specifies the steps to run.

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

The list consists of space-separated step short names and the entire list is enclosed in single quotes.

Use the `LIST STEPS` command to view the available steps.

- `SKIPSTEPS` : Optionally, specifies steps to skip.

  The list consists of space-separated step short names and the entire list is enclosed in single quotes.

  Use the `LIST STEPS` command to view the available steps.

- `DELETEDB` : Optionally, specify `true` to delete the default Oracle RAC database after cluster creation. The default value is `false`.

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 command can only create a cluster in a VM deployment. You cannot use this command to create the cluster in a physical (bare-metal) deployment.
- The simplest form of this command does not require any arguments. However, if the XML configuration file (`es.xml`) contains more than one cluster, then the `WHERE` clause is mandatory.

Example 4-5 Using the STEPS Parameter with CREATE CLUSTER

Use the following command to only create the guest domains for the `Cluster-c1` cluster:

```
CREATE CLUSTER STEPS='CREATEVM' WHERE CLUSTERNAME=Cluster-c1
```

Example 4-6 Using the SKIPSTEPS and DELETEDB Parameters with CREATE CLUSTER

Use the following command to create the `Cluster-c1` cluster, skipping the specified `SKIPSTEPS` and deleting the default database.

```
CREATE CLUSTER SKIPSTEPS='VALIDATE CALIBRATE RESECURE' DELETEDB=true WHERE
CLUSTERNAME=Cluster-c1
```

4.3.5.3 DESTROY CLUSTER

This command is used to cleanly remove a running virtual machine (VM) cluster in one operation.

Syntax

```
DESTROY CLUSTER
[ WHERE
     { CLUSTERNUMBER=cluster_number |
       CLUSTERNAME=cluster_name |
       CLUSTERID=cluster_id } ]
```


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

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

- No `WHERE` clause is required if the XML file contains one cluster only.
- The `MERGE` step indicates the machines that will be destroyed and the affected cells.

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
```

4.3.5.4 DOWNGRADE CLUSTER

Use this command downgrade a previously upgraded cluster.

Syntax

```
DOWNGRADE CLUSTER
WHERE
   { CLUSTERNAME=cluster_name [ STEPNAME=step_name ] |
     CLUSTERNUMBER=cluster_number [ STEPNAME=step_name ] |
     ID=cluster_id [ STEPNAME=step_name ] }
```

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
- `ID`: Specifies the `es.xml` ID of the cluster
- `STEPNAME` : Optionally, specifies the individual downgrade step:
  1. `RUN_ROOTSCRIPT`
  2. `CONFIG_HOME`
  3. `ADD_HOME`

Usage Notes

- To use the command, you must use the `es.xml` that contains the upgraded cluster.

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

- The command can only downgrade a cluster in a VM deployment. You cannot use this command to downgrade the cluster in a physical (bare-metal) deployment.
- The command can only run if the `DELETE_OLDHOME` step has not been run on the upgraded cluster.

Example 4-8 Cluster Downgrade

The following example shows an OEDACLI session that uses the `DOWNGRADE CLUSTER` command along with typical output at each step.

```
oedacli> downgrade cluster where clusternumber=1
oedacli> save action
oedacli> merge actions
 processMerge
 processMergeActions
 Merging Action : downgrade cluster where clusternumber=1
 Merging DOWNGRADE CLUSTER
 Action Validated and Merged OK
oedacli> deploy actions
 Deploying Action ID : 4 downgrade cluster where clusternumber=1
 Deploying DOWNGRADE CLUSTER
 Downgrading Cluster
 Validating Target cluster version 12.2.0.1
 Validating Target Clusterware Home.. /u01/app/12.2.0.1/grid
 Checking Cluster status...
 Checking status of cluster...
 Checking clustereware file system /u01/app/12.2.0.1/grid
 Validating active cluster version.
 Active cluster version found on cluster ovmClus8 is 18.0.0.0
 Validating Inventory...
 Clusterware Home /u01/app/12.2.0.1/grid validation completed successfully
 Downgrading clusterware to 12.2.0.1
 Running rootcrs.sh on node dbm01adm01vm08.example.com
 Running rootcrs.sh on node dbm01adm02vm08.example.com
 Deregistering Clustereware Home /u01/app/18.13.0.0/grid from oracle inventory
 Registering Clusterware Home /u01/app/12.2.0.1/grid with oracle inventory
 Starting Clusterware stack using software home /u01/app/12.2.0.1/grid
 Removing /u01/app/18.13.0.0/grid from Oracle inventory...
 Updating inventory on dbm01adm01vm08.example.com
 Updating inventory on dbm01adm02vm08.example.com
 Unmounting file systems.....
 Unmounting file system /u01/app/18.13.0.0/grid on dbm01adm01vm08.example.com
 Unmounting file system /u01/app/18.13.0.0/grid on dbm01adm02vm08.example.com
 Updating /etc/fstab entries...
 Active cluster version after downgraded 12.2.0.1.0
 Clusterware successfully downgraded to 12.2.0.1.190115
```

4.3.5.5 LIST CLUSTER

This command lists a specific Exadata cluster.

Syntax

```
LIST CLUSTER { NAME=cluster_name |
               ID=cluster_id }
```


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

Arguments

- `NAME`: Specifies the name of the cluster.
- `ID`: Specifies the cluster ID, as contained in the XML configuration file (`es.xml`).

4.3.5.6 LIST CLUSTERS

This command lists the existing Exadata clusters.

Syntax

```
LIST CLUSTERS
[ WHERE
     { CLUSTERNUMBER=cluster_number |
       CLUSTERNAME=cluster_name |
       ID=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
- `ID`: Specifies the `es.xml` ID of the cluster

Usage Notes

The `WHERE` clause is optional.

Example 4-9 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"
```

4.3.5.7 UPDATE CLUSTER

Use this command to update an Oracle Grid Infrastructure cluster, but only within the same major release.

Syntax

```
UPDATE CLUSTER
   GIVERSION=Grid_relnum
   GIHOMELOC=Grid_home_dir
WHERE
```


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

```
{ CLUSTERNAME=cluster_name [ STEPNAME=step_name ] |
  CLUSTERNUMBER=cluster_number [ STEPNAME=step_name ] |
  ID=cluster_id [ STEPNAME=step_name ] }
```

Arguments

- `GIVERSION`: The version of the Oracle Grid Infrastructure software you are upgrading to, including the release update (RU). For example, `19.26.0.0.250121`. The minimum version supported by the command is `19.15.0.0.220419`.
- `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
- `ID`: Specifies the `es.xml` ID of the cluster
- `STEPNAME` : Optionally, specifies the individual update step:
  1. `ADD_HOME`
  2. `CONFIG_HOME`
  3. `RUN_ROOTSCRIPT`
  4. `DELETE_OLDHOME`

Usage Notes

- The command can only update a cluster to a later version within the same major release. For example, upgrading from Oracle Grid Infrastructure release 19c version 19.15 to version 19.21.
- The command can only update a cluster in a VM deployment. You cannot use this command to update the cluster in a physical (bare-metal) deployment.
- If `STEPNAME` is not specified, then the `ADD_HOME` and `CONFIG_HOME` steps are performed.
- The `RUN_ROOTSCRIPT` step runs the `rootupgrade.sh` script sequentially in a rolling mode on all nodes. This step also restarts the database instances on each of the nodes.
- After you run the `DELETE_OLDHOME` step, you cannot later use the `DOWNGRADE CLUSTER` command to roll back the update.

4.3.5.8 UPGRADE CLUSTER

Use this command to upgrade an Oracle Grid Infrastructure cluster, either within or across major release boundaries.

Syntax

```
UPGRADE CLUSTER
   GIVERSION=Grid_relnum
   GIHOMELOC=Grid_home_dir
WHERE
   { CLUSTERNAME=cluster_name [ STEPNAME=step_name ] |
     CLUSTERNUMBER=cluster_number [ STEPNAME=step_name ] |
     ID=cluster_id [ STEPNAME=step_name ] }
```


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

Arguments

- GIVERSION: The version of the Oracle Grid Infrastructure software you are upgrading to, including the release update (RU). For example, 19.26.0.0.250121.
- 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
- ID: Specifies the es.xml ID of the cluster
- STEPNAME : Optionally, specifies the individual upgrade step:
  1. ADD_HOME
  2. CONFIG_HOME
  3. RUN_ROOTSCRIPT
  4. DELETE_OLDHOME

Usage Notes

- The command can upgrade a cluster to a later version within the same major release or to a newer major release. For example, upgrading from Oracle Grid Infrastructure release 19c to release 26ai.
- The command can only upgrade a cluster in a VM deployment. You cannot use this command to upgrade the cluster in a physical (bare-metal) deployment.
- If STEPNAME is not specified, then the ADD_HOME and CONFIG_HOME steps are performed.
- The RUN_ROOTSCRIPT step runs the rootupgrade.sh script sequentially in a rolling mode on all nodes. This step also restarts the database instances on each of the nodes.
- After you run the DELETE_OLDHOME step in conjunction with the UPGRADE CLUSTER command, you cannot later use the DOWNGRADE CLUSTER command to roll back the upgrade.

Example 4-10 Using the UPGRADE CLUSTER Command

This example shows the first three steps (ADD_HOME, CONFIG_HOME, and RUN_ROOTSCRIPT) in a step-by-step cluster upgrade. At the end of this example, the cluster is upgraded, but downgrade is still possible because the DELETE_OLDHOME step has not been run.

```
oedacli> upgrade cluster giversion=18.13.0.0.210119 gihomeloc=/u01/app/
18.13.0.0/grid where clusternumber=1 stepname=add_home
oedacli> save action
oedacli> merge actions
 processMerge
 processMergeActions
 Merging Action : upgrade cluster giversion=18.13.0.0.210119
gihomeloc=/u01/app/18.13.0.0/grid where clusternumber=1 stepname=add_home
 Merging UPGRADE CLUSTER
 Action Validated and Merged OK
oedacli> deploy actions
 Deploying Action ID : 1 upgrade cluster giversion=18.13.0.0.210119
gihomeloc=/u01/app/18.13.0.0/grid where clusternumber=1 stepname=add_home
 Deploying UPGRADE CLUSTER
 Upgrading Cluster
```


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

```
 Performing Upgrade Validations...
 Validating target Clusterware version 18.13.0.0...
 Validating source Clusterware version 12.2.0.1.190115...
 Checking Cluster status...
 Checking status of cluster...
 Active cluster version on cluster ovmClus8 is 12.2.0.1.0
 Validating Target Clusterware Home.. /u01/app/18.13.0.0/grid
 Checking to see if /u01/app/18.13.0.0/grid is already mounted
 SUCCESS: 18.13.0.0 is supported version for upgrade.
 SUCCESS: Current Cluster Version 12.2.0.1.190115 is supported for Upgrade
using this Utility.
 SUCCESS:
 SUCCESS: Required Image Version 21.2.0.0.0 for Oracle18c found on
dbm01adm01vm08
 SUCCESS: Required Image Version 21.2.0.0.0 for Oracle18c found on
dbm01adm02vm08
 SUCCESS: Cluster verification completed successfully
 SUCCESS: Find the valid upgrading version from 12.2.0.1 to 18.13.0.0
 SUCCESS: Target Cluster Version is valid..
 SUCCESS: Target Clusterware Home /u01/app/18.13.0.0/grid is not already
mounted.
 SUCCESS: Found valid target clusterware version.
 SUCCESS: Valid Bundle Patch / RU Version found.
 Validating required files...
 Extracting files..
 Copying files...
 Creating new disk image file..
 Attaching disk image to Virtual Machine dbm01adm01vm08.example.com
 Attaching disk image to Virtual Machine dbm01adm02vm08.example.com
 Completed adding new Clusterware home /u01/app/18.13.0.0/grid on  Cluster
ovmClus8
 Done...
 Done
oedacli> upgrade cluster giversion=18.13.0.0.210119 gihomeloc=/u01/app/
18.13.0.0/grid where clusternumber=1 stepname=config_home
oedacli> save action
oedacli> merge actions
 processMerge
 processMergeActions
 Merging Action : upgrade cluster giversion=18.13.0.0.210119
gihomeloc=/u01/app/18.13.0.0/grid where clusternumber=1 stepname=config_home
 Merging UPGRADE CLUSTER
 Action Validated and Merged OK
oedacli> deploy actions
 Deploying Action ID : 2 upgrade cluster giversion=18.13.0.0.210119
gihomeloc=/u01/app/18.13.0.0/grid where clusternumber=1 stepname=config_home
 Deploying UPGRADE CLUSTER
 Upgrading Cluster
 Configuring new clusterware home at /u01/app/18.13.0.0/grid
 Running Cluster Verification Utility for upgrade readiness..
 Pre-upgrade Validation completed successfully.
 Setting up new clusterware home...
 Relinking binaries with RDS /u01/app/18.13.0.0/grid
 Upgrading software..
 Completed configuring new Clusterware home /u01/app/18.13.0.0/grid
 Done...
```


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

```
 Done
oedacli> upgrade cluster giversion=18.13.0.0.210119 gihomeloc=/u01/app/
18.13.0.0/grid where clusternumber=1 stepname=RUN_ROOTSCRIPT
oedacli> save action
oedacli> merge actions
 processMerge
 processMergeActions
 Merging Action : upgrade cluster giversion=18.13.0.0.210119
gihomeloc=/u01/app/18.13.0.0/grid where clusternumber=1
stepname=RUN_ROOTSCRIPT
 Merging UPGRADE CLUSTER
 Action Validated and Merged OK
oedacli> deploy actions
 Deploying Action ID : 3 upgrade cluster giversion=18.13.0.0.210119
gihomeloc=/u01/app/18.13.0.0/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 dbm01adm01vm08.example.com
 Checking file
root_dbm01adm01vm08.example.com_2021-05-07_16-35-47-279774287.log on node
dbm01adm01vm08.example.com
 Checking status of cluster on dbm01adm01vm08.example.com
 Running rootupgrade.sh on node dbm01adm02vm08.example.com
 Checking file
root_dbm01adm02vm08.example.com_2021-05-07_16-55-09-777507261.log on node
dbm01adm02vm08.example.com
 Checking status of cluster on dbm01adm02vm08.example.com
 Running Configuration Assistants on dbm01adm01vm08.example.com
 Active cluster version after upgrade 18.0.0.0.0
 Clusterware successfully upgraded to 18.13.0.0.210119
 Applying Security Fixes...
 Done...
 Done
```

4.3.6 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.

4.3.6.1 CLONE COMPUTE

Use this command to clone an existing physical compute cluster node.

Syntax

```
CLONE COMPUTE
   { SRCNAME=source_compute_node TGTNAME=target_compute_node |
```


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

```
     TGTNAMES=target_compute_nodes }
[ 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 to add.
- TGTNAMES: Specifies a comma-separated list of new compute servers to add.
- 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.

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

- `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.
- Use `TGTNAMES` to clone multiple compute servers simultaneously in the `DEPLOY ACTIONS`. Use `TGTNAMES` without additional `SET` statements. Use `SAVE ACTION FORCE` prior to using this argument. See the example below.

Example 4-11 Cloning a Compute Node

This example clones the `dbadm13` compute server to a new server, `dbadm14`.

```
CLONE COMPUTE srcname=dbadm13 tgtname=dbadm14
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 4-12 Cloning Multiple Compute Nodes Using TGTNAMES

This example clones `dbadm23` and creates the new compute servers `dbadm24` and `dbadm25`. By using the `TGTNAMES` argument, both new nodes are deployed at the same time in the final `DEPLOY ACTIONS`.

```
CLONE COMPUTE SRCNAME='dbadm23' TGTNAME='dbadm24'
SET ADMINNET name=dbadm24 ip=10.128.161.192
SET PRIVNET name1=dbadm24-priv1 ip1=192.168.161.195 name2=dbadm24-priv2
ip2=192.168.161.196
...
SAVE ACTION FORCE

CLONE COMPUTE SRCNAME='dbadm23' TGTNAME='dbadm25'
SET ADMINNET name=dbadm25 ip=10.128.161.193
...
SAVE ACTION FORCE

CLONE COMPUTE TGTNAMES='dbadm24,dbadm25'

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS
```


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

Example 4-13 Cloning a Compute Node Using Steps

This example clones the dbadm33 compute server to a new server, dbadm34, running individual steps. The example shows only the first three steps and does not show all the steps involved in cloning the compute node.

```
CLONE COMPUTE srcname=dbadm33 tgtname=dbadm34 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=dbadm33 tgtname=dbadm34 WHERE
STEPNAME=CELL_CONNECTIVITY

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

CLONE COMPUTE srcname=dbadm33 tgtname=dbadm34 WHERE STEPNAME=ADD_NODE

SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

...
```

4.3.6.2 DELETE COMPUTE

This command deletes an existing compute node, removes the instance, home and clusterware from the cluster.

Syntax

```
DELETE COMPUTE
WHERE
   { SRCNAME=host_name [ STEPNAME=step_name ] |
     SRCNAMES=host_names [ STEPNAME=step_name ] }
```

Arguments

SRCNAME specifies the host name of the guest that you want to remove.

SRCNAMES specifies a comma-separated list host names of the guests 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 *step_name* are:

- ADD_INSTANCE

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

- EXTEND_DBHOME
- ADD_NODE
- CELL_CONNECTIVITY
- CREATE_USERS

Example 4-14 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 4-15 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
...
```

4.3.7 CURRENTACTION

You can clear or list the current action in OEDACLI.

- CLEAR CURRENTACTION
- LIST CURRENTACTION

4.3.7.1 CLEAR CURRENTACTION

Clears the current action in OEDACLI.

Syntax

```
CLEAR CURRENTACTION
```


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

4.3.7.2 LIST CURRENTACTION

Lists the current action in OEDACLI.

Syntax

```
LIST CURRENTACTION
```

4.3.8 DATABASE

You can add, alter, delete, or list the Oracle databases.

- ADD DATABASE
- ALTER DATABASE
- DELETE DATABASE
- LIST DATABASES

4.3.8.1 ADD DATABASE

This command adds an Oracle database; either a non-container database (non-CDB), container database (CDB), or pluggable database (PDB) to an existing CDB.

Syntax

```
ADD DATABASE
   DBNAME=database_name
   [ BLOCKSIZE=blocksize ]
   [ CHARSET=database_character_set ]
   [ DBLANG=database_language ]
   [ UNIQUENAME=database_unique_name ]
   [ DBTEMPLATE=database_template ]
   [ DBTYPE=database_type ]
   [ HOSTNAMES='host_names' ]
   [ DATADG=data_diskgroup RECODG=reco_diskgroup ]
   [ VAULTLIST=exascale_vault_list EXASCALECLUSTER=exascale_cluster_name
     [ STORAGETEMPLATE = exascale_storage_template_name ]]
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:

- `DBNAME` : Specifies the name of the new database.
- `BLOCKSIZE` : Optionally, specifies the block size for the new database. The default value is 8192. This argument is not required for pluggable databases.

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

- CHARSET : Optionally, specifies the character set to use for the new database. The default value is AL32UTF8. This argument is not required for pluggable databases.
- DBLANG : Optionally, specifies the language to use for the new database. The default value is all_langs. This argument is not required for pluggable databases.
- UNIQUENAME : Specifies the unique name of the new database. The default value is dbname.
- 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 : Optionally, specifies the type of database to create:
  - normal DB : Adds a non-container database (non-CDB). This is the default value.
  - CDB : Adds a container database.
  - PDB : Adds a pluggable database to an existing CDB.
- HOSTNAMES : Optionally, 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.
- DATADG : Specifies the name of the DATA disk group for the new database.

  This argument is required when adding a database using Oracle ASM storage. It does not apply to pluggable databases (PDBs) or databases using Exascale storage.
- RECODG : Specifies the name of the RECO disk group for the new database.

  This argument is required when adding a database using Oracle ASM storage. It does not apply to pluggable databases (PDBs) or databases using Exascale storage.
- VAULTLIST : Specifies a comma-separated list of Exascale vaults used to store the database.

  This argument is required when adding a database using Exascale storage. It does not apply to pluggable databases (PDBs) or databases using Oracle ASM storage.
- EXASCALECLUSTER : Specifies the name of the Exascale cluster that contains the Exascale vaults used to store the database.

  This argument is required when adding a database using Exascale storage. It does not apply to pluggable databases (PDBs) or databases using Oracle ASM storage.
- STORAGETEMPLATE : Optionally, specifes the Exascale storage template applied to the Oracle Database data files.

  You can use the OEDACLI LIST STORAGETEMPLATES command to view the list of available templates. If not specified, the default template (DATAFILE) applies.

  For example, you can specify DATA_EF_HIGHREDUNDANCY to use the built-in Exascale template that places data files on Extreme Flash (EF) storage media (if available) using high redundancy (triple mirroring).

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
- DBHOMELOCATION : Specifies the path for the target database home.

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

- `CDBNAME` : Specifies the database name for the container database. This argument is required only when `DBTYPE=PDB`

Example 4-16 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='DATAC1' RECODG='RECOC1'
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='DATAC1'
RECODG='RECOC1'
where DBHOMEID='c0_databaseHome'
```


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

```
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='DATAC1'
RECODG='
RECOC1' where DBHOMEID='c0_databaseHome'
Deploying ADD DATABASE
Running datapatch on database 'testdb'
Done...
Done
```

4.3.8.2 ALTER DATABASE

This command alters an existing Oracle database definition.

Syntax

```
ALTER DATABASE
   { { BLOCKSIZE=blocksize |
       CHARSET=database_character_set |
       DATADG=data_diskgroup |
       VAULT='vaultlist' [ EXASCALECLUSTER=exascaleclustername ] |
       DBLANG=database_language |
       DBNAME=database_name |
       UNIQUENAME=database_unique_name |
       DBTEMPLATE=database_template |
       DBTYPE=database_type |
       HOSTNAMES='host_names' |
       RECODG=reco_diskgroup } ... |
     { DBHOMELOC=new_dbhome_loc | DBHOMEID=new_dbhome_id } }
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:

- `BLOCKSIZE`: The block size for the new database. The default value is 8192. This argument is not required for pluggable databases.
- `CHARSET`: The character set to use for the new database. The default value is `AL32UTF8`. This argument is not required for pluggable databases.
- `DATADG`: The name of the DATA disk group for the new database. This argument is not required for pluggable databases.

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

- VAULT: If using Exascale, specify a vault or space-seperated list of two vaults for the database. Note that the value is also surrounded by single quotes (as displayed in the command syntax).
- EXASCALECLUSTER: Specifies the Exascale cluster that contains the specified vault or vaults. This option is required only if the vault specification cannot uniquely identify the Exascale cluster.
- DBLANG: The language to use for the new database. The default value is all_langs. This argument is not required for pluggable databases.
- DBNAME: The name of the new database.
- UNIQUENAME: The unique name of the new database. The default value is dbname.
- DBTEMPLATE : 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: The type of database to create:
  - normal DB: Specifies a non-container database (non-CDB). This is the default value.
  - CDB: Specifies a container database.
  - PDB: Specifies a pluggable database to an existing CDB.
- HOSTNAMES: 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: The name of the RECO disk group for the new database. This argument is not required for pluggable databases.
- DBHOMELOC : The location of the database home that you want to move the database to. This argument is not required for pluggable databases.
- DBHOMEID: The es.xml ID of the database home that you want to move the database to. 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

Usage Notes

For moving the database to use a different database home, specify only the new home location (DBHOMELOC) or the ID of the new database home (DBHOMEID).

4.3.8.3 DELETE DATABASE

This command deletes an existing Oracle database.

Syntax

```
DELETE DATABASE
WHERE
   { ID=database_id |
     CLUSTERNUMBER=cluster_number DATABASENAME=database_name |
```


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

```
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 4-17 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 :
```


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

```
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
```

4.3.8.4 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

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

4.3.9 DATABASEHOME

You can add, alter, delete, or list the Oracle Database homes.

- ADD DATABASEHOME
- ALTER DATABASEHOME
- DELETE DATABASEHOME
- LIST DATABASEHOMES

4.3.9.1 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 release update (RU), 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.

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

- 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 4-18 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 release update (RU). In this example, we are setting up a new Oracle database home with 12.1.0.2.170418. 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
```


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

```
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_1of2.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
node1adm01.example.com
Copying file: linuxamd64_12102_database_2of2.zip to node
node1adm01.example.com
Copying file: p26022196_121020_Linux-x86-64.zip to node
node1adm01.example.com
Copying file: p26027162_121020_Linux-x86-64.zip to node
node1adm01.example.com
Copying file: p26022196_121020_Linux-x86-64.zip to node
node1adm02.example.com
Copying file: p26027162_121020_Linux-x86-64.zip to node
node1adm02.example.com
Completed copying files...
Extracting files...
Completed extracting files...
Validating nodes for database readiness...
Installing database software ...
Running database installer on node node1adm01.example.com ... Please wait...
After running database installer...
Patching...
Done...
Done
```

4.3.9.2 ALTER DATABASEHOME

This command alters an existing Oracle database home.

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

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 release update (RU), 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

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

- 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.

4.3.9.3 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 4-19 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'
```


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

```
Deploying DELETE DATABASEHOME
Validating Oracle home..
Deinstalling database home c0_DbHome_1
Unmounting file system..
Updating /etc/fstab entries...
Done...
Done
```

4.3.9.4 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

4.3.10 DISKGROUP

You can add, alter, delete, or list the Oracle ASM disk groups.

- ADD DISKGROUP
- ALTER DISKGROUP
- ALTER DISKGROUPS
- DELETE DISKGROUP
- LIST DISKGROUPS

4.3.10.1 ADD DISKGROUP

This command adds an Oracle ASM disk group.

Syntax

```
ADD DISKGROUP
   [ ACFSNAME=acfs_name ]
   [ ACFSPATH=acfs_path ]
   [ ACFSSIZE=acfs_size ]
```


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

```
   [ CELLLIST='cell_list' ]
   DISKGROUPNAME=diskgroup_name
   [ DISKGROUPLOCATION=diskgroup_location ]
   DISKGROUPSIZE=diskgroup_size
   [ OCRVOTE=ocr_voting_disks_included ]
   [ QUORUMDISK=quorum_disks_included ]
   REDUNDANCY=redundancy_level
   SLICESIZE=slice_size
   [ SPARSE=sparse ]
   [ 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 disk group
- `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
- `DISKGROUPLOCATION` : Specifies the media type where the disk group is located. The valid values are `FLASHDISK` or `HARDDISK`. The default value is `HARDDISK`.

  The same disk group location is recommended for all non-RECO disk groups across all clusters in the Exadata configuration file (`es.xml`). If the RECO disk group location is `FLASHDISK`, then `FLASHDISK` is the recommended location for all disk groups.
- `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 valid 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.
- `TYPE` : Specifies the disk group 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

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

- `CLUSTERID` : Specifies the `es.xml` ID of the cluster

Example 4-20 Adding a Disk Group

The following example shows an OEDACLI session that uses the `ADD DISKGROUP` command along with typical output at each step.

```
oedacli> add diskgroup diskgroupname=OTHER7 diskgroupsize=1TB
redundancy=HIGH
celllist='dbm01celadm07.example.com,dbm01celadm08.example.com,dbm01celadm09.ex
ample.com'  where clustername=Cluster-c1
oedacli> save action
oedacli> merge actions
 processMerge
 processMergeActions
 Merging Action : add diskgroup diskgroupname=OTHER7 diskgroupsize=1TB
redundancy=HIGH
celllist='dbm01celadm07.example.com,dbm01celadm08.example.com,dbm01celadm09.ex
ample.com'  where clustername=Cluster-c1
 Merging ADD DISKGROUP
 Action Validated and Merged OK
oedacli> deploy actions
 Deploying Action ID : 3 add diskgroup diskgroupname=OTHER7
diskgroupsize=1TB  redundancy=HIGH
celllist='dbm01celadm07.example.com,dbm01celadm08.example.com,dbm01celadm09.ex
ample.com'  where clustername=Cluster-c1
 Deploying ADD DISKGROUP
 Diskgroup OTHER7 will be created on Storage Servers
[dbm01celadm07.example.com, dbm01celadm08.example.com,
dbm01celadm09.example.com]
 Validating free space....
 Creating Grid Disks for ASM Disk Group OTHER7
 Creating ASM Disk Group OTHER7
 Updating ASM Diskstring...
 Checking ASM Disk Group status...
Completed creation of ASM Disk Group OTHER7
Done...
Done
```

4.3.10.2 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 |
     DISKGROUPLOCATION=diskgroup_location |
     DISKGROUPSIZE=diskgroup_size |
     OCRVOTE=ocr_voting_disks_included |
     QUORUMDISK=quorum_disks_included |
     REDUNDANCY=redundancy_level |
```


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

```
     SLICESIZE=slice_size |
     SPARSE=sparse } ...
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 disk group
- `ACFSPATH` : Specifies the new path for the Oracle ACFS file system
- `ACFSSIZE` : Specifies a new size for the Oracle ACFS volume on this disk group
- `CELLLIST` : A comma separated list of cells for this disk group, enclosed in single quotes
- `DISKGROUPNAME` : Specifies a new name for the disk group
- `DISKGROUPLOCATION` : Specifies the media type where the disk group is located. The valid values are `FLASHDISK` or `HARDDISK`. The default value is `HARDDISK`.

  The same disk group location is recommended for all non-RECO disk groups across all clusters in the Exadata configuration file (`es.xml`). If the RECO disk group location is `FLASHDISK`, then `FLASHDISK` is the recommended location for all disk groups.
- `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 valid 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.

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

- With the `ALTER DISKGROUP` command, you can only deploy the following modifications to an existing disk group:
  - You can alter the size of an existing disk group by setting the `DISKGROUPSIZE` attribute.
  - You can change the storage servers that contain the disk group by setting the `CELLLIST` attribute. However, if you remove a storage server from a disk group, you must ensure that the remaining storage servers have enough space to accommodate the existing data.

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

Changes to all other disk group attributes may only be saved to the OEDA XML configuration file (`es.xml`). Consequently, you cannot use the `ALTER DISKGROUP` command in conjunction with the `DEPLOY ACTIONS` command to change the name of an existing disk group, for example.

- If you modify the redundancy, the volume size will be maintained, however the slice size will be recalculated.

Example 4-21 Resizing an Oracle ASM Disk Group 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 : "DATAC1"
redundancy : "NORMAL"
sliceSize : "2953G"
machines :
..
..
cellDisks : "ALL"
diskGroupSize : "53154G"
ocrVote : "false"
quorumDisk : "false"
id : "c0_datadg"
version : "PhysicalLayoutLegacy"
diskGroupName : "RECOC1"
redundancy : "NORMAL"
sliceSize : "ALL"
machines :
..
..
cellDisks : "ALL"
diskGroupSize : "13284G"
ocrVote : "false"
```


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

```
quorumDisk : "false"
id : "c0_recodg"
```

3. Create an action for resizing the disk group.

```
oedacli> ALTER DISKGROUP DISKGROUPSIZE='2500G' WHERE CLUSTERNUMBER=1
DISKGROUPN
AME='DATAC1'
```

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
AC1' 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='DATAC1'
Deploying ALTER DISKGROUP
Deploying Action ID : 4 alter diskgroup diskgroupsize='2500G' where
diskgroupna
me='DATAC1' clusternumber=1
Deploying ALTER DISKGROUP
Validating ASM Disk status..
Disks in ASM Disk Group DATAC1 have valid status
Shrinking Disk Group size to 2500 GB
Resizing ASM Disk Group..
Checking Reblance operations in disk group DATAC1
Rebalance operations completed in disk group DATAC1
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
```

4.3.10.3 ALTER DISKGROUPS

This command alters disk group sizing attributes for existing Oracle ASM disk groups.

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

Syntax

```
ALTER DISKGROUPS
   DATASPLIT=data_split_percent
   RECOSPLIT=reco_split_percent
   DATAREDUNDANCY=data_redundancy_level
   RECOREDUNDANCY=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. The sum of `DATASPLIT` and `RECOSPLIT` cannot exceed 100.
- `RECOSPLIT` : Specifies a percentage of the Oracle ASM volume allocated to the RECO disk group. The sum of `DATASPLIT` and `RECOSPLIT` cannot exceed 100.
- `DATAREDUNDANCY` : Specifies the redundancy level for the DATA disk group. The accepted values are `NORMAL` or `HIGH`.
- `RECOREDUNDANCY` : 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

4.3.10.4 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 }
```


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

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

Example 4-22 Deleting a Disk Group

The following example shows an OEDACLI session that uses the `DELETE DISKGROUP` command along with typical output at each step.

```
oedacli> delete diskgroup where id=10dff870-b824-cdb7-0542-8a03a88e9653
oedacli> save action
oedacli> merge actions
 processMerge
 processMergeActions
 Merging Action : delete diskgroup where id=10dff870-b824-
cdb7-0542-8a03a88e9653
 Merging DELETE DISKGROUP
 Action Validated and Merged OK
oedacli> deploy actions
 Deploying Action ID : 4 delete diskgroup where id=10dff870-b824-
cdb7-0542-8a03a88e9653
 Deploying DELETE DISKGROUP
 Dropping ASM Disk Group DATAC1
 ASM Disk Group DATAC1  does not exist..
 Grid Disks for Disk Group DATAC1 exist on cells dbm01celadm07.example.com
 Grid Disks for Disk Group DATAC1 exist on cells dbm01celadm08.example.com
 Grid Disks for Disk Group DATAC1 exist on cells dbm01celadm09.example.com
 Dropping Grid Disks on [dbm01celadm07.example.com,
dbm01celadm08.example.com, dbm01celadm09.example.com]
 Deleting quorum devices...
 Completed dropping ASM Disk Group DATAC1
 Done...
 Done
```

4.3.10.5 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 } ]
```


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

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

4.3.11 DOM0

You can list Oracle VM management domains (dom0s).

- LIST DOM0S

4.3.11.1 LIST DOM0S

This command lists Oracle VM management domains (dom0s).

Syntax

```
LIST DOM0S
```

Usage Notes

This command is deprecated. Instead, use:

```
LIST MACHINES WHERE TYPE=DOM0
```

4.3.12 DOMAIN

You can get live information from a domain.

- GET DOMAIN

4.3.12.1 GET DOMAIN

This command displays live domain information from a running system.

Syntax

```
GET DOMAIN
   ACTION=action
WHERE
   { ID=id |
     HOSTNAME=hostname |
     CLUSTERNUMBER=cluster_number { COMPUTENUMBER=compute_number |
STORAGENUMBER=storage_number } |
     CLUSTERNAME=cluster_name { COMPUTENUMBER=compute_number |
STORAGENUMBER=storage_number } |
     CLUSTERID=cluster_id { COMPUTENUMBER=compute_number |
STORAGENUMBER=storage_number } }
```


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

Arguments

`ACTION` : Specifies the information that the command returns.

If the action is applied to a guest, then the command displays specific information for the guest.

If the action is applied to the host, then the command displays information for the hypervisor.

- `LISTIMAGES`: Displays image file details.
- `LISTDOMAIN`: Displays domain details.
- `LISTMEMORY`: Displays memory details.
- `LISTVCPU`: Displays CPU details.

The following arguments are available in the `WHERE` clause:

- `ID` : Specifies the `es.xml` ID for the target machine
- `HOSTNAME` : Specifies the current hostname 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 `es.xml` compute node number in the cluster, starting at 1
- `STORAGENUMBER` : Specifies the `es.xml` storage server number in the cluster, starting at 1

4.3.13 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.
- VALIDATE ES

4.3.13.1 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.

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

- `CUSTOMERNAME` : Specifies the customer name.
- `CUSTOMERDEPT` : Specifies the customer department.

4.3.13.2 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.

4.3.13.3 VALIDATE ES

Validates certain aspects of an OEDA XML configuration file (`es.xml`).

Syntax

```
VALIDATE ES
```

4.3.14 EXASCALE

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

- DEPLOY EXASCALE
  This command performs Exascale deployment actions.

4.3.14.1 DEPLOY EXASCALE

This command performs Exascale deployment actions.

Syntax

```
DEPLOY EXASCALE
   ACTION = action
WHERE
   { CLUSTERNUMBER = clusternumber [ DATABASENAME = databasename ] |
     CLUSTERNAME = clustername [ DATABASENAME = databasename ] |
     CLUSTERID = clusterid [ DATABASENAME = databasename ] |
     DATABASENAME = databasename }
```

Arguments

- `ACTION`: Specifies the Exascale deployment action, which is one of the following: `CONFIGURECELLS`, `CONFIGURECOMPUTES`, `CREATEDBWALLET`, `CREATEGIVAULT`, `CREATEDBVAULT`, `UNDOCONFIGURECELLS`, `UNDOCONFIGURECOMPUTES`, `UNDOCREATEDBWALLET`, `UNDOCREATEGIVAULT`, `UNDOCREATEDBVAULT`.

The following arguments are available in the `WHERE` clause:

- `CLUSTERNUMBER` : Specifies the cluster number in the OEDA configuration file (`es.xml`), starting at 1.
- `CLUSTERNAME` : Specifies the name of the cluster.
- `CLUSTERID` : Specifies the OEDA configuration file (`es.xml`) ID of the cluster.
- `DATABASENAME` : Specifies the name of the database.

Usage Notes

- The `CONFIGURECELLS` action must be performed before the `CONFIGURECOMPUTES` action.
- The `CONFIGURECELLS` and `CONFIGURECOMPUTES` actions must be performed before any wallet or vault actions.
- The `CONFIGURECOMPUTES` action creates an Exascale wallet for the Oracle Grid Infrastructure operating system (OS) owner. This wallet also supports Oracle Database in a non-role-separated Exadata configuration, where the Oracle Grid Infrastructure OS owner is the same as the Oracle Database OS owner.
- The `CREATEDBWALLET` action only creates an Exascale wallet for an Oracle Database OS owner in a role-separated Exadata configuration, where the Oracle Database OS owner is not the same as the Oracle Grid Infrastructure OS owner.
- Undo actions must be performed in the following order: `UNDOCREATEDBVAULT`, `UNDOCREATEGIVAULT`, `UNDOCREATEDBWALLET`, `UNDOCONFIGURECOMPUTES` and `UNDOCONFIGURECELLS`.

4.3.15 EXASCALECLUSTER

- ADD EXASCALECLUSTER
  This command adds a new Exascale cluster.

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

- ALTER EXASCALECLUSTER
  This command alters an existing Exascale cluster.
- DELETE EXASCALECLUSTER
  This command deletes an existing Exascale cluster.
- LIST EXASCALECLUSTER
  This command displays details for an existing Exascale cluster.
- LIST EXASCALECLUSTERS
  This command displays details for all Exascale clusters.

4.3.15.1 ADD EXASCALECLUSTER

This command adds a new Exascale cluster.

Syntax

```
ADD EXASCALECLUSTER
   NAME = exascaleclustername
   [ VIP = ersvip
     IP = ersip ]
```

Arguments

- `NAME`: Specifies the name of the Exascale cluster.
- `VIP`: Specifies the DNS name for the virtual IP (VIP) interface that provides access to Exascale control services (also known as Exascale RESTful Services or ERS).
- `IP`: Specifies the IP address for the ERS VIP.

Usage Notes

- The network interface specified by `VIP` and `IP` is associated with a highly-available virtual IP (VIP) network interface, which provides a consistent network end point for Exascale control services. The VIP is hosted by one of the Exadata storage servers, which also runs a front-end ERS instance. If the storage server or ERS instance becomes unavailable, then the VIP moves to another server hosting a front-end ERS instance.

  You should define the ERS VIP if possible. The option to define an Exascale cluster without an ERS VIP primarily exists to support Exascale configurations on existing Exadata systems that have no spare IP addresses to support the ERS VIP.

4.3.15.2 ALTER EXASCALECLUSTER

This command alters an existing Exascale cluster.

Syntax

```
ALTER EXASCALECLUSTER
   [ NAME = newexascaleclustername ]
   [ MODE = deploymentmode ]
   [ VIP = newersvip ]
   [ IP = newersip ]
   [ RESTNETWORKID = newersnetworkid ]
   [ BSWHAVIPIPS = bswhavipips ]
   [ BSWHAVIPNETMASK = bswhavipnetmask ]
```


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

```
   [ ADDCELLS = 'addcells' | DROPCELLS = 'dropcells' ]
WHERE
   NAME = exascaleclustername
   [ STEPNAME = stepname ]
```

Arguments

- `NAME`: Changes the name of the Exascale cluster.
- `MODE`: Specifies the Exascale cluster deployment mode.
- `VIP`: Specifies the DNS name for the virtual IP (VIP) interface that provides access to Exascale control services (also known as Exascale RESTful Services or ERS).
- `IP`: Specifies the IP address for the ERS VIP.
- `RESTNETWORKID`: Specifies the ERS network ID.
- `BSWHAVIPIPS`: Specifies a comma-separated list of IP addresses to use for the Exascale Block Store Worker (BSW) virtual IP (VIP) interfaces.
- `BSWHAVIPNETMASK`: Specifies the netmask to use for the BSW VIPs.
- `ADDCELLS`: Specifies a comma-separated list of cells to add to the Exascale cluster. The list must be enclosed in single quotes.
- `DROPCELLS`: Specifies a comma-separated list of cells to drop from the Exascale cluster. The list must be enclosed in single quotes.

The following arguments are available in the `WHERE` clause:

- `NAME`: Specifies the name for the Exascale cluster that you want to alter.
- `STEPNAME`: Specifies one of the following steps for an `ADDCELLS` or `DROPCELLS` operation:
  - `CONFIG_CELL`
  - `CONFIG_COMPUTES`

  For an `ADDCELLS` operation, the steps are performed in the listed order. For a `DROPCELLS` operation, the steps are performed in the reverse order.

Usage Notes

- Do not specify the Exascale cluster deployment mode using any value other than `onPremises`. All other values are reserved for Oracle internal use.
- The network interface specified by `VIP` and `IP` is associated with a highly-available virtual IP (VIP) network interface, which provides a consistent network end point for Exascale control services. The VIP is hosted by one of the Exadata storage servers, which also runs a front-end ERS instance. If the storage server or ERS instance becomes unavailable, then the VIP moves to another server hosting a front-end ERS instance.

4.3.15.3 DELETE EXASCALECLUSTER

This command deletes an existing Exascale cluster.

Syntax

```
DELETE EXASCALECLUSTER
WHERE
   NAME = exascaleclustername
```


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

Arguments

The following argument is available in the `WHERE` clause:

- `NAME`: Specifies the name for the Exascale cluster that you want to delete.

Usage Notes

- You cannot delete an Exascale cluster while it provides resources to Oracle Grid Infrastructure or an Oracle database.

4.3.15.4 LIST EXASCALECLUSTER

This command displays details for an existing Exascale cluster.

Syntax

```
LIST EXASCALECLUSTER
WHERE
   NAME = exascaleclustername
```

Arguments

The following argument is available in the `WHERE` clause:

- `NAME`: Specifies the name for the Exascale cluster that you want to display.

4.3.15.5 LIST EXASCALECLUSTERS

This command displays details for all Exascale clusters.

Syntax

```
LIST EXASCALECLUSTERS
```

4.3.16 EXITONERROR

You can set the exit status for OEDACLI.

- SET EXITONERROR
  This command controls the exit status for OEDACLI when an error in encountered.

4.3.16.1 SET EXITONERROR

This command controls the exit status for OEDACLI when an error in encountered.

Syntax

```
SET EXITONERROR ENABLE={ true | false }
```

Arguments

- `ENABLE`: When `ENABLE=true`, a run-time error causes OEDACLI to terminate with exit code 1. This option may be useful when OEDACLI is used inside a script.

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

4.3.17 FILE

You can manage the Engineered System XML file (`es.xml`).

- LOAD FILE
- MERGE FILES
- SAVE FILE
- SAVE FILES

4.3.17.1 LOAD FILE

Use this command to load an Engineered System XML file (`es.xml`) into OEDACLI memory for further use.

Syntax

```
LOAD FILE NAME=file_name
```

Arguments

- `NAME` : Specifies the name of an existing Engineered System XML file

Example 4-23 Loading a File into OEDACLI for Editing

```
oedacli> LOAD FILE NAME=Oracle-test12.xml
SUCCESS - file loaded OK
Customer : testadm01 - Oracle
```

4.3.17.2 MERGE FILES

Use this command to merge multiple Engineered System XML files into a new file.

Syntax

```
MERGE FILES
   INPUT='input_files'
   OUTPUT=output_file
```

Arguments

- `INPUT`: Specifies a space-separated list of Engineered System XML files to merge. The list of input files is surrounded by single quotes (`'`).
- `OUTPUT`: Specifies the new target Engineered System XML file. The file is overwritten if it already exists.

4.3.17.3 SAVE FILE

Use this command to save a current set of actions to an Engineered System XML file.

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

Syntax

```
SAVE FILE
   [ NAME=file_name ]
   [ MAC={ true | false } ]
```

Arguments

- `NAME`: Optionally specifies the name of a new Engineered System XML file. If you do not specify the `NAME` argument, then the command overwrites the source file.
- `MAC`: When `MAC=true`, OEDACLI updates the Engineered System XML file with MAC addresses for the system components. The ILOM interfaces must be available for this operation.

4.3.17.4 SAVE FILES

Use this command to generate a full set of OEDA configuration files.

Syntax

```
SAVE FILES
   LOCATION=dir_name
   [ MAC={ true | false } ]
```

Arguments

- `LOCATION`: Specifies the name of a target directory for file creation.
- `MAC`: When `MAC=true`, OEDACLI updates the Engineered System XML file with MAC addresses for the system components. The ILOM interfaces must be available for this operation.

Usage Notes

The specified `LOCATION` directory is created if it does not already exist.

4.3.18 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.
- CREATE GUEST
  This command creates a new virtual machine (VM) guest.
- 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.
- MIGRATE GUEST
  This command moves an Oracle Linux KVM guest to another KVM host.

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

4.3.18.1 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 a 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 VM host where the new guest clone will be created.

Syntax

```
CLONE GUEST
   { SRCNAME=source_guest TGTNAME=target_guest |
     TGTNAMES=target_guests }
[ WHERE
     STEPNAME=stepname ]


SET PARENT NAME=parent_name
[ SET ADMINNET NAME=admin_name IP=admin_ip [ GATEWAY=gateway
NETMASK=netmask ] ]
[ SET BACKUPNET 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 [ NATHOSTNAME=nat_host_name
NATIP=nat_ip NATDOMAINNAME=nat_domain_name, NATNETMASK=nat_netmask ]
SET VIPNET NAME=vip_name,IP=vip_ip
[ SET DNSSERVERS IP1=DNS_server_1 [ IP2=DNS_server_2 [ IP3=DNS_server_3 ]]]
[ SET NTPSERVERS IP1=NTP_server_1 [ IP2=NTP_server_2 [ IP3=NTP_server_3 ]]]
[ SET VCPU COUNT=guest_cpu ]
[ SET VMEM SIZE=guest_memory ]
[ SET VDISK SIZE=guest_disk ]
[ SET IMAGEVERSION VERSION=image_version ]
```

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.
- `TGTNAMES`: Specifies a comma-separated list of new guest VMs to add.
- `STEPNAME`: Specifies the deployment step to perform, which is one of the following:
  - `CREATE_GUEST`
  - `CREATE_USERS`
  - `CELL_CONNECTIVITY`
  - `CONFIG_CLUSTERWARE`
  - `RUN_ROOTSCRIPT`

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

  - 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. Also, optionally specifies the gateway and netmask for the administration network.
- BACKUPNET : 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.
- DNSSERVERS: Specifies the Domain Name System (DNS) servers used by the guest. You can specify up to three DNS server IP addresses. If none are specified, the guest uses the same DNS servers as the host.
- NTPSERVERS: Specifies the Network Time Protocol (NTP) servers used by the guest. You can specify up to three NTP server IP addresses. If none are specified, the guest uses the same NTP servers as the host.
- VCPU : Specifies the number of virtual CPU cores allocated to the new guest VM.
- VMEM : Specifies the amount of RAM (in GB) allocated to the new guest VM.
- VDISK : Specifies the amount of disk space (in GB) allocated to the new guest VM.
- IMAGEVERSION : Specifies the software image version used to deploy the new guest VM.

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.
- You can use the CLONE GUEST command to add a VM to a server that already contains a VM in the same cluster. Consequently, you can create a configuration where one VM host contains multiple VMs that are members of the same cluster. However, this capability is only available on systems not configured with Exadata Secure RDMA Fabric Isolation.

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

- During deployment, the cloned guest inherits various configuration attributes from the source guest, including the client network configuration and the backup network configurations (if present).

  If all of the KVM hosts have the same network configuration, then the inherited attributes work as expected.

  However, if the new KVM host uses a different physical network configuration, deployment of the cloned guest will fail. This situation is most likely when an Exadata system contains different versions of compute node hardware. For example, when adding an X10M server to an X8M-2 rack.

  In this case, you must manually adjust the relevant network definition by using the `ALTER NETWORK` command before deployment. Contact Oracle Support for details.

Example 4-24 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'
TGTNAME='exa01adm03vm01.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
```

Example 4-25 Cloning Multiple Guests Using TGTNAMES

This example clones `dbguest23` and creates the new guests `dbguest24` and `dbguest25`. By using the `TGTNAMES` argument, both new guests are deployed at the same time in the final `DEPLOY ACTIONS`.

```
CLONE GUEST SRCNAME='dbguest23' TGTNAME='dbguest24'
SET ADMINNET name='dbguest24' ip='10.128.161.202'
...
SAVE ACTION FORCE

CLONE GUEST SRCNAME='dbguest23' TGTNAME='dbguest25'
SET ADMINNET name='dbguest25' ip='10.128.161.203'
...
SAVE ACTION FORCE

CLONE GUEST TGTNAMES='dbguest24,dbguest25'
```


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

```
SAVE ACTION
MERGE ACTIONS
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

4.3.18.2 CREATE GUEST

This command creates a new virtual machine (VM) guest.

Syntax

```
CREATE GUEST
   NAME=guest_name


SET PARENT NAME=parent_name
SET ADMINNET NAME=admin_name [ IP=admin_ip GATEWAY=gateway NETMASK=netmask
VLANID=vlan_ID ]
SET CLIENTNET NAME=client_name [ IP=client_ip GATEWAY=gateway NETMASK=netmask
VLANID=vlan_ID ]
[ SET CLIENTNET NATHOSTNAME=nat_host_name NATIP=nat_ip
NATDOMAINNAME=nat_domain_name, NATNETMASK=nat_netmask ]
[ SET BACKUPNET NAME=backup_name [ IP=backup_ip GATEWAY=gateway
NETMASK=netmask VLANID=vlan_ID ] ]
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 DNSSERVERS IP1=DNS_server_1 [ IP2=DNS_server_2 [ IP3=DNS_server_3 ]]]
[ SET NTPSERVERS IP1=NTP_server_1 [ IP2=NTP_server_2 [ IP3=NTP_server_3 ]]]
[ SET VCPU COUNT=guest_cpu ]
[ SET VMEM SIZE=guest_memory ]
[ SET VDISK SIZE=guest_disk ]
```

Arguments

- `NAME`: Specifies the name of the new guest VM.

Additional settings in conjunction with the `CREATE GUEST` command:

- `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 to use for the administration network. Also, optionally specifies the IP address, gateway, netmask, and VLAN ID for the administration network.
- `CLIENTNET` : Specifies the DNS name to use for the client network. Also, optionally specifies the IP address, gateway, netmask, and VLAN ID for the client network. May also optionally specify the host name, IP address, domain name, and netmask details to support Network Address Translation (NAT).
- `BACKUPNET` : Specifies the DNS name to use for the backup network. Also, optionally specifies the IP address, gateway, netmask, and VLAN ID for the backup network.

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

- `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.
- `DNSSERVERS`: Specifies the Domain Name System (DNS) servers used by the guest. You can specify up to three DNS server IP addresses. If none are specified, the guest uses the same DNS servers as the host.
- `NTPSERVERS`: Specifies the Network Time Protocol (NTP) servers used by the guest. You can specify up to three NTP server IP addresses. If none are specified, the guest uses the same NTP servers as the host.
- `VCPU` : Specifies the number of virtual CPU cores allocated to the new guest VM.
- `VMEM` : Specifies the amount of RAM (in GB) allocated to the new guest VM.
- `VDISK` : Specifies the amount of disk space (in GB) allocated to the new guest VM.

Usage Notes

- For configurations with Exadata Secure RDMA Fabric Isolation or InfiniBand security (PKEYS), `INTERCONNECT` defines the private network that connects the processing nodes and `PRIVNET` defines the private storage network.
- 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 4-26 Creating a Guest VM

This example shows how to create a new guest VM. In this example, the new guest is `exa01adm03vm04.example.com`. A series of `SET` commands are used to specify the configuration for the new guest VM.

```
oedacli> LOAD FILE NAME=exa01adm03-pre-createing-vm04.xml
oedacli> CREATE GUEST NAME='exa01adm03vm04.example.com'
oedacli> SET PARENT NAME='exa01adm03.example.com'
oedacli> SET ADMINNET NAME='exa01adm03vm04.example.com' IP='10.xxx.xx.x'
oedacli> SET CLIENTNET NAME='exa01client03vm04.example.com' IP='10.xxx.xx.x'
oedacli> SET PRIVNET NAME1='exa01adm03vm04-priv' IP1='192.168.16.8'
NAME2='exa01adm01vm03-priv2' IP2='192.168.16.9'
oedacli> SET VCPU COUNT=4
oedacli> SET VMEM SIZE=16GB
oedacli> SET VDISK SIZE=50GB
oedacli> SAVE ACTION
oedacli> MERGE ACTIONS
oedacli> SAVE FILE NAME='exa01adm03-after-createing-vm04.xml'
oedacli> DEPLOY ACTIONS
```

Example 4-27 Creating Multiple Guests at the Same Time

This example creates the new guests `guest04` and `guest05` on the host named `host01`. Both new guests are deployed at the same time in the final `DEPLOY ACTIONS`.

```
CREATE GUEST NAME='guest04'
SET PARENT name='host01'
...
SAVE ACTION FORCE
```


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

```
CLONE GUEST NAME='guest05'
SET PARENT name='host01'
...
SAVE ACTION FORCE

MERGE ACTIONS
DEPLOY ACTIONS
```

4.3.18.3 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 ] |
     SRCNAMES='node_names' [ STEPNAME=stepname ] }
```

Arguments

- SRCNAME specifies the host name of an existing DomU that you want to delete.
- SRCNAMES specifies a comma-separated list of host names for existing guests that you want to delete. The specified list of host names must be surrounded by quotation marks.
- STEPNAME allows you to undo individual steps of the clone operation.

  The value of *stepname* can be one of the following:

  - ADD_INSTANCE
  - EXTEND_DBHOME
  - RUN_ROOTSCRIPT
  - CONFIG_CLUSTERWARE
  - CELL_CONNECTIVITY
  - CREATE_USERS
  - CREATE_GUEST

  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
  ```


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

Example 4-28 Deleting a Guest

The following example shows an OEDACLI session that uses the `DELETE GUEST` command along with typical output at each step.

```
oedacli> delete guest where srcname='dbm01vm103.example.com'
oedacli> save action
oedacli> merge actions
 processMerge
 processMergeActions
 Merging Action : delete guest where srcname='dbm01vm103.example.com'
 Merging DELETE GUEST
 Action Validated and Merged OK
oedacli> deploy actions
 Deploying Action ID : 1 delete guest where srcname='dbm01vm103.example.com'
 Deploying DELETE GUEST
 Delete Guest
 Removing Cell Connectivity...
 Deleting cellip.ora and cellinit.ora on node dbm01vm103.example.com
 Done deleting cellip.ora and cellinit.ora on dbm01vm103.example.com
 Successfully completed removing cell connectivity [elapsed Time [Elapsed =
5018 mS [0.0 minutes] Wed May 05 10:29:51 PDT 2021]]
 Deleting users on dbm01vm103.example.com ...
 Deleting cluster users on node dbm01vm103.example.com
 Deleting cluster users...
 Deleting groups...
 Done deleting users and groups on node dbm01vm103.example.com
 Successfully completed deleting users on node(s)
dbm01vm103.example.com[elapsed Time [Elapsed = 6043 mS [0.0 minutes] Wed May
05 10:29:57 PDT 2021]]
 Deleting KVM Guest dbm01vm103.example.com from dom0 dbm01adm03.example.com
 KVM Guest dbm01vm103.example.com deleted successfully.
 Done...
 Done
```

4.3.18.4 MIGRATE GUEST

This command moves an Oracle Linux KVM guest to another KVM host.

Prerequisites

- The Exadata system must use 2-socket Oracle Exadata system hardware with RoCE Network Fabric (X8M-2 or later).
- The Exadata system must use Oracle Exadata System Software release 25.1.0 or later.
- You must have an OEDA-generated Engineered System XML configuration file (`es.xml`) that accurately reflects the current state of the Exadata system.
- The Exadata system must be configured to use Exascale storage to host the guest image files.
- The environment must be configured with SSH equivalency, so that the OS user and the server running OEDACLI can perform actions as the `root` user on both the source and target KVM hosts.
- The source and target KVM hosts must reside in the same Exadata system configuration and have the same network visibility.

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

- The source and target KVM hosts must have access to the same Oracle Exadata storage servers.
- The source and target KVM hosts must not have any outstanding work from a previous Exadata Live Update.
- The target KVM host must use Oracle Exadata System Software packages that are the same version as, or newer than, those on the source KVM host.
- The target KVM host must have sufficient free CPU and memory resources to accommodate the guest.
  - It is possible to over-commit virtual CPUs such that the total number of virtual CPUs assigned to all guests exceeds the number of physical CPUs on the system. Over-committing CPUs can be done only when the competing workloads for over-subscribed resources are well understood and the concurrent demand does not exceed physical capacity.
  - It is not possible to over-commit memory.
- The guest being migrated must not have any outstanding work from a previous Exadata Live Update operation that could be applied when the guest restarts during the migration process. Either clear all outstanding work before migration, or reconfigure the guest so that no outstanding work occurs during the migration process.
- The guest name must not already be in use on the target KVM host.
- Additionally, for Exadata VM Live Migration:
  - The Exadata system hardware must be at least Oracle Exadata X10M.
  - The Exadata system must use Oracle Exadata System Software release 26.1.0 or later.
  - The Exadata system must be configured with Exadata Secure RDMA Fabric Isolation (Secure Fabric) enabled.
  - The source and target KVM hosts must use the same model of server hardware, with the same number and type of CPUs.

Syntax

```
MIGRATE GUEST
     HOSTNAME=guest_name
     MODE={OFFLINE|OFFLINEFORCE|LIVE}
     SRCHOST=source_host
     TGTHOST=target_host
[ WHERE
     STEPNAME=stepname ]
```

Arguments

- `HOSTNAME`: Specifies the hostname of the guest being migrated.
- `MODE`: Specifies the migration mode.
  - `OFFLINE`: Performs an automated migration process where the guest is gracefully shut down, moved to the target KVM host, and then restarted.

    Use this option in most offline migration situations to gracefully migrate a working guest.
  - `OFFLINEFORCE`: Performs a forced offline migration process without gracefully shutting down the guest.

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

Use this option only in situations where the guest or source KVM host are not accessible. For example, to resurrect a guest from a failed KVM host.

  - `LIVE`: Performs the migration using Exadata VM Live Migration. In this case, the guest remains available throughout the operation and normal processing is uninterrupted, apart from a brief switchover to the target KVM host.

- `SRCHOST`: Specifies the hostname or IP address of the source KVM host. This is where the guest resides before the migration operation.
- `TGTHOST`: Specifies the hostname or IP address of the target KVM host. This is where you intent to move the guest to.
- `STEPNAME`: Specifies the migration step to perform.

  The following list shows the migration steps (in order) for an offline migration (`MODE=OFFLINE` or `MODE=OFFLINEFORCE`).

  1. `STOP_GUEST`
  2. `CREATE_BRIDGES`
  3. `DETACH_INTERFACES`
  4. `DETACH_VOLUMES`
  5. `ATTACH_VOLUMES`
  6. `MIGRATE_GUEST`
  7. `ATTACH_INTERFACES`
  8. `STARTUP_GUEST`

  The following list shows the migration steps (in order) for Exadata VM Live Migration (`MODE=LIVE`).

  1. `ATTACH_VOLUMES`
  2. `CREATE_BRIDGES`
  3. `MIGRATE_GUEST`
  4. `DETACH_VOLUMES`

  Additionally, you can perform a precheck for Exadata VM Live Migration using the `PRECHECK_LIVEMIGRATION` step. This step serves as a standalone command that is only performed when the `MIGRATE GUEST` command uses Exadata VM Live Migration (`MODE=LIVE`) and includes `WHERE STEPNAME=PRECHECK_LIVEMIGRATION`.

Example 4-29 Migrating a Guest VM Using Offline Migration

This example shows the command sequence to migrate a guest using offline migration.

In this example, the guest being moved is named `exa01vm01.example.com`, the KVM host where the guest resides before the migration operation is `exa01adm02.example.com`, and the target KVM host is `exa01adm06.example.com`.

```
oedacli> LOAD FILE NAME=exa01.xml
oedacli> MIGRATE GUEST HOSTNAME=exa01vm01.example.com MODE=offline
SRCHOST=exa01adm02.example.com TGTHOST=exa01adm06.example.com
oedacli> SAVE ACTION
oedacli> MERGE ACTIONS
...
Merging MIGRATE GUEST
Action Validated and Merged OK
```


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

```
oedacli> SAVE FILE NAME=exa01.migration-merged.xml
File : exa01.migration-merged.xml saved OK
oedacli> DEPLOY ACTIONS
Deploying Action ID : 28 MIGRATE GUEST HOSTNAME=exa01vm01.example.com
MODE=offline SRCHOST=exa01adm02.example.com TGTHOST=exa01adm06.example.com
Deploying MIGRATE GUEST
...
Done...
Done [Elapsed = 193346 mS [3.0 minutes] Wed Jul 30 16:14:59 PDT 2025]]
oedacli> SAVE FILE NAME=exa01.migration-deployed.xml
File : exa01.migration-deployed.xml saved OK
```

Example 4-30 Performing a Precheck for Exadata VM Live Migration

This example shows how to perform a precheck before migrating a guest using Exadata VM Live Migration.

In this example, the `MIGRATE GUEST` command checks that the guest named `exa01vm01` can be migrated from `exa01adm04` (the source KVM host) to `exa01adm06` (the target KVM host) using Exadata VM Live Migration.

```
oedacli> LOAD FILE NAME=exa01.xml
oedacli> MIGRATE GUEST HOSTNAME=exa01vm01 MODE=live SRCHOST=exa01adm04
TGTHOST=exa01adm06 WHERE STEPNAME=PRECHECK_LIVEMIGRATION
oedacli> SAVE ACTION
oedacli> MERGE ACTIONS
oedacli> DEPLOY ACTIONS
```

Example 4-31 Migrating a Guest VM Using Exadata VM Live Migration

This example shows how to migrate a guest using Exadata VM Live Migration.

In this example, the guest being moved is named `exa01vm01`, the KVM host where the guest resides before the migration operation is `exa01adm04`, and the target KVM host is `exa01adm06`. The example shows only a brief sample of the output from various migration steps.

```
oedacli> LOAD FILE NAME=exa01.xml
oedacli> MIGRATE GUEST HOSTNAME=exa01vm01 MODE=live SRCHOST=exa01adm04
TGTHOST=exa01adm06
oedacli> SAVE ACTION
oedacli> MERGE ACTIONS
oedacli> SAVE FILE NAME=exa01.migration-merged.xml
oedacli> DEPLOY ACTIONS

...
Attaching EDV Volume exa01vm01_sys using device
exa01vm01_sys_c7db815ef4cc457b86a3015d1c467ace to host exa01adm06
...
Mounting Guest Configuration Files volume /dev/exc/
exa01vm01_cfg_1412f96d1bd742248723f203438c5fad for guest exa01vm01 on kvmhost
exa01adm06
...
Migrate guest exa01vm01 from exa01adm04 to exa01adm06. Step MIGRATE_GUEST
completed successfully...
...
Unmounting Guest Configuration Files volume /dev/exc/
```


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

```
exa01vm01_cfg_1412f96d1bd742248723f203438c5fad for guest exa01vm01 on kvmhost
exa01adm04.
...
Detaching EDV volume device exa01vm01_sys_c7db815ef4cc457b86a3015d1c467ace
from source KVM host exa01adm04
...
Done...
Done [Elapsed = 92755 mS [1.5 minutes] Wed Apr 29 15:27:06 PDT 2026]]
oedacli> SAVE FILE NAME=exa01.migration-deployed.xml
File : exa01.migration-deployed.xml saved OK
```

Related Topics

- Moving a Guest Using Exadata VM Live Migration
- Moving a Guest Using Automated Offline Migration

4.3.19 ILOM

You can alter or list the ILOMS in the loaded `es.xml` file.

- ALTER ILOM
- LIST ILOMS

4.3.19.1 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=host_name |
     ILOMNAME=ilom_host_name }
```

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

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

4.3.19.2 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.

4.3.20 LOGIN

You can check remote login to a machine or cluster.

- CHECK LOGIN

4.3.20.1 CHECK LOGIN

This command checks remote login to a machine or cluster.

Syntax

```
CHECK LOGIN
WHERE
   { CLUSTERNUMBER=cluster_number |
     CLUSTERNAME=cluster_name |
     CLUSTERID=cluster_id |
     HOSTNAME=host_name }
```

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
- `HOSTNAME` : Specifies the name of an individual host

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

4.3.21 MACHINE

You can alter or list the machines in the `es.xml` file.

- ALTER MACHINE
- ALTER MACHINES
- LIST MACHINES

4.3.21.1 ALTER MACHINE

This command modifies attributes of an Engineered System machine, which is a physical server or virtual machine (VM) guest.

Syntax

```
ALTER MACHINE
   { { DNSSERVERS='dnsservers' |
       GATEWAYADAPTER=gatewayadapter |
       HOSTNAMEADAPTER=hostnameadapter |
       NTPSERVERS='ntpservers' |
       TIMEZONE=timezone |
       VSWITCH={ true | false } |
       TYPE=newmachinetype |
       ORADISKSIZE=oradisksize |
       ORADISKPATH=oradiskpath }...
     | ACTION={ START [ WAIT={ true | false } ] |
                 STOP [ WAIT={ true | false } ] |
                 RESTART [ WAIT={ true | false } ] |
                 CREATEIMAGE  IMAGEFILE=imagefile SIZE=size
[ FILESYSTEM=filesystem ] [ ZIPFILE=zipfile ] |
                 ATTACHDISK IMAGEFILE=imagefile MOUNTPATH=mountpath
[ FILESYSTEM=filesystem ] |
                 DETACHDISK IMAGEFILE=imagefile [ DELETE=delete ] |
                 SETVCPU [ VCPU=vcpu ] [ MAXVCPU=maxvcpu ] [ REBOOT={ true |
false } ] |
                 SETVMEM VMEM=vmem [ REBOOT={ true | false } ] } }
WHERE
   { ID=id |
     HOSTNAME=hostname |
     CLUSTERNUMBER=clusternumber { COMPUTENUMBER=computenumber |
STORAGENUMBER=storagenumber }  |
     CLUSTERNAME=clustername { COMPUTENUMBER=computenumber |
STORAGENUMBER=storagenumber } |
     CLUSTERID=clusterid { COMPUTENUMBER=computenumber |
STORAGENUMBER=storagenumber } }
```

Arguments

You can specify the following arguments for the machine:

- `DNSSERVERS`: Comma-separated list of DNS servers for the machine. Enclose the list in single quotes (`'`).
- `GATEWAYADAPTER`: Network adapter to use for the machine gateway: `ADMIN` or `CLIENT`

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

- `HOSTNAMEADAPTER`: Network adapter to use for the machine hostname: `ADMIN` or `CLIENT`
- `NTPSERVERS`: Comma-separated list of NTP servers for the machine. Enclose the list in single quotes (`'`).
- `TIMEZONE`: Timezone specification for the machine
- `VSWITCH`: Specifies if the machine is connected to a vswitch. Default is `false`.
- `TYPE`: Sets a machine as new type. For example: `X8XT`
- `ORADISKSIZE`: Size (in GB) of the Oracle software volume (typically `/u01`). Default is 20, minimum is 10.
- `ORADISKPATH`: Mount point of the Oracle software volume. Default is `/u01`. For non-default values, take great care to ensure other paths are compatible with the setting.
- `ACTION`: Action to perform against the specified VM guest machine. One of the following:
  - `START`: Starts the specified guest. The action completes when the machine is ready for SSH connections. There is no error if the machine is already started.
  - `STOP`: Stops the specified guest. The action completes when the machine is shut down. There is no error if the machine is already stopped.
  - `RESTART`: Stops and then starts the specified guest. The action completes when the machine is ready for SSH connections. There is no error if the machine is already stopped.
  - `CREATEIMAGE`: Creates a disk image in the hypervisor for subsequent `ATTACHDISK` commands.
  - `ATTACHDISK`: Attaches a disk image from the hypervisor to a running guest and mounts the file system in the guest. `FILESYSTEM` must be specified if the default (`xfs`) is not used.
  - `DETACHDISK`: Unmount the file system and detach a disk image from a guest, and optionally delete the disk image.
  - `SETVCPU`: Changes the CPU resource allocation for a guest. The action allows allocation increases and reductions within the limit specified by the `maxcpu` value in the `vm.cfg` file. The action can operate on a running guest. A reboot is not required to enable changes.
  - `SETVMEM`: Changes the RAM resource allocation for a guest. A reboot is required to enable a change.

  The following `ACTION` arguments are available:

  - `WAIT`: Optionally specifies whether to wait for the associated action to complete. Default is `true`. With `WAIT=false`, the action starts asynchronously, and the command returns immediately with no indication or guarantee of success.
  - `IMAGEFILE`: Name of the disk image file
  - `SIZE`: Size (in GB) of the disk image file
  - `FILESYSTEM`: File system type used in the disk image. Default is `xfs`. Alternative is `ext4`.
  - `ZIPFILE`: Zip file used to optionally populate a new disk image
  - `MOUNTPATH`: Guest directory used to mount the disk image
  - `DELETE`: Optionally specifies whether to delete the disk image. Default is `false`.
  - `VCPU`: Number of Virtual CPUs allocated to the guest machine

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

  - MAXVCPU: Maximum number of Virtual CPUs allocated to the guest machine
  - VMEM: Amount of RAM (in GB) allocated to the guest machine
  - REBOOT: Optionally specifies whether to automatically reboot the guest. Default is false.

The following arguments are available in the WHERE clause:

- ID: Specifies the identifier for the target machine in the OEDA XML file (es.xml)
- HOSTNAME: Specifies the current host name of the target machine
- CLUSTERNUMBER: Cluster number associated with target machine in the OEDA XML file (es.xml), starting at 1
- CLUSTERNAME: Name of the cluster associated with target machine in the OEDA XML file (es.xml)
- CLUSTERID: ID of the cluster associated with target machine in the OEDA XML file (es.xml)
- COMPUTENUMBER: The compute node number associated with target machine in the OEDA XML file (es.xml). The numbering of compute nodes starts at 1 for each cluster.
- STORAGENUMBER: The storage server number associated with target machine in the OEDA XML file (es.xml). The numbering of storage servers starts at 1 for each cluster.

Example 4-32 ACTION Examples

The following examples show commands that use the ACTION argument:

- Start a guest:

```
ALTER MACHINE ACTION=START WHERE HOSTNAME=node1hostname
```

- Stop a guest:

```
ALTER MACHINE ACTION=STOP WHERE HOSTNAME=node1hostname
```

- Restart a guest:

```
ALTER MACHINE ACTION=RESTART WHERE HOSTNAME=node1hostname
```

- Create a disk image:

```
ALTER MACHINE ACTION=CREATEIMAGE IMAGEFILE=example.img SIZE=2 WHERE
HOSTNAME=node1hostname


ALTER MACHINE ACTION=CREATEIMAGE IMAGEFILE=example.img SIZE=2
ZIPFILE=data.zip WHERE HOSTNAME=node1hostname


ALTER MACHINE ACTION=CREATEIMAGE IMAGEFILE=example.img SIZE=2
FILESYTEM=ext4 WHERE HOSTNAME=node1hostname


ALTER MACHINE ACTION=CREATEIMAGE IMAGEFILE=example.img SIZE=2
FILESYTEM=ext4 ZIPFILE=data.zip WHERE HOSTNAME=node1hostname
```


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

- Attach a disk image:

```
ALTER MACHINE ACTION=ATTACHDISK IMAGEFILE=/EXAVMIMAGES/example.img
MOUNTPATH=/u02 WHERE HOSTNAME=node1hostname


ALTER MACHINE ACTION=ATTACHDISK IMAGEFILE=/EXAVMIMAGES/example.img
MOUNTPATH=/u02 FILESYSTEM=ext4 WHERE HOSTNAME=node1hostname
```

- Detach a disk image:

```
ALTER MACHINE ACTION=DETACHDISK IMAGEFILE=example.img WHERE
HOSTNAME=node1hostname


ALTER MACHINE ACTION=DETACHDISK IMAGEFILE=example.img DELETE=true WHERE
HOSTNAME=node1hostname
```

- Change the CPU resource allocation:

```
ALTER MACHINE ACTION=SETVCPU MAXVCPU=12 VCPU=12 WHERE
HOSTNAME=node1hostname
```

- Change the RAM resource allocation:

```
ALTER MACHINE ACTION=SETVMEM VMEM=9 WHERE HOSTNAME=node1hostname
```

4.3.21.2 ALTER MACHINES

This command modifies attributes of a collection of Engineered System machines.

Syntax

```
ALTER MACHINES
   { ORADISKSIZE=oradisksize |
     ORADISKPATH=oradiskpath |
     IMAGEVERSION=imageversion }...
WHERE
   { CLUSTERNUMBER=clusternumber |
     CLUSTERNAME=clustername |
     CLUSTERID=clusterid |
     TYPE=GUESTS }
```

Arguments

You can specify the following arguments for the machine:

- `ORADISKSIZE`: Size (in GB) of the Oracle software volume (typically `/u01`). Default is 20, minimum is 10.
- `ORADISKPATH`: Mount point of the Oracle software volume. Default is `/u01`. For non-default values, take great care to ensure other paths are compatible with the setting.
- `IMAGEVERSION`: The version of the image to deploy on guest machines. This argument is only applicable to guest machines (`WHERE TYPE=GUEST`).

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

The following arguments are available in the WHERE clause:

- CLUSTERNUMBER: Cluster number associated with target machine in the OEDA XML file (es.xml), starting at 1
- CLUSTERNAME: Name of the cluster associated with target machine in the OEDA XML file (es.xml)
- CLUSTERID: ID of the cluster associated with target machine in the OEDA XML file (es.xml)
- TYPE=GUEST: Applies the change to all guest machines. This option is only applicable when setting the IMAGEVERSION.

4.3.21.3 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 |
       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, DOMU, KVM, or KVMGUEST.
- 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.

4.3.22 NETWORK

You can add, alter, delete, or list the networks defined in the XML configuration file.

- ADD NETWORK
- ALTER NETWORK
- ALTER NETWORKS
- DELETE NETWORK
- LIST NETWORKS

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

4.3.22.1 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
   [ MTU=mtu ]
   [ GATEWAY=gateway ]
   [ SSHENABLED={ true | false } ]
   [ MAC=mac_addr ]
   [ LINKSPEED=nic_speed ]
   [ SLAVE='slave' ]
   [ PKEY=pkey ]
   [ PKEYNAME=pkey_name ]
   [ STATUS = status ]
   [ LACP={ true | false } ]
   [ 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`: Specifies the type of network, which can be `ADMIN`, `CLIENT`, `BACKUP`, `PRIVATE`, `OTHER`, `INGESTVIP`, `REPLICATIONVIP`, or `ILOM`.
- `HOSTNAME`: Specifies the short host name for the network
- `IP`: Specifies the IP address of the network
- `NETMASK`: Specifies the network netmask
- `DOMAINNAME`: Specifies the domain name of the network
- `MASTER`: Specifies the primary adapter name
- `MTU`: Specifies the network MTU (maximum transmission unit). Not applicable for `PRIVATE` or `ILOM` networks.

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

- GATEWAY: Specifies the network subnet gateway
- SSHENABLED: TRUE indicates that SSH should be enabled for this adapter
- MAC: Specifies the MAC address
- LINKSPEED: Either 10000 or 25000. Only for X7-2 or later 2-socket compute nodes.
- SLAVE: Specifies a space-separated list of secondary devices for a bonded network, enclosed in single quotes
- PKEY: Specifies the pkey, for PRIVATE networks only
- PKEYNAME: Specifies the pkey name, for PRIVATE networks only
- STATUS: Specifies the status of the network
- LACP: TRUE indicates that LACP should be enabled
- VLANID: Specifies the VLAN identifier. Applicable if the network is a non-private VLAN network, or a VLAN network on a KVM guest.
- 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

4.3.22.2 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={ true | false } |
     MAC=mac_addr |
     MTU=mtu |
```


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

```
      LINKSPEED=nic_speed |
      SLAVE='slave' |
      PKEY=pkey |
      PKEYNAME=pkey_name |
      STATUS=status |
      LACP={ true | false } |
      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 primary adapter name
- `GATEWAY`: Specifies the network subnet gateway
- `SSHENABLED`: `TRUE` indicates that SSH should be enabled for this adapter
- `MAC`: Specifies the MAC address
- `MTU`: Specifies the network MTU (maximum transmission unit). Not applicable for `PRIVATE` or `ILOM` networks.
- `LINKSPEED`: Either `10000` or `25000`. Only for X7-2 or later 2-socket compute nodes.
- `SLAVE`: Specifies a space-separated list of secondary devices for a bonded network, enclosed in single quotes
- `PKEY`: Specifies the pkey, for `PRIVATE` networks only
- `PKEYNAME`: Specifies the pkey name, for `PRIVATE` networks only
- `STATUS`: Specifies the status of the network
- `LACP`: `TRUE` indicates that LACP should be enabled

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

- VLANID: Specifies the VLAN identifier. Applicable if the network is a non-private VLAN network, or a VLAN network on a KVM guest.
- 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 an 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

4.3.22.3 ALTER NETWORKS

This command modifies the VLAN number for the specified machines.

Syntax

```
ALTER NETWORKS
   VLANID=vlan_id
WHERE
   HOSTNAMES=hostnames
   NETWORKTYPE=network_type
```

Arguments

- VLANID: Specifies the VLAN identifier.
- HOSTNAMES: Specifies a comma-separated list of host names
- NETWORKTYPE: Specifies the type of network, which can be ADMIN, CLIENT, BACKUP, PRIVATE (storage private network), or CLUSTERPRIVATE.

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

Usage Notes

- During corresponding `MERGE ACTIONS` processing, the VLAN identifier is set even if the specified host does not have an existing VLAN.
- During corresponding `DEPLOY ACTIONS` processing, the storage cells reboot in parallel to enable the new VLAN identifier.

4.3.22.4 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

4.3.22.5 LIST NETWORKS

This command lists the network details.

Syntax

```
LIST NETWORKS
[ WHERE
     { ID=network_id |
       HOSTNAMES=ALL [ NETWORKTYPE=network_type ] |
       HOSTID=host_id [ NETWORKTYPE=network_type ] |
```


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

```
       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 ] } ]
```

Arguments

The following arguments are available in the `WHERE` clause:

- `ID` : Specifies the ID for the network object in the XML configuration file
- `HOSTNAMES=ALL` : Specifies all hosts
- `NETWORKTYPE` : Specifies the type of network, which can be `ADMIN`, `CLIENT`, `BACKUP`, `PRIVATE`, or `ILOM`.
- `HOSTID` : Specifies the host ID in the XML configuration file
- `HOSTNAME` : Specifies the host name for an existing compute node
- `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

4.3.23 PASSWORD

You can specify the default passwords.

- ALTER PASSWORD
  Modifies the OEDA properties that specifies the default passwords.

4.3.23.1 ALTER PASSWORD

Modifies the OEDA properties that specifies the default passwords.

Syntax

```
ALTER PASSWORD
   [ PASSWORD=new_password ]
WHERE
   USER={ ROOT | NONROOT }
```

Arguments

- `PASSWORD`: Specifies the new password.
- `USER`: Specifies the user type; either `ROOT` or `NONROOT`.

Usage Notes

If the `PASSWORD` argument is not specified, then the user is prompted to specify and verify the password interactively.

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

4.3.24 PROPERTY

You can set OEDA properties.

- ALTER PROPERTY
  Modifies an OEDA properties parameter.
- LIST PROPERTIES
  Displays information about OEDA properties.

4.3.24.1 ALTER PROPERTY

Modifies an OEDA properties parameter.

Syntax

```
ALTER PROPERTY
   NAME=property_name
   VALUE=value
WHERE
   [ OS=os_name ]
```

Arguments

- `NAME`: Specifies the name of the property to alter. See `LIST PROPERTIES` to display the available properties.
- `VALUE`: Specifies the new property value
- `OS`: Optionally specifies the operating system that is associated with the property. Permitted values are: `LINUX`, `SPARC`, `WINDOWS`, `MAC`, `OVM`, or `WEB`.

Usage Notes

If the `OS` argument is not specified, then the property is updated in the main `es.properties` file is updated.

4.3.24.2 LIST PROPERTIES

Displays information about OEDA properties.

Syntax

```
LIST PROPERTIES
[ WHERE
     OS=os_name ]
```

Arguments

- `OS`: Optionally specifies the operating system that is associated with the properties. Permitted values are: `LINUX`, `SPARC`, `WINDOWS`, `MAC`, `OVM`, or `WEB`.

Usage Notes

If the `OS` argument is not specified, then the properties in the main `es.properties` are displayed.

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

4.3.25 RACK

You can manage an Engineered System rack.

- ADD RACK
  Add a new rack to the Engineered System XML file.
- ALTER RACK
  Modify a rack in the Engineered System XML file.
- LIST RACKS
  Displays information about racks in the Engineered System XML file.
- LIST RACKDESCRIPTION
  Displays information about rack types supported by OEDA.

4.3.25.1 ADD RACK

Add a new rack to the Engineered System XML file.

Syntax

```
ADD RACK
   TYPE=ELASTIC
   MODEL={ X7 | X8 | X8M }
   SRCRACK=srcrack


SET ADMIN NAME=admin_name IP=admin_ip
SET LEAF1 NAME=leaf1_name IP=leaf1_ip
SET LEAF2 NAME=leaf2_name IP=leaf2_ip
[ SET SPINE NAME=spine_name IP=spine_ip ]
[ SET PDU1 NAME=pdu1_name IP=pdu1_ip ]
[ SET PDU2 NAME=pdu2_name IP=pdu2_ip ]
```

Arguments

- `TYPE`: Specifies the `ELASTIC` rack type.
- `MODEL`: Specifies the rack model. Permitted values are: `X7`, `X8`, or `X8M`.
- `SRCRACK`: Specifies the `es.xml` identifier of the existing rack, which is used as the template for the new rack.

You can set the following options for the new rack:

- `ADMIN` : Specifies the DNS name and IP address to use for the administration network switch. The DNS name may optionally include the domain name.
- `LEAF1` : Specifies the DNS name and IP address to use for the first leaf switch. The DNS name may optionally include the domain name.
- `LEAF2` : Specifies the DNS name and IP address to use for the second leaf switch. The DNS name may optionally include the domain name.
- `SPINE` : Optionally specifies the DNS name and IP address to use for the spine switch. The DNS name may optionally include the domain name.

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

- PDU1 : Optionally specifies the DNS name and IP address to use for the first PDU. The DNS name may optionally include the domain name.
- PDU2 : Optionally specifies the DNS name and IP address to use for the second PDU. The DNS name may optionally include the domain name.

4.3.25.2 ALTER RACK

Modify a rack in the Engineered System XML file.

Syntax

```
ALTER RACK
   ULOC={ ulocation | 0 | GENERATE }
WHERE
   RACKNUMBER=rack_number
   [ HOSTNAME=host_name ]


ALTER RACK
   { DESCRIPTION=description |
     RACKID=rack_type } ...
WHERE
   RACKNUMBER=rack_number
```

Arguments

You can modify the following attributes:

- ULOC: Sets the rack location (uloc) as follows:
  - ULOC=*ulocation*: Specifies a rack location for a specific host or switch
  - ULOC=0: Sets the rack location to 0 (zero) for all hosts and switches in the specified rack.
  - ULOC=GENERATE: Assigns valid rack locations to all hosts and switches in the specified rack.
- DESCRIPTION: Sets the rack description
- RACKID: Sets the rack type

The following arguments are available in the WHERE clause:

- RACKNUMBER: Specifies the number of the rack in the Engineered System XML file, starting at 1
- HOSTNAME: Specifies the host name for a specific host or switch

Usage Notes

If you use the ULOC argument and supply a specific rack location (ULOC=*ulocation*), then you must specify HOSTNAME=*host_name* in the WHERE clause.

4.3.25.3 LIST RACKS

Displays information about racks in the Engineered System XML file.

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

Syntax

```
LIST RACKS
```

4.3.25.4 LIST RACKDESCRIPTION

Displays information about rack types supported by OEDA.

Syntax

```
LIST RACKDESCRIPTION
[ WHERE
     TYPE=rack ]
```

Arguments

`TYPE`: Optionally specifies the rack model that you want to use to limit the command output.

Usage Notes

If the `WHERE` clause is omitted, all rack types are displayed.

4.3.26 REQUIREDFILES

You can display information about required files.

- LIST REQUIREDFILES
  Displays information about required installation files for Oracle Grid Infrastructure or Oracle Database.

4.3.26.1 LIST REQUIREDFILES

Displays information about required installation files for Oracle Grid Infrastructure or Oracle Database.

Syntax

```
LIST REQUIREDFILES
   { GRID |
     DATABASE |
     GIVERSION='<giversion>' |
     DBVERSION='<dbversion>' }
```

Arguments

- `GRID`: Displays information about required files for Oracle Grid Infrastructure
- `DATABASE`: Displays information about required files for Oracle Database
- `GIVERSION`: Specifies the Oracle Grid Infrastructure version including the release update (RU). For example: `12.1.0.2.170418`.
- `DBVERSION`: Specifies the Oracle Database version including the release update (RU). For example: `12.1.0.2.170418`.

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

Usage Notes

- Use `GIVERSION` or `DBVERSION` when no `es.xml` file is loaded.
- Use `LIST SOFTWAREVERSIONS { GRID | DATABASE }` to display information about software versions.

4.3.27 REMOTE

You can run a command on a remote host.

- EXEC REMOTE
  Runs a command on remote hosts.

4.3.27.1 EXEC REMOTE

Runs a command on remote hosts.

Syntax

```
EXEC REMOTE
   COMMAND='command'
WHERE
   HOSTNAMES='host_names'
```

Section Title

- `COMMAND`: Specifies the command to run on the specified hosts.
- `HOSTNAMES`: Specifies the comma-separated list of remote hosts.

4.3.28 SCAN

You can add, alter, delete, or list the SCANs for a cluster.

- ADD SCAN
- ALTER SCAN
- DELETE SCAN
- LIST SCANS

4.3.28.1 ADD SCAN

This command adds a SCAN to the an Oracle Grid Infrastructure cluster.

Syntax

```
ADD SCAN
   SCANNAME=scan_name
   [ SCANPORT=scan_port ]
   SCANIPS='scan_ip_addrs'
WHERE
   { CLUSTERNAME=cluster_name |
```


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

```
   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` : Optionally 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

4.3.28.2 ALTER SCAN

This command modifies the attributes of a SCAN for an Oracle Grid Infrastructure 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

4.3.28.3 DELETE SCAN

This command removes the SCAN for an Oracle Grid Infrastructure cluster.

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

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

4.3.28.4 LIST SCANS

This command lists the SCAN details for all Oracle Grid Infrastructure clusters or for an individual cluster.

Syntax

```
LIST SCANS
[ WHERE
     { ID=scan_id |
       CLUSTERNAME=cluster_name |
       CLUSTERNUMBER=cluster_number |
       CLUSTERID=cluster_id } ]
```

Arguments

The following arguments are available in the `WHERE` clause:

- `ID` : 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.

4.3.29 SECURITY

You can manage security settings.

- SET SECURITY
  This command performs security-related actions, like setting the non-default root password for the OEDACLI session or controlling password-based authentication.

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

4.3.29.1 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

Usage Notes

If the `HOSTNAMES` argument is not specified, then the command acts on the currently loaded Engineered System XML file.

4.3.30 SOFTWAREVERSIONS

You can check supported software versions.

- LIST SOFTWAREVERSIONS
  Displays information about supported versions of Oracle Grid Infrastructure and Oracle Database.

4.3.30.1 LIST SOFTWAREVERSIONS

Displays information about supported versions of Oracle Grid Infrastructure and Oracle Database.

Syntax

```
LIST SOFTWAREVERSIONS [ GRID | DATABASE ]
```

Arguments

- `GRID`: Displays information about supported versions of Oracle Grid Infrastructure
- `DATABASE`: Displays information about supported versions of Oracle Database

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

Usage Notes

- If neither argument (`GRID` or `DATABASE`) is specified, then `LIST SOFTWAREVERSIONS` displays information about supported software versions for Oracle Grid Infrastructure and Oracle Database.

4.3.31 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.

4.3.31.1 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.
- If the `HOSTNAMES` argument is not specified, then the command acts on the currently loaded Engineered System XML file.

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

Example 4-33 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 4-34 Deploy SSH keys for specified hosts

```
oedacli> DEPLOY SSHKEYS ENABLEPASSWORDLOGIN=false PASSWORD="password"
HOSTNAMES="dbm01adm05.example.com,dbm01adm06.example.com"
```

4.3.31.2 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.

Usage Notes

- If the `HOSTNAMES` argument is not specified, then the command acts on the currently loaded Engineered System XML file.

Example 4-35 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 4-36 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]
```


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

4.3.31.3 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.

4.3.32 STEPS

You can display the steps for a virtual `CREATE CLUSTER` deployment.

- LIST STEPS
  This command lists the steps for a virtual `CREATE CLUSTER` deployment.

4.3.32.1 LIST STEPS

This command lists the steps for a virtual `CREATE CLUSTER` deployment.

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` in other commands, such as 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

```
LIST STEPS
[ WHERE
   { CLUSTERNUMBER=cluster_number |
     CLUSTERNAME=cluster_name |
     ID=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

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

- `ID` : Specifies the `es.xml` ID of the cluster

Example 4-37 Listing the Steps for a Specific Cluster

```
oedacli> LIST STEPS WHERE CLUSTERNAME = Cluster-c1
{
  "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",
```


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

```
    "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"
  } ]
}
```

4.3.33 STORAGEPOOL

- ADD STORAGEPOOL
  This command adds a new Exascale storage pool.
- ALTER STORAGEPOOL
  This command alters an existing Exascale storage pool.
- DELETE STORAGEPOOL
  This command deletes an existing Exascale storage pool.
- LIST STORAGEPOOL
  This command displays details for an existing Exascale storage pool.
- LIST STORAGEPOOLS
  This command displays details for existing Exascale storage pools.

4.3.33.1 ADD STORAGEPOOL

This command adds a new Exascale storage pool.

Syntax

```
ADD STORAGEPOOL
   NAME = storagepoolname
   [ SIZE = size ]
   [ TYPE = storagetype ]
   CELLLIST = 'celllist'
[ WHERE
     EXASCALECLUSTER = exascaleclustername ]
```

Arguments

- `NAME`: Specifies the name of the Exascale storage pool.

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

- SIZE: Optionally, specifies the size of the storage pool in GB. If not specified, then the storage pool receives all of the available space in the specified cells.
- TYPE: Specify EF to create a flash-based storage pool using High Capacity Exadata storage servers. Otherwise, the storage pool media type is derived from the Exadata storage server type.
- CELLLIST: Specifies the list of Exadata storage servers to use for the storage pool. The value is a space-seperated list of cell names, or use ALL to specify all available cells in the OEDA configuration file (es.xml). Note that the value is also surrounded by single quotes (as displayed in the command syntax).

The following argument is available in the WHERE clause:

- EXASCALECLUSTER: Specifies the name for the Exascale cluster that contains the new storage pool.

Usage Notes

- When you use the ADD STORAGEPOOL command, there is no immediate check to validate the specified storage pool size. However, deployment fails with an error if there is insufficient space.
- The WHERE clause is optional if the OEDA configuration file (es.xml) contains only one Exascale cluster. Otherwise, the WHERE clause is mandatory.

4.3.33.2 ALTER STORAGEPOOL

This command alters an existing Exascale storage pool.

Syntax

```
ALTER STORAGEPOOL
   [ NAME = storagepoolname ]
   [ SIZE = size ]
   [ TYPE = storagetype ]
   [ ADDCELLS = 'addcelllist' ]
   [ DROPCELLS = 'dropcelllist' ]
WHERE
   [ EXASCALECLUSTER = exascaleclustername ]
   NAME = storagepoolname
```

Arguments

- NAME: Changes the name of the Exascale storage pool.
- SIZE: Changes the size (in GB) of the storage pool.
- TYPE: Specify EF to create a flash-based storage pool using High Capacity Exadata storage servers. Otherwise, the storage pool media type is derived from the Exadata storage server type.
- ADDCELLS: Specifies the list of Exadata storage servers to add to the storage pool. The value is a space-seperated list of cell names. Note that the value is also surrounded by single quotes (as displayed in the command syntax).
- DROPCELLS: Specifies the list of Exadata storage servers to remove from the storage pool. The value is a space-seperated list of cell names. Note that the value is also surrounded by single quotes (as displayed in the command syntax).

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

The following arguments are available in the `WHERE` clause:

- `EXASCALECLUSTER`: Specifies the name of the Exascale cluster that contains the storage pool. This argument is optional if the OEDA configuration file (`es.xml`) contains only one Exascale cluster.
- `NAME`: Specifies the name for the Exascale storage pool that you want to alter.

Usage Notes

- When you use the `ALTER STORAGEPOOL` command, there is no immediate check to validate the specified storage pool size. However, deployment fails with an error if there is insufficient space.

4.3.33.3 DELETE STORAGEPOOL

This command deletes an existing Exascale storage pool.

Syntax

```
DELETE STORAGEPOOL
WHERE
   [ EXASCALECLUSTER = exascaleclustername ]
   NAME = storagepoolname
```

Arguments

The following arguments are available in the `WHERE` clause:

- `EXASCALECLUSTER`: Specifies the name of the Exascale cluster that contains the storage pool. This argument is optional if the OEDA configuration file (`es.xml`) contains only one Exascale cluster.
- `NAME`: Specifies the name for the Exascale storage pool that you want to delete.

Usage Notes

- You cannot delete an Exascale storage pool while it provides resources to Oracle Grid Infrastructure or an Oracle database.

4.3.33.4 LIST STORAGEPOOL

This command displays details for an existing Exascale storage pool.

Syntax

```
LIST STORAGEPOOL
WHERE
   [ EXASCALECLUSTER = exascaleclustername ]
   NAME = storagepoolname
```

Arguments

The following arguments are available in the `WHERE` clause:

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

- `EXASCALECLUSTER`: Specifies the name of the Exascale cluster that contains the storage pool. This argument is optional if the OEDA configuration file (`es.xml`) contains only one Exascale cluster.
- `NAME`: Specifies the name for the Exascale storage pool that you want to display.

4.3.33.5 LIST STORAGEPOOLS

This command displays details for existing Exascale storage pools.

Syntax

```
LIST STORAGEPOOLS
[ WHERE
     EXASCALECLUSTER = exascaleclustername ]
```

Arguments

The following arguments are available in the `WHERE` clause:

- `EXASCALECLUSTER`: Specifies the name of an existing Exascale cluster.

4.3.34 STORAGETEMPLATES

- LIST STORAGETEMPLATES
  This command displays details for the Exascale storage templates.

4.3.34.1 LIST STORAGETEMPLATES

This command displays details for the Exascale storage templates.

Syntax

```
LIST STORAGEPOOLTEMPLATES
```

4.3.35 SWITCH

You can display and manage the switches.

- ADD SWITCH
  Add a new switch to a rack in the Engineered System XML file.
- ALTER SWITCH
- LIST SWITCHES

4.3.35.1 ADD SWITCH

Add a new switch to a rack in the Engineered System XML file.

Syntax

```
ADD SWITCH
WHERE
```


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

```
   { RACKID=rack_id |
     RACKNUMBER=rack_number }


SET ADMINNET NAME=admin_name, IP=admin_ip [ , GATEWAY=gateway,
NETMASK=netmask ]
SET DESCRIPTION VALUE=description
[ SET IBPARTITIONMEMBERSHIP VALUE=pkey_membership ]
[ SET ULOC VALUE=ulocation ]
[ SET TIMEZONE VALUE=timezone ]
[ SET DNSSERVERS IP1=dns_ip1, IP2=dns_ip2, IP3=dns_ip3 ]
[ SET NTPSERVERS IP1=ntp_ip1, IP2=ntp_ip2, IP3=ntp_ip3 ]
```

Arguments

The following arguments are available in the `WHERE` clause:

- `RACKID`: Specifies the rack identifier in the XML configuration file
- `RACKNUMBER`: Specifies the rack number in the XML configuration file

You can set the following options for the new rack:

- `ADMINNET`: Specifies the details for the administration network.
- `DESCRIPTION`: Specifies a description for the switch
- `IBPARTITIONMEMBERSHIP`: Specifies the switch pkey membership
- `ULOC`: Specifies physical rack location for the switch
- `TIMEZONE`: Specifies a valid time zone for the switch
- `DNSSERVERS`: Specifies the DNS servers to set in the switch configuration
- `NTPSERVERS`: Specifies the NTP servers to set in the switch configuration

4.3.35.2 ALTER SWITCH

This command alters attributes of an individual switch.

Syntax

```
ALTER SWITCH
   { DNSSERVERS='dns_servers' |
     NTPSERVERS='ntp_servers' |
     TIMEZONE=time_zone |
     DESCRIPTION=description |
     ULOC=ulocation |
     IBPARTITIONMEMBERSHIP=pkey_membership } ...
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

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

- NTPSERVERS: Specifies a comma-separated list of NTP servers that is enclosed in single quotes
- TIMEZONE: Specifies a valid time zone for the switch
- DESCRIPTION: Specifies a description for the switch
- ULOC: Specifies physical rack location for the switch
- IBPARTITIONMEMBERSHIP: Specifies the switch pkey membership

The following arguments are available in the WHERE clause:

- ID: Specifies the switch identifier in the XML configuration file
- HOSTNAME: Specifies the host name of the switch

4.3.35.3 LIST SWITCHES

This command displays 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 identifier in the XML configuration file
- HOSTNAME : Specifies the host name of the switch

Usage Notes

If you do not specify a WHERE clause, then the command displays the details for all switches.

4.3.36 SU

Controls use of the su command.

- SET SU
  This command enables and disables using su from the root user to run remote non-root commands from OEDACLI.

4.3.36.1 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 }
```


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

Arguments

- `ENABLE`: Specifies whether or not OEDACLI uses `su` from the `root` user to run remote non-`root` commands.

4.3.37 ULOC

You can display the location of rack items.

- LIST ULOC

4.3.37.1 LIST ULOC

This command displays the location of defined rack items and new rack items.

Syntax

```
LIST ULOC
   [ NEWCOMPUTES={ new_computes_count | 'new_computes_list' } ]
   [ NEWCELLS={ new_cells_count | 'new_cells_list' } ]
[ WHERE
     RACKUMBER=rack_number ]
```

Arguments

- `NEWCOMPUTES`: Optionally specifies the number of new compute nodes to be added, or a comma-separated list of new compute node host names. Default is 0.
- `NEWCELLS`: Optionally specifies the number of new storage cells to be added, or a comma-separated list of new storage cell host names. Default is 0.

The following arguments are available in the `WHERE` clause:

- `RACKNUMBER`: Specifies the rack number in the XML configuration file

Usage Notes

If you do not specify a `WHERE` clause, then the command displays the details for rack number 1.

4.3.38 VAULT

- ADD VAULT
  This command adds a new Exascale vault.
- ALTER VAULT
  This command alters an existing Exascale vault.
- DELETE VAULT
  This command deletes an existing Exascale vault.
- LIST VAULT
  This command displays details for an existing Exascale vault.
- LIST VAULTS
  This command displays details for existing Exascale vaults.

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

4.3.38.1 ADD VAULT

This command adds a new Exascale vault.

Syntax

```
ADD VAULT
   NAME = vaultname
   [ HC = hcsize ]
   [ EF = efsize ]
   [ XT = xtsize ]
[ WHERE
     EXASCALECLUSTER = exascaleclustername ]
```

Arguments

- `NAME`: Specifies the name of the Exascale vault.
- `HC`: Provisions the vault with the specified amount of storage space from the High Capacity (HC) storage pool.
- `EF`: Provisions the vault with the specified amount of storage space from the Extreme Flash (EF) storage pool.
- `XT`: Provisions the vault with the specified amount of storage space from the Extended (XT) storage pool.

The following argument is available in the `WHERE` clause:

- `EXASCALECLUSTER`: Specifies the name for the Exascale cluster that contains the new vault.

Usage Notes

- Each space provisioning argument can specify either a fixed amount of space, in GB or TB, or a percentage of the total size of the corresponding storage pool. The storage pool must already exist.
- To use a vault for Oracle Database or Oracle Grid Infrastructure file storage, you must specify at least one space provisioning argument. Vaults with no space provisioning are reserved for system usage.
- The `WHERE` clause is optional if the OEDA configuration file (`es.xml`) contains only one Exascale cluster. Otherwise, the `WHERE` clause is mandatory.

4.3.38.2 ALTER VAULT

This command alters an existing Exascale vault.

Syntax

```
ALTER VAULT
   [ NAME = vaultname ]
   [ HC = hcsize ]
   [ EF = efsize ]
   [ XT = xtsize ]
WHERE
```


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

```
   [ EXASCALECLUSTER = exascaleclustername ]
   NAME = vaultname
```

Arguments

- `NAME`: The new name of the Exascale vault (for renaming a vault).
- `HC`: Provisions the vault with the specified amount of storage space from the High Capacity (HC) storage pool.
- `EF`: Provisions the vault with the specified amount of storage space from the Extreme Flash (EF) storage pool.
- `XT`: Provisions the vault with the specified amount of storage space from the Extended (XT) storage pool.

The following arguments are available in the `WHERE` clause:

- `EXASCALECLUSTER`: Specifies the name of the Exascale cluster that contains the vault. This argument is optional if the OEDA configuration file (`es.xml`) contains only one Exascale cluster.
- `NAME`: Specifies the name of the Exascale vault that you want to alter.

Usage Notes

- Each space provisioning argument can specify either a fixed amount of space, in GB or TB, or a percentage of the total size of the corresponding storage pool.
- If you alter any vault space provisioning arguments, then ensure that you specify all of the required space provisioning arguments (previous settings are not retained).

4.3.38.3 DELETE VAULT

This command deletes an existing Exascale vault.

Syntax

```
DELETE VAULT
WHERE
   [ EXASCALECLUSTER = exascaleclustername ]
   NAME = vaultname
```

Arguments

The following arguments are available in the `WHERE` clause:

- `EXASCALECLUSTER`: Specifies the name of the Exascale cluster that contains the vault. This argument is optional if the OEDA configuration file (`es.xml`) contains only one Exascale cluster.
- `NAME`: Specifies the name of the Exascale vault that you want to delete.

Usage Notes

- You cannot delete an Exascale vault while it provides resources to Oracle Grid Infrastructure or an Oracle database.

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

4.3.38.4 LIST VAULT

This command displays details for an existing Exascale vault.

Syntax

```
LIST VAULT
WHERE
   [ EXASCALECLUSTER = exascaleclustername ]
   NAME = vaultname
```

Arguments

The following arguments are available in the `WHERE` clause:

- `EXASCALECLUSTER`: Specifies the name of the Exascale cluster that contains the vault. This argument is optional if the OEDA configuration file (`es.xml`) contains only one Exascale cluster.
- `NAME`: Specifies the name of the Exascale vault that you want to display.

4.3.38.5 LIST VAULTS

This command displays details for existing Exascale vaults.

Syntax

```
LIST VAULTS
[ WHERE
     EXASCALECLUSTER = exascaleclustername
     [ STORAGEPOOL = storagepoolname ]]
```

Arguments

The following arguments are available in the `WHERE` clause:

- `EXASCALECLUSTER`: Specifies the name of the Exascale cluster that contains the vault.
- `STORAGEPOOL`: Specifies the name of the Exascale storage pool for which you want to display associated vaults.

4.3.39 VIP

You can add, alter, delete, or list the Oracle Clusterware VIPs.

- ADD VIP
- ALTER VIP
- DELETE VIP
- LIST VIPS

4.3.39.1 ADD VIP

This command adds a VIP for an individual node in a cluster.

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

Syntax

```
ADD VIP
   NAME=vip_name
   DOMAINNAME=domain_name
   IP=vip_ip_addr
   [ NETMASK=netmask ]
   [ INTERFACE_NAME=interaface_name ]
   [ IPADDR_TYPE={ipv4|ipv6} ]
   [ NETNUM=clusterware_network_number ]
WHERE
   { HOSTNAME=host_name |
     CLUSTERNAME=cluster_name COMPUTENUMBER=compute_number |
     CLUSTERNUMBER=cluster_number COMPUTENUMBER=compute_number |
     CLUSTERID=cluster_id COMPUTENUMBER=compute_number }
```

Arguments

You must 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
- `NETMASK`: Optionally specifies the netmask (network mask) for the VIP
- `INTERFACE_NAME`: Optionally specifies the interface name associated with the VIP
- `IPADDR_TYPE`: Optionally specifies the IP address type for the VIP. Either `ipv4` or `ipv6`.
- `NETNUM`: Optionally specifies the clusterware network number associated with 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

This command does not add multiple VIPs to a node. It is typically used to re-create a VIP after a `DELETE VIP` command.

When adding a VIP to a compute node, the node must not already have a VIP defined.

4.3.39.2 ALTER VIP

This command alters the attributes of an existing VIP for an individual node in a cluster.

Syntax

```
ALTER VIP
   { NAME=vip_name |
```


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

```
     DOMAINNAME=domain_name |
     NETMASK=netmask |
     IP=vip_ip_addr } ...
WHERE
   { HOSTNAME=host_name |
     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
- `NETMASK`: Specifies the netmask (network mask) 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

4.3.39.3 DELETE VIP

This command removes a VIP from an individual node in a cluster.

Syntax

```
DELETE VIP
WHERE
   { ID=vip_id |
     VIPNAME=vip_name |
     VIPIP=vip_ip |
     VIPDOMAIN=vip_domain |
     HOSTNAME=host_name }
```

Arguments

The following arguments are available in the `WHERE` clause:

- `ID`: Specifies the ID for the VIP in the XML configuration file
- `VIPNAME`: Specifies the host name for the VIP
- `VIPIP`: Specifies the IP address associated with the VIP
- `VIPDOMAIN`: Specifies the domain name for the VIP
- `HOSTNAME`: Specifies the host name for the compute node in the cluster

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

4.3.39.4 LIST VIPS

This command lists the VIPs for all clusters or for an individual node in a cluster.

Syntax

```
LIST VIPS
[ WHERE
     { HOSTNAME=host_name |
       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.

4.3.40 VOLUME

You can display and manage storage volumes.

- ADD VOLUME
- DELETE VOLUME
- LIST VOLUMES

4.3.40.1 ADD VOLUME

This command adds a storage volume to a KVM guest.

Syntax

```
ADD VOLUME
   MOUNTPATH=mount_path
   SIZE=size
WHERE
   HOSTNAME=host_name
```

Arguments

You must specify the following attributes when adding a volume:

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

- `MOUNTPATH`: Specifies the mount point directory for the volume
- `SIZE`: Specifies the volume size in GB. Must be an integer value.

The following arguments are available in the `WHERE` clause:

- `HOSTNAME`: Specifies the host name for the KVM guest

Usage Notes

You can use this command to add a custom sized volume for the `/u01` filesystem.

4.3.40.2 DELETE VOLUME

This command deletes a storage volume from a KVM guest.

Syntax

```
DELETE VOLUME
WHERE
   HOSTNAME=host_name
   MOUNTPATH=mount_path
```

Arguments

The following arguments are required in the `WHERE` clause:

- `HOSTNAME`: Specifies the host name for the KVM guest
- `MOUNTPATH`: Specifies the mount point directory for the volume

4.3.40.3 LIST VOLUMES

This command lists the volumes defined for an individual KVM guest.

Syntax

```
LIST VOLUMES
WHERE
   HOSTNAME=host_name
```

Arguments

The following arguments are required in the `WHERE` clause:

- `HOSTNAME`: Specifies the host name for the KVM guest

4.3.41 XMLACTION

You can manage actions relating to the Engineered System XML file.

- ALTER XMLACTION
- DELETE XMLACTION
- LIST XMLACTION
- LIST XMLACTIONS

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

4.3.41.1 ALTER XMLACTION

This command alters a specific OEDACLI action.

Syntax

```
ALTER XMLACTION
   ID=action_identifer
   CMDID=command_identifer
   CMD=command
```

Arguments

- `ID`: Specifies the action identifier. See output from `LIST XMLACTIONS`.
- `CMDID`: Specifies the command identifier. See output from `LIST XMLACTIONS`.
- `CMD`: Specifies the command for the action.

4.3.41.2 DELETE XMLACTION

This command removes a specific OEDACLI action.

Syntax

```
DELETE XMLACTION
   ID=action_identifer
```

Arguments

`ID`: Specifies the action identifier. See output from `LIST XMLACTIONS`.

4.3.41.3 LIST XMLACTION

This command lists a specific OEDACLI action.

Syntax

```
LIST XMLACTION
   ID=action_identifer
```

Arguments

`ID`: Specifies the action identifier. See output from `LIST XMLACTIONS`.

4.3.41.4 LIST XMLACTIONS

This command lists the current actions created by OEDACLI.

Syntax

```
LIST XMLACTIONS
```


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

5

Installing Oracle Exadata 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 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 and Installing Oracle Exadata Rack
- Acclimating Oracle Exadata
  Oracle recommends an acclimation period of 24 hours for Oracle Exadata.
- 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.

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

- 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 photocopy machine, 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*

5.2 Unpacking and Installing Oracle Exadata Rack

Oracle Rack Cabinet 1242 is the rack cabinet used for all Oracle Exadata Rack models since Oracle Exadata X7.

Use the following resources for information on unpacking and installing Oracle Exadata Rack:

- For an outline of the recommended procedures, see the *Oracle Rack Cabinet 1242 Getting Started Guide* at https://docs.oracle.com/cd/E85660_01/html/E85665/index.html.
- For additional details, consult the *Oracle Rack Cabinet 1242 User's Guide* at https://docs.oracle.com/en/servers/options/rack-cabinet-1242/user-guide/index.html.

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

5.3 Acclimating Oracle Exadata

Oracle recommends an acclimation period of 24 hours for Oracle Exadata.

The recommended acclimation period applies to all hardware, including racks and cables. Acclimation to the surrounding environment prevents damage as a result of condensation.

5.4 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.

5.4.1 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.

5.4.2 Connecting Power Cords

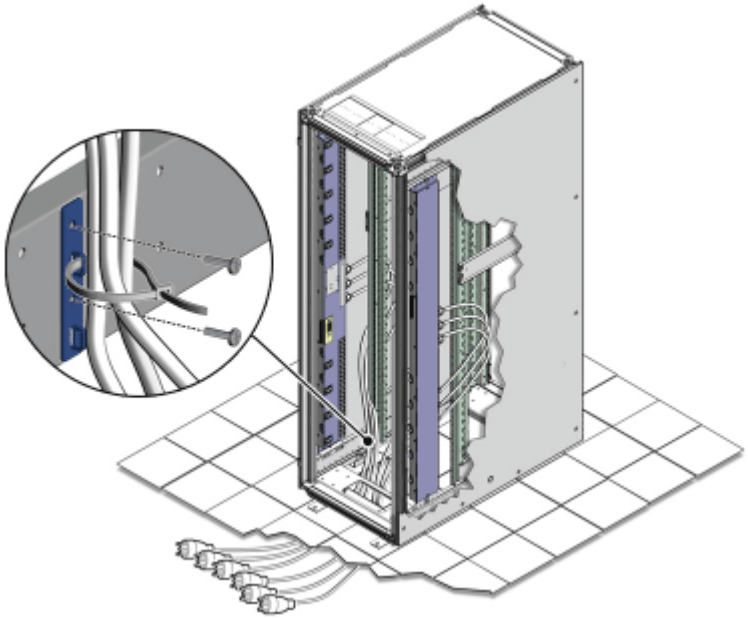
The following procedure describes how to connect power cords to Oracle Exadata Rack:

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 5-1.

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

5. Secure the power cords in bundles as shown in Figure 5-1.
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 5-1 Power Cord Routing Example



5.4.3 Powering on Oracle Exadata Rack

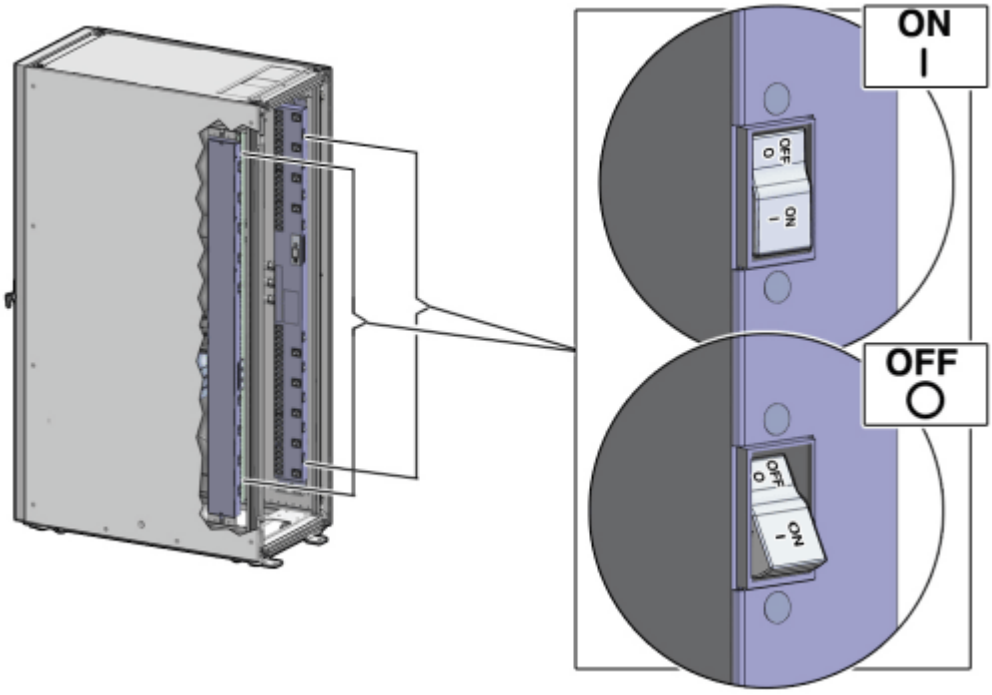
This procedure describes how to power on Oracle Exadata Rack.

> **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 (|) side of the toggle switch. When the breaker is in the ON position, the breakers are flush with the side of the PDU.

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

Figure 5-2 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 Management Network 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
   - Cisco 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.

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

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.

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

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 Management Network Switch
  You must perform an initial configuration of the Management Network 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.

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

- Verifying the InfiniBand Network Fabric Network
  This procedure describes how to verify the InfiniBand Network Fabric network.
- Customizing the Database Server File System Layout
  You can customize the file system layout on Exadata KVM hosts and bare-metal database servers on systems that run Oracle Exadata System Software release 26.1.0 or later and use system hardware with RoCE Network Fabric (X8M-2 or later).
- Customizing the Guest Image
  You can customize the guest image for Exadata guest VMs that run Oracle Exadata System Software release 26.1.0 or later.
- Imaging a New System
  You have different options for preparing an Oracle Exadata system for deployment.
- Performing Initial Elastic Configuration of Oracle Exadata
  This topic provides background information on elastic configuration and describes how to perform the initial configuration of Oracle Exadata.
- Adding Additional Elastic Nodes to an Existing Rack
  You can add nodes to an existing rack using elastic configuration.
- Using the OEDA genPasswordHash Utility
  You can configure OEDA to use customized OS user passwords by using the `genPasswordHash.sh` utility.
- 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.
- Using IPv4 and IPv6 in an Existing Deployment
  You can use OEDACLI to reconfigure an existing system with IPv4, IPv6, or both.

Related Topics

- Oracle Exadata Database Machine Setup/Configuration Best Practices (My Oracle Support Doc ID KB176594)

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

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

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. OEDA supports IPv6 addresses.

   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 the Cisco Nexus 9336C-FX2 Switch
  The RoCE Network Fabric switch supplied with the engineered system rack is minimally configured during installation.
- Configuring the RoCE Network Fabric Switches to Enable Exadata Secure RDMA Fabric Isolation
  You can configure your RoCE Network Fabric switches to enable Exadata Secure RDMA Fabric Isolation.
- Configuring Sun Datacenter InfiniBand Switch 36 Switch

6.2.1 Configuring the Cisco Nexus 9336C-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.

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

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 msecs
   ...
   ```

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
   ```


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

```
  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
```


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

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]:
```


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

```
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 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.

    To find the available RDMA Network Fabric patches, search for '**RDMA network switch**' in My Oracle Support document KB153930. Download and use the latest patch for your Oracle Exadata System Software release.

    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 (`leaf`): `roce_leaf_switch.cfg`
    - Multi-rack spine (`mspine`): `roce_spine_switch_multi.cfg`
    - Multi-rack leaf (`mleaf`): `roce_leaf_switch_multi.cfg`

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

- Single rack leaf with Secure Fabric support (`sfleaf`): `roce_sf_leaf_switch.cfg`
- Multi-rack leaf with Secure Fabric support (`msfleaf`): `roce_sf_leaf_switch_multi.cfg`
- Single rack leaf configured with 23 host ports (`leaf23`): `roce_leaf_switch_23hosts.cfg`
- Multi-rack leaf configured with 23 host ports (`mleaf23`): `roce_leaf_switch_23hosts_multi.cfg`
- Multi-rack leaf configured with 14 inter-switch links (`mleaf_u14`): `roce_leaf_switch_14uplinks_multi.cfg`
- Multi-rack leaf configured with 14 inter-switch links and with Secure Fabric support (`msfleaf_u14`): `roce_sf_leaf_switch_14uplinks_multi.cfg`
- Multi-rack leaf configured with 23 host ports and 13 inter-switch links (`mleaf23_u13`): `roce_leaf_switch_23hosts_13uplinks_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 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.

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

Related Topics

- Exadata Database Machine and Exadata Storage Server Supported Versions (My Oracle Support Doc ID KB153930)

6.2.2 Configuring the RoCE Network Fabric 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 9336C-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 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.

6.2.3 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.

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

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 document KB153930 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
    ```


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

```
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)
```


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

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 KB153930)

6.3 Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack

Oracle Exadata X3-2 systems and Oracle Exadata X2-2 systems have three Sun Datacenter InfiniBand Switch 36 switches. Starting with Oracle Exadata X4-2, Oracle Exadata Systems have two Sun Datacenter InfiniBand Switch 36 switches.

> **Note**
>
> This procedure does not apply to Oracle Exadata 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 Two-Socket Systems (X3-2 and later): rack unit 20 (U20) and rack unit 22 (U22)
- Oracle Exadata X2-2 racks: rack unit 20 (U20) and rack unit 24 (U24)
- Oracle Exadata 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:

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

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 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.

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

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 Half Racks with Sun Fire X4170 Oracle Database Servers, and Oracle Exadata 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 System Overview* for hardware component information
- Cabling tables in *Oracle Exadata System Overview*

6.4 Configuring the Management Network Switch

You must perform an initial configuration of the Management Network 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

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

- 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
   zdlra1sw-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, Fortooine Revision 1.40
   ```


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

```
...

Configuration register is 0x2102

Switch#
```

The version of the Cisco 4948E-F switch firmware purchased and shipped by Oracle with Oracle Exadata 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.
```


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

```
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
exa1sw-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.

```
exa1sw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
exa1sw-ip(config)# enable password password
exa1sw-ip(config)# enable secret password
exa1sw-ip(config)# end
exa1sw-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.

```
exa1sw-ip#configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
exa1sw-ip(config)#line vty 0 15
exa1sw-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
exa1sw-ip(config)#end
```


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

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.

```
exa1sw-ip(config-line)#no password
exa1sw-ip(config-line)#end
exa1sw-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:

```
exa1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
exa1sw-ip(config)# crypto key generate rsa
% You already have RSA keys defined named exa1sw-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]
exa1sw-ip(config)# username admin password 0 welcome1
exa1sw-ip(config)# line vty 0 15
exa1sw-ip(config-line)# transport input ssh
exa1sw-ip(config-line)# exit
exa1sw-ip(config)# aaa new-model

exa1sw-ip(config)# ip ssh time-out 60
exa1sw-ip(config)# ip ssh authentication-retries 3
exa1sw-ip(config)# ip ssh version 2
exa1sw-ip(config)# end
*Sep 15 14:26:37.045: %SYS-5-CONFIG_I: Configured from console by console
exa1sw-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 in effect, *hours-offset* is the hours offset from UTC, and *minutes-offset* is the minutes offset from UTC.

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

- 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 monthhh: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:

```
exa1sw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
exa1sw-ip(config)# clock timezone EST -5
exa1sw-ip(config)# clock summer-time EDT recurring
exa1sw-ip(config)# end
exa1sw-ip# clock set 21:00:00 August 09 2018
exa1sw-ip# write memory
Building configuration...
Compressed configuration from 3784 bytes to 1465 bytes [OK ]
exa1sw-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:

```
exa1sw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
exa1sw-ip(config)# ntp server 10.7.7.32 prefer
exa1sw-ip(config)# ntp server 198.51.100.19
exa1sw-ip(config)# end
*Sep 15 14:51:08.665:%SYS-5-CONFIG_I:Configured from console by console
exa1sw-ip# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes [OK ]
exa1sw-ip# show ntp status
<output will vary per network>
     .
exa1sw-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.

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

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:

```
exa1sw-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:
>
> ```
> exa1sw-ip#configure terminal
> Enter configuration commands, one per line. End with CNTL/Z.
> exa1sw-ip(config)# no ip default-gateway 10.7.7.1
> exa1sw-ip(config)# end
> exa1sw-ip#
> *Sep 15 14:13:26.013: %SYS-5-CONFIG_I: Configured from console by
> console
> exa1sw-ip(config)# write memory
> Building configuration...
> Compressed configuration from 2502 bytes to 1085 bytes[OK]
> ```

16. Save the current configuration.

```
exa1sw-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:

```
exa1sw-ip# exit

exa1sw-ip con0 is now available
```


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

```
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 Rack
- 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.

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

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 msecs
   ...
   ```

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 ----
   ```


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

```
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
```


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

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
```


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

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
```

4. View the modified time.

```
test123sw-adm0(config)# show clock
12:46:22.692 PST Thu Aug 31 2017
Time source is NTP
```

5. Save the configuration.

```
test123sw-adm0# copy running-config startup-config
[########################################] 100%
Copy complete.
```


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

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 Management Network Switch. Note that for full rack systems, you cannot connect PDUs to the Management Network 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:

1. Power off the PDU.
2. Disconnect the power leads from the power source.
3. Select an unused LAN network connection from the available network connections on the laptop as follows:
   - For Oracle Solaris:
     a. Log in as a super user.
     b. Use the `dladm show-link` command to find an unused Ethernet interface.
     c. Use the `ifconfig -a` command to determine which interface is being used.
     d. 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.

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

> **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 Management Network Switch. The equipment currently using the Management Network 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 Management Network Switch.
13. Reconnect the two cables that were originally connected to the Management Network Switch.

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

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

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

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 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

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

- 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
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

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

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.

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

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

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

- 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
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.

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

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.

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

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.

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

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

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

Table 6-18 (Cont.) 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 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
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

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

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
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.

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

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.

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

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
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.

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

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

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

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 2	0	17	22
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

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

- 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.

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

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
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.

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

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
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.

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

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

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

Table 6-42 (Cont.) 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 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
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

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

- 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
B	Module 1, phase 1	0	0	0
B	Module 1, phase 2	0	5	7

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

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 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.

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

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
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

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

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
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
A	Module 1, phase 2	0	35	37
A	Module 1, phase 3	0	35	37

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

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
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

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

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

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

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
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

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

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

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

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

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

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

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

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.

Perform the following checks on each Exadata Storage Server:

1. Connect to the storage server console.

   Use one of the following connection options:

   - Attach a serial cable between a console and the SER MGT port on the server.

     Connections through the modular jack serial port must use shielded cables to comply with radio frequency emission limits.

     Set the terminal device communication properties to these values: 9600 baud, 8 bit, no parity, 1 stop bit.

     After you connect to ILOM, run the `start /SP/console` command to connect to the host console.

     For more details, see Connect to Oracle ILOM.

   - Attach an Ethernet cable to the NET MGT port on the server and then connect using a workstation on the network.

     After you connect to ILOM, run the `start /SP/console` command to connect to the host console.

     For more details, see Connect to Oracle ILOM.

   - Connect using the KVM switch, if one is available.

2. Log in to the storage server as the `root` user.

> **Note**
>
> If you do not have the password for the `root` user, then contact Oracle Support Services.

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

3. 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
```

4. 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.

5. 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
```


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

```
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
```


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

```
Slot Number:11
Firmware state:Online, Spun Up
```

6. If applicable, verify the PMEM modules are visible.

   ```
   # cellcli -e list physicaldisk where disktype=PMEM
   ```

7. 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 re-seat the module and flash card. If the problem continues, then replace the card with the on-site spare.

8. 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 re-seat 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.

9. Log out of the system using the following command:

   ```
   # logout
   ```

6.8 Checking Oracle Exadata Database Servers

Before installing the software, check the configuration of the Oracle Exadata Database Servers.

Perform the following checks on each Exadata database server:

1. Connect to the database server console.

   Use one of the following connection options:

   - Attach a serial cable between a console and the SER MGT port on the server.

     Connections through the modular jack serial port must use shielded cables to comply with radio frequency emission limits.

     Set the terminal device communication properties to these values: 9600 baud, 8 bit, no parity, 1 stop bit.

     After you connect to ILOM, run the `start /SP/console` command to connect to the host console.

     For more details, see Connect to Oracle ILOM.

   - Attach an Ethernet cable to the NET MGT port on the server and then connect using a workstation on the network.

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

After you connect to ILOM, run the `start /SP/console` command to connect to the host console.

For more details, see Connect to Oracle ILOM.

- Connect using the KVM switch, if one is available.

2. Log in as the `root` user to the database server.

> **Note**
>
> If you do not have the password for the `root` user, then contact Oracle Support Services.

3. 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
```

4. Change to the `/opt/oracle.SupportTools/firstconf` directory.

```
cd /opt/oracle.SupportTools/firstconf
```

5. 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.

6. 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/checkhwfw.out

# more /tmp/checkhwfw.out
```


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

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.

7. 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  -
     ```


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

```
------------------------------------------------------------------------
------

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
```

8. 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
     ```


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

9. (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.

10. 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
    ```

11. Log out of the system using the following command:

    ```
    # logout
    ```

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 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_identifier'" | grep "system_identifier =" > /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_identifier= \
   "Exadata Database Machine rack_type serial_number"'
   ```


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

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 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 document KB452291.

Related Topics

- Verify RoCE Cabling on Oracle Exadata Database Machine Servers (My Oracle Support Doc ID KB452291)

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

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 Half Rack or Oracle Exadata 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)
   ```


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

```
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 Customizing the Database Server File System Layout

You can customize the file system layout on Exadata KVM hosts and bare-metal database servers on systems that run Oracle Exadata System Software release 26.1.0 or later and use system hardware with RoCE Network Fabric (X8M-2 or later).

Customizing the file system layout allows for more flexible storage configurations to meet diverse application requirements. To do this, you must design a custom layout and apply it during the system imaging process.

- Creating the Custom File System Layout
- Applying the Custom File System Layout

6.12.1 Creating the Custom File System Layout

You can create a custom file system layout for Exadata database servers using a JSON-formatted file named `fslayout.json`. The file contains a list of file system mount points and corresponding sizes conforming to the following general format:

```
{
  "mount-point1": {"size": "size-value1"}[,
  "mount-pointN": {"size": "size-valueN"}]...
}
```


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

For example:

```
{
  "/": {"size": "90G"},
  "/tmp": {"size": "3G"},
  "/var": {"size": "50G"},
  "/var/log": {"size": "50G"},
  "/var/log/audit": {"size": "20G"},
  "/home": {"size": "30G"},
  "/usr": {"size": "30G"},
  "/etc": {"size": "30G"},
  "/root": {"size": "30G"},
  "/opt": {"size": "200G"},
  "/EXAVMIMAGES": {"size": "300G"},
  "/project/data": {"size": "250G"}
}
```

The custom file system layout defined in `fslayout.json` file is subject to the following requirement and restrictions:

- Each mount point (*mount-pointN*) must be an absolute mount path that starts with a slash character (/). The path must not contain empty path elements (no // segments). For example, `/var/log` is valid, but `var/log` and `/var//log` are not.
- The system canonicalizes mount paths by converting elements to lowercase (except for `/EXAVMIMAGES`) and removing trailing slashes (except for the single root path `/`). After canonicalization, duplicate mount paths are not permitted.
- The following mount paths are not permitted: `/proc`, `/sys`, `/dev`, `/run`, `/srv`, `/sbin`, `/media`, `/lib`, `/lib64`, `/bin`, `/selinux`, `/boot`, `/boot/efi`.
- Nested mount paths are not permitted under: `/usr`, `/etc`, `/root`, `/opt`, `/EXAVMIMAGES`.
- The following nested mount paths are explicitly allowed `/var`, `/var/log`, and `/var/log/audit`.
- Each size value (*size-valueN*) must be a positive number and unit suffix, either `K` (kilobyte), `M` (megabyte), or `G` (gigabyte).
- If you specify an entry for `/EXAVMIMAGES`, the minimum permitted size is `300G`.
- During validation, a warning is generated if you specify a size value that is smaller than expected. If you proceed beyond the warning, the system automatically adjusts the file system size if required.

6.12.2 Applying the Custom File System Layout

After you create the JSON file (`fslayout.json`) containing the custom file system layout, you can apply it to your Exadata database servers during system imaging. In all imaging modes (PXE, NFS, and ISO imaging), `fslayout.json` is fetched and validated during the `initrd` stage. If validation succeeds, imaging continues to apply the custom file system layout. If validation fails, imaging stops and prints errors to the serial console so you can fix the file and retry.

The following describes how to apply `fslayout.json` using each imaging mode.

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

PXE imaging

1. Place `fslayout.json` in the same directory with the `preconf.csv` file on an NFS or HTTP server that is accessible from the server being imaged using PXE boot.
2. Instruct `initrd` to fetch `fslayout.json` at boot time by adding `fs-layout=.../fslayout.json` to the PXE append line.

   For example:

   ```
   append initrd=initrd.img root=/dev/ram0 ramdisk_size=150000 \
   preconf=http://198.51.100.91/export/exadata_images/config/preconf.csv \
   fs-layout=http://198.51.100.91/export/exadata_images/config/fslayout.json
   ```

3. Boot the server using PXE.

NFS-based imaging (using DHCP)

1. Place `preconf.csv` and `fslayout.json` in the same NFS directory as the imaging ISO.

   For example:

   ```
   /export/exadata_images/config/
     preconf.csv
     fslayout.json
   ```

2. When the imaging process detects `preconf.csv`, it automatically looks for `fslayout.json` in the same directory. If found, the file is validated and applied.

ISO-based imaging (no DHCP)

1. Include `preconf.csv` and `fslayout.json` in the ISO image.

   For example:

   ```
   # mkisofs -l -input-charset ascii -o preconf.iso preconf.csv fslayout.json
   ```

2. Attach the ISO as virtual media and boot the server. During boot, `initrd` reads the ISO, detects `fslayout.json`, and then proceeds to validate and apply it.

6.13 Customizing the Guest Image

You can customize the guest image for Exadata guest VMs that run Oracle Exadata System Software release 26.1.0 or later.

Oracle Exadata System Software release 26.1.0 introduces a gold guest image workflow allowing Exadata administrators to create, customize, seal, clone, and deploy customized guest images for Exadata database server virtual machines (VMs).

Using `libguestfs` tools (such as `guestfish`, `virt-resize`, and `virt-customize`), administrators can clone and modify a baseline system image, install application binaries and data, create or resize file systems, and perform other offline customizations. After sealing the gold image and making it available across the Exadata environment using Exascale storage, administrators can use Oracle Exadata Deployment Assistant (OEDA) to quickly and consistently provision new guests using the customized images.

Before you begin, consider the following requirements, recommendations, and restrictions:

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

- The Exadata system must use Oracle Exadata System Software release 26.1.0 or later. Likewise, the OEDA tools and baseline system image must be from Oracle Exadata System Software release 26.1.0 or later.
- The Exadata system must be configured to use Exascale storage to host the guest image files.
- Image modification is performed using `guestfish` and other `libguestfs` tools. To avoid potential issues with Oracle Exadata System Software updates on Exadata servers, Oracle recommends using a separate (non-Exadata) Linux host to perform image modification.
- Exadata system images contain a logical volume (LV) named `LVDoNotRemoveOrUse`. This LV reserves space for maintenance operations performed by Exadata tools. If you add or alter a LV associated with a file system mounted on `/`, `/var`, `/opt`, `/etc`, `/root`, or `/usr`, recalculate and resize `LVDoNotRemoveOrUse`. First-boot validation can fail if `LVDoNotRemoveOrUse` is missing or undersized.

  Calculate the required size of `LVDoNotRemoveOrUse` from the eligible mounted file systems in the root volume group. The eligible mount points are `/`, `/var`, `/opt`, `/etc`, `/root`, and `/usr`. For each eligible mounted LV, the required reservation for that LV is 4% of the LV size rounded up to the next whole GiB (minimum 1 GiB). Then, sum the required reservations for each LV to determine the total required size of `LVDoNotRemoveOrUse`. Round each LV's reserve up before adding the values together. Do not calculate 4% of the combined LV sizes and round only once.

  ```
  Required LVDoNotRemoveOrUse size =
  sum(max(ceil(size(LV) * 0.04), 1 GiB)) for each eligible mounted LV
  ```

- Determine whether the image uses BIOS or UEFI boot before running bootloader commands under `guestfish`. For BIOS images, mount `/boot` and run `grub2-install /dev/sda`. For UEFI images, mount `/boot` and `/boot/efi`; do not run BIOS-only bootloader commands.
- Do not create nested mounts under `/opt`, `/etc`, `/root`, or `/usr`.
- Do not make separate file systems for `/proc`, `/sys`, `/dev`, `/run`, `/srv`, `/sbin`, `/media`, `/lib`, `/lib64`, `/bin`, or `/selinux`.

The following topics show representative examples of common guest image customizations. Use the commands and examples contained therein as the basis for your customizations:

- Customizing the Guest Image to Expand an Existing File System
- Customizing the Guest Image to Add a Storage Volume and File System

6.13.1 Customizing the Guest Image to Expand an Existing File System

The following example workflow shows how to customize a guest image to expand an existing file system. You can adapt the following to perform similar customizations in your environment.

1. Prepare a host with `libguestfs` packages and choose a staging server.

   Modification of the baseline Exadata system image is performed using the `libguestfs` tools (such as `guestfish`, `virt-resize`, and `virt-customize`). To avoid potential issues with Oracle Exadata System Software updates, Oracle recommends using a separate Linux host to perform image modification. You can use `dnf` to install the required `libguestfs` packages. For example:

   ```
   # dnf install libguestfs-xfs libguestfs-tools
   ```


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

After image modification, the rest of the workflow is performed as `root` on an Exadata KVM host, also known as the staging server. Choose a suitable KVM host as the staging server. In this example workflow, the staging server is `exa01adm05`.

2. Select the Exadata base system image being customized and copy it to the `libguestfs` server.

> **Note**
>
> The base system image must be from Oracle Exadata System Software release 26.1.0 or later.

In this example workflow, the baseline Exadata system image is `System.first.boot.26.1.0.0.0.260515.rtg.img`.

3. On the `libguestfs` server, examine the baseline Exadata system image.

> **Note**
>
> On the `libguestfs` server, set `LIBGUESTFS_BACKEND=direct` in the Linux shell environment to avoid potential configuration and permission issues with `libguestfs` tools.

```
# ls -la ./System.first.boot.26.1.0.0.0.260515.rtg.img
-rw-r-----. 1 root root 106622563328 Apr 27 12:03
System.first.boot.26.1.0.0.0.260515.rtg.img
# export LIBGUESTFS_BACKEND=direct
# virt-filesystems -a ./System.first.boot.26.1.0.0.0.260515.rtg.img --all -
h -l
Name                                Type       VFS     Label MBR Size
Parent
/dev/sda1                           filesystem xfs     BOOT  -   412M -
/dev/sda2                           filesystem unknown -     -   256M -
/dev/VGExaDbDomU/LVDbHome           filesystem xfs     HOME  -   3.9G -
/dev/VGExaDbDomU/LVDbKdump          filesystem xfs     KDUMP -   20G  -
/dev/VGExaDbDomU/LVDbSwap1          filesystem swap    SWAP  -   16G  -
/dev/VGExaDbDomU/LVDbSys1           filesystem xfs     DBSYS -   15G  -
/dev/VGExaDbDomU/LVDbSys2           filesystem unknown -     -   15G  -
/dev/VGExaDbDomU/LVDbTmp            filesystem xfs     TMP   -   2.9G -
/dev/VGExaDbDomU/LVDbVar1           filesystem xfs     VAR   -   1.9G -
/dev/VGExaDbDomU/LVDbVar2           filesystem unknown -     -   2.0G -
/dev/VGExaDbDomU/LVDbVarLog         filesystem xfs     DIAG  -   18G  -
/dev/VGExaDbDomU/LVDbVarLogAudit    filesystem xfs     AUDIT -   924M -
/dev/VGExaDbDomU/LVDoNotRemoveOrUse filesystem unknown -     -   2.5G -
/dev/VGExaDbDomU/LVDbHome           lv         -       -     -   4.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbKdump          lv         -       -     -   20G  /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbSwap1          lv         -       -     -   16G  /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbSys1           lv         -       -     -   15G  /dev/
VGExaDbDomU
```


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

```
/dev/VGExaDbDomU/LVDbSys2           lv         -       -      -   15G  /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbTmp            lv         -       -      -   3.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbVar1           lv         -       -      -   2.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbVar2           lv         -       -      -   2.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbVarLog         lv         -       -      -   18G  /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbVarLogAudit    lv         -       -      -   1.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDoNotRemoveOrUse lv         -       -      -   2.5G /dev/
VGExaDbDomU
/dev/VGExaDbDomU                    vg         -       -      -   99G  /dev/
sda3
/dev/sda3                           pv         -       -      -   99G  -
/dev/sda1                           partition  -       -      -
512M /dev/sda
/dev/sda2                           partition  -       -      -
256M /dev/sda
/dev/sda3                           partition  -       -      -
99G  /dev/sda
/dev/sda                            device     -       -      -   99G  -

# virt-cat -a ./System.first.boot.26.1.0.0.0.260515.rtg.img /etc/fstab
libguestfs: error: mount_ro: mount: /crashfiles: mount point is not a
directory
virt-cat: some filesystems could not be mounted (ignored)
LABEL=BOOT              /boot           xfs
defaults,nodev,nosuid           0 0
LABEL=DBSYS             /               xfs
defaults                        0 0
LABEL=SWAP              swap            swap
defaults                        0 0
LABEL=KDUMP             /crashfiles     xfs
defaults,nodev                  0 0
LABEL=TMP               /tmp            xfs
defaults,nodev,nosuid           0 0
LABEL=HOME              /home           xfs
defaults,nodev,nosuid           0 0
LABEL=VAR               /var            xfs
defaults,nodev                  0 0
LABEL=DIAG              /var/log        xfs
defaults,nodev,nosuid           0 0
LABEL=AUDIT             /var/log/audit  xfs
defaults,nodev,noexec,nosuid    0 0
tmpfs                   /dev/shm        tmpfs
defaults,nodev,size=30831m      0 0

# qemu-img info ./System.first.boot.26.1.0.0.0.260515.rtg.img
image: System.first.boot.26.1.0.0.0.260515.rtg.img
file format: raw
virtual size: 99.3 GiB (106622563328 bytes)
disk size: 5.81 GiB
```


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

Take note of the internal structure of the baseline image. In particular, note the size of the image and location of the physical volume (`/dev/sda3`), which will be expanded in order to accommodate the desired file system expansion.

4. Create an expanded clone of the baseline system image:

```
# qemu-img create -f raw custom.img 150G
Formatting 'custom.img', fmt=raw size=161061273600

# virt-resize --expand /dev/sda3 ./
System.first.boot.26.1.0.0.0.260515.rtg.img  ./custom.img
[   0.0] Examining System.first.boot.26.1.0.0.0.260515.rtg.img
  25% ...
 100% ...
**********

Summary of changes:

/dev/sda1: This partition will be left alone.

/dev/sda2: This partition will be left alone.

/dev/sda3: This partition will be resized from 98.5G to 149.2G.  The LVM
PV
on /dev/sda3 will be expanded using the ‘pvresize’ method.

**********
[  11.0] Setting up initial partition table on custom.img
[  22.2] Copying /dev/sda1
[  23.1] Copying /dev/sda2
[  23.4] Copying /dev/sda3
...
[ 137.1] Expanding /dev/sda3 using the ‘pvresize’ method

Resize operation completed with no errors.  Before deleting the old disk,
carefully check that the resized disk boots and works correctly
```

In the example, the size of the expanded baseline image is 150 GB (`150G`). You can choose any size that accommodates the base image and desired customizations.

5. Customize the base image.

You can use `guestfish` to customize a guest image. It is recommended to prepare a script with the desired customizations and then run `guestfish` to apply the scripted customizations.

> **Note**
>
> You are responsible for all image customizations that you perform. Oracle will not support questions or issues with specific image customizations. If a guest fails, and Oracle suspects the failure may be caused by an image customization, then Oracle may ask for the issue to be reproduced on a standard (uncustomized) system.

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

The `guestfish` utility is a powerful and flexible tool for accessing and customizing a guest image. You can use it to install application binaries and data, create or resize file systems, and perform numerous other offline customizations.

The following example shows a `guestfish` script that starts by mounting the guest image and making it bootable. This initialization is a typical requirement for system images. Following the initialization phase, the script continues to perform a set of customizations to expand the root file system. You can use the following example as a template for similar customizations that you want to make.

> **Note**
>
> Consider the need to recalculate and resize `LVDoNotRemoveOrUse`. First-boot validation can fail if `LVDoNotRemoveOrUse` is missing or undersized.

```
# cat modify.fish
# Add target disk/volume to guestfs environment
add ./custom.img

# Start execution
run

# Mount root filesystem (LVM root) as /
mount /dev/VGExaDbDomU/LVDbSys1 /

# Mount /boot partition
mount /dev/sda1 /boot

# Install GRUB bootloader to disk (BIOS mode)
# NOTE: Not applicable for UEFI systems
sh "grub2-install /dev/sda"

# NOTE: Alternative for UEFI systems
# Mount EFI System Partition (required for UEFI boot files)
# mount /dev/sda2 /boot/efi

# End of Initialization

# Expand root file system lv to 40G. We need to expand both active/
inactive partitions
lvresize /dev/VGExaDbDomU/LVDbSys1 40960
lvresize /dev/VGExaDbDomU/LVDbSys2 40960

# Resize root filesystem.
xfs_growfs /

# LVDoNotRemoveOrUse is the snapshot reserve LV for KVM guest custom
filesystem layouts.
# The required reserve is calculated from the mounted system filesystems:
# /, /var, /opt, /etc, /root, /usr
# For each of those mounts that exists and is backed by an LV in the root
VG:
# reserve_for_mount = max(ceil(LV_size_in_GiB * 0.04), 1 GiB)
# The final required size is the sum of the reserves for each mount.
```


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

```
Rounding is done for each LV before summing (not once at the end).
#
# Examples:
# 40G LV -> ceil(4% of 40G) = ceil(1.6G) = 2GiB
# 10G LV -> ceil(0.4G) = 1GiB minimum
# 100G LV -> 4GiB
#
# If / is 40G, /var is 10G, and /opt is 20G,
# then LVDoNotRemoveOrUse required size = 2GiB + 1GiB + 1GiB = 4GiB
#
# In this case, / is 40G, /var is 2G, and no /opt.
# So, LVDoNotRemoveOrUse required size = 2GiB + 1GiB = 3GiB

lvresize /dev/VGExaDbDomU/LVDoNotRemoveOrUse 3072

# Exit guestfs session
exit
```

> **Note**
>
> For further `guestfish` commands and examples see the `libguestfs` help at https://libguestfs.org/.

After you prepare the customization script, you can run `guestfish` to apply it. For example, the following `guestfish` command runs the script at `./modify.fish`:

```
# guestfish --rw -f ./modify.fish
```

Examine any output from the `guestfish` command to ensure that the customizations are performed as planned.

6. Create a checksum for the customized image and copy the customized image back to the Exadata staging server.

   You can use the Linux `sha256sum` utility to create a checksum for a customized image. For example:

   ```
   # sha256sum ./custom.img
   b796a1f4b38d442f8ce1ab405525c38b92bddf1f37031580b46b7cd4af348bac ./
   custom.img
   ```

   The checksum is used later to validate the integrity of the customized image after it is copied to the Exadata staging server and loaded into the Exascale volume used for deployment.

7. On the Exadata staging server, provision a volume for the customized image.

   To provision a volume for a customized image, use the Exascale command line interface (ESCLI) to create a new empty volume and then attach it to the staging server. For example:

   ```
   [root@exa01adm05 ~]# escli --wallet <wallet-location> --ctrl <ERS-endpoint>
   Wallet: <wallet-location>
   User: admin
   ```


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

```
@> lsinitiator
id                                     hostName      giClusterName
giClusterId
d8fd0073-59ac-e40b-d8fd-007359ace40b exa01adm06
f7c2bd74-08c7-b02c-f7c2-bd7408c7b02c exa01adm05

@> mkvolume 150g --attributes vault=sysvault_exa01,name=goldvol_exa01
Created volume with id vol0021_99b9510268b04b74a8157d662ade5e4f

@> mkvolumeattachment  vol0021_99b9510268b04b74a8157d662ade5e4f
goldvol_exa01 --attributes initiator=f7c2bd74-08c7-b02c-f7c2-bd7408c7b02c
Created edv attachment with id att0021_a72a1ab7cec84c1dbdaf5a700354410a

@> exit
```

In the example:

- The volume size is 150 GB (`150g`), matching the size of the customized image it will accommodate.
- The volume is named `goldvol_exa01`. You can choose your own name for each customized image. Also take note of the volume ID (for example, `vol0021_99b9510268b04b74a8157d662ade5e4f`) as you will require it later to reference the volume.
- The volume is created in an Exascale vault named `sysvault_exa01`. You can choose an existing vault in your Exascale system or create a specific vault to house your customized images.

After volume provisioning, you can verify a volume by examining the corresponding block device at `/dev/exc/<volume-name>`. For example:

```
[root@exa01adm05 ~]# ls -la /dev/exc/goldvol_exa01
brw-rw----. 2 root disk 251, 11 May  7 07:40 /dev/exc/goldvol_exa01
```

8. Copy the customized image into the Exascale volume.

   For example:

   ```
   [root@exa01adm05 ~]# dd if=./custom.img of=/dev/exc/goldvol_exa01 bs=16M
   status=progress oflag=direct
   ...
   ```

9. Catalog the gold-image volume.

   As an operational best practice, you should maintain a catalog of your customized gold-images. For each image, you should maintain metadata to help you identify the image and its contents. You should also keep checksum information to validate the integrity of each image.

   To ensure an enduring read-only copy of the customized image, you can create a snapshot of the underlying Exascale volume. To identify the Exascale volume, use the volume ID reported during volume creation. For example:

   ```
   [root@exa01adm05 ~]# escli --wallet <wallet-location> --ctrl <ERS-endpoint>

   @> mkvolumesnapshot vol0021_99b9510268b04b74a8157d662ade5e4f --attributes
   name=goldvol_exa01_snap1
   ```


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

```
Created volume snapshot with id
vol0021_snap0001_50e52177583f4be4bad68ac20b65001e
```

Take note of the volume snapshot ID (for example, `vol0021_snap0001_50e52177583f4be4bad68ac20b65001e`) as you will require it to reference the volume snapshot.

After you create a volume snapshot, you can continue to use the original volume for further customizations. At any point you can make another snapshot for a different generation of the customized image.

After you have a snapshot for a validated and tested image, you can lock the snapshot to prevent accidental deletion. For example:

```
@> chvolumesnapshot vol0021_snap0001_50e52177583f4be4bad68ac20b65001e --
attributes lockExpiryTime=2050-06-30T13:00:00,lockImmutable=true
```

To use a validated and tested image, make a volume clone and associated attachment. For example:

```
@> mkvolume --attributes
volumeSnapshot=vol0021_snap0001_50e52177583f4be4bad68ac20b65001e,name=goldv
ol_exa01_clone1
Created volume with id vol0039_e9867af0dc03439698e5f6d7edf43a04

@> mkvolumeattachment vol0039_e9867af0dc03439698e5f6d7edf43a04
goldvol_exa01_clone1 --attributes initiator=f7c2bd74-08c7-b02c-f7c2-
bd7408c7b02c
Created edv attachment with id att0031_690f91c85ad64573a93485426467977b
```

The volume clone behaves just like the original Exascale volume, with the added convenience that it is easy to drop and re-create it to return the clone back to the snapshot state.

You can use the Linux `sha256sum` utility to create a checksum for a customized image. For example:

```
[root@exa01adm05 ~]# sha256sum /dev/exc/goldvol_exa01_clone1
b796a1f4b38d442f8ce1ab405525c38b92bddf1f37031580b46b7cd4af348bac /dev/exc/
goldvol_exa01_clone1
```

Compare this checksum with the one created on the `libguestfs` server to validate that the image has not changed. You can also save the checksum in the image catalog and check it every time you use the image.

10. Clone the customized image to the required KVM hosts.

    Before you can reference a customized image in an Exadata guest VM, you must clone the image to make it available on the required KVM hosts.

    You can use the Oracle Exadata Deployment Assistant (OEDA) Command Line Utility (OEDACLI) to perform the required actions.

    Following is an example of the required command sequence. In the example:

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

- `exa01adm05` and `exa01adm06` identify the KVM hosts that will expose the customized image. The command sequence assumes that the customized image is being cloned to two hosts. However, you can extend this to any required number of hosts.
- `goldvol_exa01adm05` is the volume name used to identify the volume on `exa01adm05`, and `goldvol_exa01adm06` is used to identify the volume on `exa01adm06`.
- `vol0039_e9867af0dc03439698e5f6d7edf43a04` is the Exascale volume ID associated with the customized image.
- `/root/es.xml` represents the path to the Exadata XML configuration file for your system.

```
[root@exa01adm05 linux-x64]# cat ./clone.cmd
CLONE EDVVOLUME volumename=goldvol_exa01adm05 where
volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04 hostname=exa01adm05
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

CLONE EDVVOLUME volumename=goldvol_exa01adm06 where
volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04 hostname=exa01adm06
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

SAVE FILE

[root@exa01adm05 linux-x64]# ./oedacli -c /root/es.xml -f ./clone.cmd
Version : 260504
SUCCESS - file loaded OK
Customer : Example - Example
processMerge
processMergeActions
Merging Action : CLONE EDVVOLUME volumename=goldvol_exa01adm05 where
volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04 hostname=exa01adm05
Merging CLONE EDVVOLUME
Action Validated and Merged OK
Deploying Action ID : 10 CLONE EDVVOLUME volumename=goldvol_exa01adm05
where volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04
hostname=exa01adm05
Deploying CLONE EDVVOLUME
Found the source EDV volume goldvol_exa01_clone1 for volume source
vol0039_e9867af0dc03439698e5f6d7edf43a04
Cloned an EDV volume with Exascale Volume ID
vol0041_e5f6d7edf43a04e9867af0dc03439698
Attached volume vol0041_e5f6d7edf43a04e9867af0dc03439698 into KVM host
exa01adm05.example.com for guest exa01adm05.example.com in /dev/exc/
goldvol_exa01adm05_vol0041_e5f6d7edf43a04e9867af0dc03439698
Done...
Done [Elapsed = 17235 mS [0.0 minutes] Thu May 07 10:00:30 PDT 2026]]
processMerge
processMergeActions
Merging Action : CLONE EDVVOLUME volumename=goldvol_exa01adm06 where
volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04 hostname=exa01adm06
Merging CLONE EDVVOLUME
Action Validated and Merged OK
```


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

```
Deploying Action ID : 11 CLONE EDVVOLUME volumename=goldvol_exa01adm06
where volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04
hostname=exa01adm06
Deploying CLONE EDVVOLUME
Found the source EDV volume goldvol_exa01_clone1 for volume source
vol0039_e9867af0dc03439698e5f6d7edf43a04
Cloned an EDV volume with Exascale Volume ID
vol0042_690f91c85ad64573a93485426467977b
Attached volume vol0042_690f91c85ad64573a93485426467977b into KVM host
exa01adm06.example.com for guest exa01adm06.example.com in /dev/exc/
goldvol_exa01adm06_vol0042_690f91c85ad64573a93485426467977b
Done...
Done [Elapsed = 10085 mS [0.0 minutes] Thu May 07 10:00:40 PDT 2026]]
File : /root/es.xml saved OK
[root@exa01adm05 linux-x64]#
```

Examine the OEDACLI output to determine the corresponding block device for the image volume on each KVM host. For example, `/dev/exc/goldvol_exa01adm05_vol0041_e5f6d7edf43a04e9867af0dc03439698` on `exa01adm05`, and `/dev/exc/goldvol_exa01adm06_vol0042_690f91c85ad64573a93485426467977b` on `exa01adm06`. Take note of these block devices as you will require them to reference the customized image in your guest VMs.

11. Alter the configuration of the Exadata guest VMs to use the customized image.

    To use a customized image in Exadata guest VMs, you must first define the VMs using OEDA as you would for any Exadata VM cluster. Then, before deploying the cluster, you must alter the VM definition in the Exadata XML configuration file using OEDACLI.

    Following is an example of the required command sequence. In the example:

    - `exa01adm05vm01` and `exa01adm06vm01` identify the guest VMs being altered.

      It is assumed that `exa01adm05vm01` resides on the KVM host `exa01adm05`, and `exa01adm06vm01` resides on `exa01adm06`.

      The command sequence assumes that the customized image is being used across two guests. However, you can extend this to any required number of guests.

    - The `imagefile` settings reference the block devices for the cloned image volume on each KVM host (shown in the previous step). Each guest VM must reference the cloned image volume that is attached to the KVM host where that guest VM will run.

      For `exa01adm05vm01`, the system image is set to `/dev/exc/goldvol_exa01adm05_vol0041_e5f6d7edf43a04e9867af0dc03439698`, which is the reference to the gold-image volume clone on the KVM host `exa01adm05`.

      For `exa01adm06vm01`, the system image is set to `/dev/exc/goldvol_exa01adm06_vol0042_690f91c85ad64573a93485426467977b`, which is the reference to the gold-image volume clone on `exa01adm06`.

    - `/root/cluster01-es.xml` represents the path to the Exadata XML configuration file containing the VMs being altered.

```
[root@exa01adm05 linux-x64]# cat ./update.cmd
ALTER MACHINE imagefile=/dev/exc/
goldvol_exa01adm05_vol0041_e5f6d7edf43a04e9867af0dc03439698 where
hostname=exa01adm05vm01
SAVE ACTION
MERGE ACTIONS
```


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

```
ALTER MACHINE imagefile=/dev/exc/
goldvol_exa01adm06_vol0042_690f91c85ad64573a93485426467977b where
hostname=exa01adm06vm01
SAVE ACTION
MERGE ACTIONS

SAVE FILE

[root@exa01adm05 linux-x64]# ./oedacli -c /root/cluster01-es.xml -f ./
update.cmd
Version : 260504
SUCCESS - file loaded OK
Customer : Example - Example
processMerge
processMergeActions
Merging Action : ALTER MACHINE imagefile=/dev/exc/
goldvol_exa01adm05_vol0041_e5f6d7edf43a04e9867af0dc03439698 where
hostname=exa01adm05vm01
Merging ALTER MACHINE
Action Validated and Merged OK
processMerge
processMergeActions
Merging Action : ALTER MACHINE imagefile=/dev/exc/
goldvol_exa01adm06_vol0042_690f91c85ad64573a93485426467977b where
hostname=exa01adm06vm01
Merging ALTER MACHINE
Action Validated and Merged OK
File : /root/cluster01-es.xml saved OK
[root@exa01adm05 linux-x64]#
```

12. Deploy the Exadata guest VMs using the customized image.

    Use the OEDA deployment script (`install.sh`) to deploy the customized guest VMs. For example:

```
[root@exa01adm05 linux-x64]# ./install.sh -cf /root/cluster01-es.xml -s 4
Initializing
===== Executing Create Virtual Machine =====
Checking Secure Fabric setup in [exa01celadm09.example.com,
exa01celadm07.example.com, exa01celadm08.example.com]
Checking status of remote files...
Checking status of existing files on remote nodes...
Getting status of local files...
Creating KVM hosts vaults: sysvault_exa01adm06,sysvault_exa01adm05
Vault sysvault_exa01adm06 was already created
Vault sysvault_exa01adm05 was already created
Creating EDV volumes and attachments...
Setting up exascale on KVM hosts if not already setup
Setting up exascale with lock on host: exa01adm05.example.com
Setting up exascale with lock on host: exa01adm06.example.com
Created 10 EDV volumes in 3163 mS
Found 10 volume attachments
Created EDV volume attachments in 3875 mS
Successfully created EDV volumes and attachments.
Creating guests for cluster Cluster-cluster01
```


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

```
Creating KVM guests in cluster Cluster-cluster01
File name is /EXAVMIMAGES/upgrade/linux-x64/WorkDir/
exa01adm05vm01.example.com-vm.xml
File name is /EXAVMIMAGES/upgrade/linux-x64/WorkDir/
exa01adm06vm01.example.com-vm.xml
Creating KVM guest using /EXAVMIMAGES/upgrade/linux-x64/WorkDir/
exa01adm05vm01.example.com-vm.xml on node exa01adm05.example.com
Creating KVM guest using /EXAVMIMAGES/upgrade/linux-x64/WorkDir/
exa01adm06vm01.example.com-vm.xml on node exa01adm06.example.com
KVM Guests created, waiting for guests to start...
No config Keys in the configuration file..
Creating databasemachine.xml for EM discovery
Done Creating databasemachine.xml for EM discovery
===== Successfully completed execution of step Create Virtual Machine
[elapsed Time [Elapsed = 314480 mS [5.0 minutes] Thu May 07 10:08:05 PDT
2026]] =====

[root@exa01adm05 linux-x64]#
```

Note that the deployment does not specifically identify when a customized image is used.

To confirm that the customized image was used, examine the newly created guest VMs and verify that they contain the expected customizations. For example, the following confirms that guest VM `exa01adm05vm01` has the customized 40GB root file system:

```
[root@exa01adm05 ~]# ssh -l root exa01adm05vm01
root@exa01adm05vm01's password:
...
[root@exa01adm05vm01 ~]# df -h
Filesystem
                       Size  Used Avail Use% Mounted on
devtmpfs
                        40G     0   40G   0% /dev
tmpfs
                        41G  5.2M   41G   1% /run
tmpfs
                        41G     0   41G   0% /sys/fs/cgroup
/dev/mapper/VGExaDbDomU-
LVDbSys1
 40G  3.2G   37G   8% /
/dev/mapper/VGExaDbDomU-
LVDbKdump
 20G  175M   20G   1% /crashfiles
/dev/
sda1
                  412M  170M  243M  42% /boot
/dev/mapper/VGExaDbDomU-
LVDbTmp
3.0G   54M  2.9G   2% /tmp
/dev/mapper/VGExaDbDomU-
LVDbVar1
2.0G  146M  1.8G   8% /var
/dev/mapper/VGExaDbDomU-
LVDbHome
4.0G   61M  3.9G   2% /home
/dev/mapper/VGExaDbDomU-
```


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

```
LVDbVarLog
 18G  191M   18G   2% /var/log
/dev/mapper/VGExaDbDomU-
LVDbVarLogAudit
924M   40M  885M   5% /var/log/audit
tmpfs
                        81G  4.0K   81G   1% /dev/shm
/dev/mapper/
VGExaDbDisk.exa01adm05vm01_u01_vol0027_fb9191ec508b4816899553e32973e053-
LVDBDisk        18G  161M   18G   1% /u01
/dev/mapper/
VGExaDbDisk.exa01adm05vm01_gih01_vol0030_70de90b21cde4ab5afefb5e49fe0d66a-
LVDBDisk      48G  3.1G   45G   7% /u01/app/23.0.0.0/grid
/dev/mapper/
VGExaDbDisk.exa01adm05vm01_dbh01_vol0033_4fa5dfdedf3c45c480060c508ddde9a4-
LVDBDisk      48G  8.6G   40G  18% /u01/app/oracle/product/23.0.0.0/
dbhome_4
tmpfs
                      8.1G     0  8.1G   0% /run/user/0
[root@exa01adm05vm01 ~]#
```

6.13.2 Customizing the Guest Image to Add a Storage Volume and File System

The following example workflow shows how to customize a guest image to add a storage volume with a file system. You can adapt the following to perform similar customizations in your environment.

1. Prepare a host with `libguestfs` packages and choose a staging server.

   Modification of the baseline Exadata system image is performed using the `libguestfs` tools (such as `guestfish`, `virt-resize`, and `virt-customize`). To avoid potential issues with Oracle Exadata System Software updates, Oracle recommends using a separate Linux host to perform image modification. You can use `dnf` to install the required `libguestfs` packages. For example:

   ```
   # dnf install libguestfs-xfs libguestfs-tools
   ```

   After image modification, the rest of the workflow is performed as `root` on an Exadata KVM host, also known as the staging server. Choose a suitable KVM host as the staging server. In this example workflow, the staging server is `exa01adm05`.

2. Select the Exadata base system image being customized and copy it to the `libguestfs` server.

   > **Note**
   >
   > The base system image must be from Oracle Exadata System Software release 26.1.0 or later.

   In this example workflow, the baseline Exadata system image is `System.first.boot.26.1.0.0.0.260515.rtg.img`.

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

3. On the `libguestfs` server, examine the baseline Exadata system image.

> **Note**
>
> On the `libguestfs` server, set `LIBGUESTFS_BACKEND=direct` in the Linux shell environment to avoid potential configuration and permission issues with `libguestfs` tools.

```
# ls -la ./System.first.boot.26.1.0.0.0.260515.rtg.img
-rw-r-----. 1 root root 106622563328 Apr 27 12:03
System.first.boot.26.1.0.0.0.260515.rtg.img
# export LIBGUESTFS_BACKEND=direct
# virt-filesystems -a ./System.first.boot.26.1.0.0.0.260515.rtg.img --all -
h -l
Name                                Type       VFS     Label MBR Size
Parent
/dev/sda1                           filesystem xfs     BOOT  -   412M -
/dev/sda2                           filesystem unknown -     -   256M -
/dev/VGExaDbDomU/LVDbHome           filesystem xfs     HOME  -   3.9G -
/dev/VGExaDbDomU/LVDbKdump          filesystem xfs     KDUMP -   20G  -
/dev/VGExaDbDomU/LVDbSwap1          filesystem swap    SWAP  -   16G  -
/dev/VGExaDbDomU/LVDbSys1           filesystem xfs     DBSYS -   15G  -
/dev/VGExaDbDomU/LVDbSys2           filesystem unknown -     -   15G  -
/dev/VGExaDbDomU/LVDbTmp            filesystem xfs     TMP   -   2.9G -
/dev/VGExaDbDomU/LVDbVar1           filesystem xfs     VAR   -   1.9G -
/dev/VGExaDbDomU/LVDbVar2           filesystem unknown -     -   2.0G -
/dev/VGExaDbDomU/LVDbVarLog         filesystem xfs     DIAG  -   18G  -
/dev/VGExaDbDomU/LVDbVarLogAudit    filesystem xfs     AUDIT -   924M -
/dev/VGExaDbDomU/LVDoNotRemoveOrUse filesystem unknown -     -   2.5G -
/dev/VGExaDbDomU/LVDbHome           lv         -       -     -   4.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbKdump          lv         -       -     -   20G  /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbSwap1          lv         -       -     -   16G  /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbSys1           lv         -       -     -   15G  /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbSys2           lv         -       -     -   15G  /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbTmp            lv         -       -     -   3.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbVar1           lv         -       -     -   2.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbVar2           lv         -       -     -   2.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbVarLog         lv         -       -     -   18G  /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDbVarLogAudit    lv         -       -     -   1.0G /dev/
VGExaDbDomU
/dev/VGExaDbDomU/LVDoNotRemoveOrUse lv         -       -     -   2.5G /dev/
VGExaDbDomU
/dev/VGExaDbDomU                    vg         -       -     -   99G  /dev/
sda3
```


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

```
/dev/sda3                              pv         -       -     -   99G  -
/dev/sda1                              partition  -       -     -
512M /dev/sda
/dev/sda2                              partition  -       -     -
256M /dev/sda
/dev/sda3                              partition  -       -     -
99G  /dev/sda
/dev/sda                               device     -       -     -   99G  -

# virt-cat -a ./System.first.boot.26.1.0.0.0.260515.rtg.img /etc/fstab
libguestfs: error: mount_ro: mount: /crashfiles: mount point is not a
directory
virt-cat: some filesystems could not be mounted (ignored)
LABEL=BOOT              /boot   xfs     defaults,nodev,nosuid   0 0
LABEL=DBSYS             /       xfs     defaults        0 0
LABEL=SWAP              swap    swap    defaults        0 0
LABEL=KDUMP             /crashfiles     xfs     defaults,nodev  0 0
LABEL=TMP               /tmp    xfs     defaults,nodev,nosuid   0 0
LABEL=HOME              /home   xfs     defaults,nodev,nosuid   0 0
LABEL=VAR               /var    xfs     defaults,nodev  0 0
LABEL=DIAG              /var/log        xfs     defaults,nodev,nosuid   0 0
LABEL=AUDIT             /var/log/audit  xfs
defaults,nodev,noexec,nosuid    0 0
tmpfs           /dev/shm        tmpfs   defaults,nodev,size=30831m 0 0

# qemu-img info ./System.first.boot.26.1.0.0.0.260515.rtg.img
image: System.first.boot.26.1.0.0.0.260515.rtg.img
file format: raw
virtual size: 99.3 GiB (106622563328 bytes)
disk size: 5.81 GiB
```

Take note of the size and internal structure of the baseline image.

4. Create a clone of the baseline system image and create an image for the new storage volume and file system:

```
# qemu-img create -f raw custom.img 100G
Formatting 'custom.img', fmt=raw size=107374182400

# virt-resize --expand /dev/sda3 ./
System.first.boot.26.1.0.0.0.260515.rtg.img  ./custom.img
[   0.0] Examining System.first.boot.26.1.0.0.0.260515.rtg.img
  25% ...
 100% ...
**********

Summary of changes:

/dev/sda1: This partition will be left alone.

/dev/sda2: This partition will be left alone.

/dev/sda3: This partition will be resized from 98.5G to 99.2G.  The LVM PV
on /dev/sda3 will be expanded using the ‘pvresize’ method.

**********
```


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

```
[   4.5] Setting up initial partition table on custom.img
[  15.7] Copying /dev/sda1
[  16.5] Copying /dev/sda2
[  16.8] Copying /dev/sda3
...
[ 137.1] Expanding /dev/sda3 using the ‘pvresize’ method

Resize operation completed with no errors.  Before deleting the old disk,
carefully check that the resized disk boots and works correctly

# qemu-img create -f raw extra.img 200G
Formatting 'extra.img', fmt=raw size=214748364800
```

In the example, the size of the cloned baseline image is 100 GB (`100G`) and the size of the image for the new storage volume and file system is 200 GB (`200G`).

5. Customize the base image and new storage volume.

   You can use `guestfish` to customize a guest image. It is recommended to prepare a script with the desired customizations and then run `guestfish` to apply the scripted customizations.

   > **Note**
   >
   > You are responsible for all image customizations that you perform. Oracle will not support questions or issues with specific image customizations. If a guest fails, and Oracle suspects the failure may be caused by an image customization, then Oracle may ask for the issue to be reproduced on a standard (uncustomized) system.

   The `guestfish` utility is a powerful and flexible tool for accessing and customizing a guest image. You can use it to install application binaries and data, create or resize file systems, and perform numerous other offline customizations.

   The following example shows a `guestfish` script that starts by mounting the guest image and making it bootable. This initialization is a typical requirement for system images. Following the initialization phase, the script continues to perform a set of customizations to add a storage volume and configure it as an XFS file system mounted on `/extra`. You can use the following example as a template for similar customizations that you want to make.

```
# cat modify.fish
# Add target disk/volume to guestfs environment
add ./custom.img
add ./extra.img

# Start execution
run

# Mount root filesystem (LVM root) as /
mount /dev/VGExaDbDomU/LVDbSys1 /

# Mount /boot partition
mount /dev/sda1 /boot

# Install GRUB bootloader to disk (BIOS mode)
```


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

```
# NOTE: Not applicable for UEFI systems
sh "grub2-install /dev/sda"

# NOTE: Alternative for UEFI systems
# Mount EFI System Partition (required for UEFI boot files)
# mount /dev/sda2 /boot/efi

# End of Initialization

# Create PV on extra volume
pvcreate /dev/sdb

# Create VG on extra volume
vgcreate vg_extra /dev/sdb

# Create LV
lvcreate-free lv_extra vg_extra 100

# Format LV with XFS and label
mkfs xfs /dev/vg_extra/lv_extra label:EXTRA

# Create mount directory
mkdir /extra

# Append fstab entry
write-append /etc/fstab "LABEL=EXTRA  /extra  xfs  defaults  0  2\n"

# Exit guestfs session
exit
```

> **Note**
>
> For further `guestfish` commands and examples see the `libguestfs` help at https://libguestfs.org/.

After you prepare the customization script, you can run `guestfish` to apply it. For example, the following `guestfish` command runs the script at `./modify.fish`:

```
# guestfish --rw -f ./modify.fish
```

Examine any output from the `guestfish` command to ensure that the customizations are performed as planned.

6. Create a checksum for each customized image and copy the customized images back to the Exadata staging server.

   You can use the Linux `sha256sum` utility to create a checksum for each customized image. For example:

```
# sha256sum ./custom.img
b796a1f4b38d442f8ce1ab405525c38b92bddf1f37031580b46b7cd4af348bac ./
custom.img
# sha256sum ./extra.img
```


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

```
cd4af348bacb796a1f4bf37031580b46b7cd4af348bacd442f8ce1ab405525c3 ./
extra.img
```

These checksums are used later to validate the integrity of the customized images after they are copied to the Exadata staging server and loaded into the Exascale volumes used for deployment.

7. On the Exadata staging server, provision volumes for the customized images.

   To provision volumes for customized images, use the Exascale command line interface (ESCLI) to create new empty volumes and then attach them to the staging server. For example:

```
[root@exa01adm05 ~]# escli --wallet <wallet-location> --ctrl <ERS-endpoint>
Wallet: <wallet-location>
User: admin
@> lsinitiator
id                                     hostName      giClusterName
giClusterId
d8fd0073-59ac-e40b-d8fd-007359ace40b exa01adm06
f7c2bd74-08c7-b02c-f7c2-bd7408c7b02c exa01adm05

@> mkvolume 100g --attributes vault=sysvault_exa01,name=goldvol_exa01
Created volume with id vol0021_0804568e719c4824bcdb146d4535c5f1

@> mkvolumeattachment vol0021_0804568e719c4824bcdb146d4535c5f1
goldvol_exa01 --attributes initiator=f7c2bd74-08c7-b02c-f7c2-bd7408c7b02c
Created edv attachment with id att0021_903462044500465c954358bca5e17510

@> mkvolume 200g --attributes vault=sysvault_exa01,name=extravol_exa01
Created volume with id vol0022_dd1a0d5d3fe94b97b95545d763e4324b

@> mkvolumeattachment vol0022_dd1a0d5d3fe94b97b95545d763e4324b
extravol_exa01 --attributes initiator=f7c2bd74-08c7-b02c-f7c2-bd7408c7b02c
Created edv attachment with id att0022_dc2974e3331745099a05fb8e939b8b80

@> exit
```

   In the example:

   - The first volume size is 100 GB (`100g`), matching the size of the customized system image it will accommodate. The second volume size is 200 GB (`200g`), matching the size of the new storage volume containing the `/extra` file system.
   - The volumes are named `goldvol_exa01` and `extravol_exa01`. You can choose your own name for each customized image. Also take note of the volume IDs (for example, `vol0021_0804568e719c4824bcdb146d4535c5f1` and `vol0022_dd1a0d5d3fe94b97b95545d763e4324b`) as you will require them later to reference the volumes.
   - The volumes are created in an Exascale vault named `sysvault_exa01`. You can choose an existing vault in your Exascale system or create a specific vault to house your customized images.

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

After volume provisioning, you can verify each volume by examining the corresponding block device at `/dev/exc/<volume-name>`. For example:

```
[root@exa01adm05 ~]# ls -la /dev/exc/*_exa01
brw-rw----. 2 root disk 251, 11 May  7 10:17 /dev/exc/goldvol_exa01
brw-rw----. 2 root disk 251, 12 May  7 10:18 /dev/exc/extravol_exa01
```

8. Copy the customized images into the prepared Exascale volumes.

   For example:

   ```
   [root@exa01adm05 ~]# dd if=./custom.img of=/dev/exc/goldvol_exa01 bs=16M
   status=progress oflag=direct
   ...

   [root@exa01adm05 ~]# dd if=./extra.img of=/dev/exc/extravol_exa01 bs=16M
   status=progress oflag=direct
   ...
   ```

9. Catalog the gold-image volumes.

   As an operational best practice, you should maintain a catalog of your customized gold-images. For each image, you should maintain metadata to help you identify the image and its contents. You should also keep checksum information to validate the integrity of each image.

   To ensure an enduring read-only copy of the customized image, you can create a snapshot of the underlying Exascale volume. To identify the Exascale volume, use the volume ID reported during volume creation. For example:

   ```
   [root@exa01adm05 ~]# escli --wallet <wallet-location> --ctrl <ERS-endpoint>

   @> mkvolumesnapshot vol0021_0804568e719c4824bcdb146d4535c5f1 --attributes
   name=goldvol_exa01_snap1
   Created volume snapshot with id
   vol0021_snap0001_50e52177583f4be4bad68ac20b65001e

   @> mkvolumesnapshot vol0022_dd1a0d5d3fe94b97b95545d763e4324b --attributes
   name=extravol_exa01_snap1
   Created volume snapshot with id
   vol0022_snap0001_5217ad68ac207583fb65001e50e4be4b
   ```

   Take note of the volume snapshot IDs (for example, `vol0021_snap0001_50e52177583f4be4bad68ac20b65001e` and `vol0022_snap0001_5217ad68ac207583fb65001e50e4be4b`) as you will require them to reference the volume snapshots.

   After you create a volume snapshot, you can continue to use the original volume for further customizations. At any point you can make another snapshot for a different generation of the customized image.

   After you have a snapshot for each validated and tested image, you can lock them to prevent accidental deletion. For example:

   ```
   @> chvolumesnapshot vol0021_snap0001_50e52177583f4be4bad68ac20b65001e --
   attributes lockExpiryTime=2050-06-30T13:00:00,lockImmutable=true
   ```


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

```
@> chvolumesnapshot vol0022_snap0001_5217ad68ac207583fb65001e50e4be4b --
attributes lockExpiryTime=2050-06-30T13:00:00,lockImmutable=true
```

To use a set of validated and tested images, make volume clones and associated attachments. For example:

```
@> mkvolume --attributes
volumeSnapshot=vol0021_snap0001_50e52177583f4be4bad68ac20b65001e,name=goldv
ol_exa01_clone1
Created volume with id vol0039_e9867af0dc03439698e5f6d7edf43a04

@> mkvolumeattachment vol0039_e9867af0dc03439698e5f6d7edf43a04
goldvol_exa01_clone1 --attributes initiator=f7c2bd74-08c7-b02c-f7c2-
bd7408c7b02c
Created edv attachment with id att0031_690f91c85ad64573a93485426467977b

@> mkvolume --attributes
volumeSnapshot=vol0022_snap0001_5217ad68ac207583fb65001e50e4be4b,name=extra
vol_exa01_clone1
Created volume with id vol0040_5f6d7ed43dc690398ef43a04e9867af0

@> mkvolumeattachment vol0040_5f6d7ed43dc690398ef43a04e9867af0
extravol_exa01_clone1 --attributes initiator=f7c2bd74-08c7-b02c-f7c2-
bd7408c7b02c
Created edv attachment with id att0032_64573a93c85a7b4d2646797854690f91
```

The volume clones behave just like the original Exascale volumes, with the added convenience that it is easy to drop and re-create them to return back to the snapshot state.

You can use the Linux `sha256sum` utility to create a checksum for each customized image. For example:

```
[root@exa01adm05 ~]# sha256sum /dev/exc/goldvol_exa01_clone1
b796a1f4b38d442f8ce1ab405525c38b92bddf1f37031580b46b7cd4af348bac /dev/exc/
goldvol_exa01_clone1
[root@exa01adm05 ~]# sha256sum /dev/exc/extravol_exa01_clone1
cd4af348bacb796a1f4bf37031580b46b7cd4af348bacd442f8ce1ab405525c3 /dev/exc/
extravol_exa01_clone1
```

Compare the checksums with the ones created on the `libguestfs` server to validate that the images have not changed. You can also save the checksums in the image catalog and check them every time you use the images.

10. Clone the customized images to the required KVM hosts.

    Before you can reference a customized image in an Exadata guest VM, you must clone the image to make it available on the required KVM hosts.

    You can use the Oracle Exadata Deployment Assistant (OEDA) Command Line Utility (OEDACLI) to perform the required actions.

    Following is an example of the required command sequence. In the example:

    - `exa01adm05` and `exa01adm06` identify the KVM hosts that will expose the customized images. The command sequence assumes that the customized images are being cloned to two hosts. However, you can extend this to any required number of hosts.

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

- `goldvol_exa01adm05` and `extravol_exa01adm05` are the volume names used to identify the volumes on `exa01adm05`. `goldvol_exa01adm06` and `extravol_exa01adm06` are used to identify the volumes on `exa01adm06`.
- `vol0039_e9867af0dc03439698e5f6d7edf43a04` and `vol0040_5f6d7ed43dc690398ef43a04e9867af0` are the Exascale volume IDs associated with the customized images.
- `/root/es.xml` represents the path to the Exadata XML configuration file for your system.

```
[root@exa01adm05 linux-x64]# cat ./clone.cmd
CLONE EDVVOLUME volumename=goldvol_exa01adm05 where
volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04 hostname=exa01adm05
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

CLONE EDVVOLUME volumename=extravol_exa01adm05 where
volumesource=vol0040_5f6d7ed43dc690398ef43a04e9867af0 hostname=exa01adm05
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

CLONE EDVVOLUME volumename=goldvol_exa01adm06 where
volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04 hostname=exa01adm06
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

CLONE EDVVOLUME volumename=extravol_exa01adm06 where
volumesource=vol0040_5f6d7ed43dc690398ef43a04e9867af0 hostname=exa01adm06
SAVE ACTION
MERGE ACTIONS
DEPLOY ACTIONS

SAVE FILE

[root@exa01adm05 linux-x64]# ./oedacli -c /root/es.xml -f ./clone.cmd
Version : 260504
SUCCESS - file loaded OK
Customer : Example - Example
processMerge
processMergeActions
Merging Action : CLONE EDVVOLUME volumename=goldvol_exa01adm05 where
volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04 hostname=exa01adm05
Merging CLONE EDVVOLUME
Action Validated and Merged OK
Deploying Action ID : 12 CLONE EDVVOLUME volumename=goldvol_exa01adm05
where volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04
hostname=exa01adm05
Deploying CLONE EDVVOLUME
Found the source EDV volume goldvol_exa01_clone1 for volume source
vol0039_e9867af0dc03439698e5f6d7edf43a04
Cloned an EDV volume with Exascale Volume ID
vol0053_4bb5b73865774dcd8da0bec3ee73c176
Attached volume vol0053_4bb5b73865774dcd8da0bec3ee73c176 into KVM host
```


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

```
exa01adm05.example.com for guest exa01adm05.example.com in /dev/exc/
goldvol_exa01adm05_vol0053_4bb5b73865774dcd8da0bec3ee73c176
Done...
Done [Elapsed = 16457 mS [0.0 minutes] Thu May 07 11:41:12 PDT 2026]]
processMerge
processMergeActions
Merging Action : CLONE EDVVOLUME volumename=extravol_exa01adm05 where
volumesource=vol0040_5f6d7ed43dc690398ef43a04e9867af0 hostname=exa01adm05
Merging CLONE EDVVOLUME
Action Validated and Merged OK
Deploying Action ID : 13 CLONE EDVVOLUME volumename=extravol_exa01adm05
where volumesource=vol0040_5f6d7ed43dc690398ef43a04e9867af0
hostname=exa01adm05
Deploying CLONE EDVVOLUME
Found the source EDV volume extravol_exa01_clone1 for volume source
vol0040_5f6d7ed43dc690398ef43a04e9867af0
Cloned an EDV volume with Exascale Volume ID
vol0054_b804ab74ccd7400ca5aef861f588be54
Attached volume vol0054_b804ab74ccd7400ca5aef861f588be54 into KVM host
exa01adm05.example.com for guest exa01adm05.example.com in /dev/exc/
extravol_exa01adm05_vol0054_b804ab74ccd7400ca5aef861f588be54
Done...
Done [Elapsed = 10285 mS [0.0 minutes] Thu May 07 11:41:22 PDT 2026]]
processMerge
processMergeActions
Merging Action : CLONE EDVVOLUME volumename=goldvol_exa01adm06 where
volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04 hostname=exa01adm06
Merging CLONE EDVVOLUME
Action Validated and Merged OK
Deploying Action ID : 14 CLONE EDVVOLUME volumename=goldvol_exa01adm06
where volumesource=vol0039_e9867af0dc03439698e5f6d7edf43a04
hostname=exa01adm06
Deploying CLONE EDVVOLUME
Found the source EDV volume goldvol_exa01_clone1 for volume source
vol0039_e9867af0dc03439698e5f6d7edf43a04
Cloned an EDV volume with Exascale Volume ID
vol0055_1f7d044d6a93435d9bf79c788a811e3e
Attached volume vol0055_1f7d044d6a93435d9bf79c788a811e3e into KVM host
exa01adm06.example.com for guest exa01adm06.example.com in /dev/exc/
goldvol_exa01adm06_vol0055_1f7d044d6a93435d9bf79c788a811e3e
Done...
Done [Elapsed = 10305 mS [0.0 minutes] Thu May 07 11:41:32 PDT 2026]]
processMerge
processMergeActions
Merging Action : CLONE EDVVOLUME volumename=extravol_exa01adm06 where
volumesource=vol0040_5f6d7ed43dc690398ef43a04e9867af0 hostname=exa01adm06
Merging CLONE EDVVOLUME
Action Validated and Merged OK
Deploying Action ID : 15 CLONE EDVVOLUME volumename=extravol_exa01adm06
where volumesource=vol0040_5f6d7ed43dc690398ef43a04e9867af0
hostname=exa01adm06
Deploying CLONE EDVVOLUME
Found the source EDV volume extravol_exa01_clone1 for volume source
vol0040_5f6d7ed43dc690398ef43a04e9867af0
Cloned an EDV volume with Exascale Volume ID
vol0056_11b7c5feb9c74ca6b1e1e9b45578ea93
```


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

```
Attached volume vol0056_11b7c5feb9c74ca6b1e1e9b45578ea93 into KVM host
exa01adm06.example.com for guest exa01adm06.example.com in /dev/exc/
extravol_exa01adm06_vol0056_11b7c5feb9c74ca6b1e1e9b45578ea93
Done...
Done [Elapsed = 10120 mS [0.0 minutes] Thu May 07 11:41:43 PDT 2026]]
File : /root/es.xml saved OK
[root@exa01adm05 linux-x64]#
```

Examine the OEDACLI output to determine the corresponding block device for the image volumes on each KVM host. For example, `/dev/exc/goldvol_exa01adm05_vol0053_4bb5b73865774dcd8da0bec3ee73c176` and `/dev/exc/extravol_exa01adm05_vol0054_b804ab74ccd7400ca5aef861f588be54` on `exa01adm05`, plus `/dev/exc/goldvol_exa01adm06_vol0055_1f7d044d6a93435d9bf79c788a811e3e` and `/dev/exc/extravol_exa01adm06_vol0056_11b7c5feb9c74ca6b1e1e9b45578ea93` on `exa01adm06`. Take note of these block devices as you will require them to reference the customized images in your guest VMs.

11. Alter the configuration of the Exadata guest VMs to use the customized images.

    To use a customized image in Exadata guest VMs, you must first define the VMs using OEDA as you would for any Exadata VM cluster. Then, before deploying the cluster, you must alter the VM definition in the Exadata XML configuration file using OEDACLI.

    Following is an example of the required command sequence. In the example:

    - `exa01adm05vm01` and `exa01adm06vm01` identify the guest VMs being altered.

      It is assumed that `exa01adm05vm01` resides on the KVM host `exa01adm05`, and `exa01adm06vm01` resides on `exa01adm06`.

      The command sequence assumes that the customized image is being used across two guests. However, you can extend this to any required number of guests.

    - The `imagefile` settings reference the block devices for the cloned image volumes on each KVM host (shown in the previous step).

      `exa01adm05vm01` references volumes on the KVM host `exa01adm05`, where `/dev/exc/goldvol_exa01adm05_vol0053_4bb5b73865774dcd8da0bec3ee73c176` is the system gold-image and `/dev/exc/extravol_exa01adm05_vol0054_b804ab74ccd7400ca5aef861f588be54` is the extra volume image.

      `exa01adm06vm01` references volumes on the KVM host `exa01adm06`, where `/dev/exc/goldvol_exa01adm06_vol0055_1f7d044d6a93435d9bf79c788a811e3e` is the system gold-image and `/dev/exc/extravol_exa01adm06_vol0056_11b7c5feb9c74ca6b1e1e9b45578ea93` is the extra volume image.

    - The `imagefileoption=ADDEXTRAEXCVOL` setting adds an extra (non-standard) volume to the guest VM.

    - `/root/cluster01-es.xml` represents the path to the Exadata XML configuration file containing the VMs being altered.

```
[root@exa01adm05 linux-x64]# cat ./update.cmd
ALTER MACHINE imagefile=/dev/exc/
goldvol_exa01adm05_vol0053_4bb5b73865774dcd8da0bec3ee73c176 where
hostname=exa01adm05vm01
SAVE ACTION
MERGE ACTIONS
```


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

```
ALTER MACHINE imagefileoption=ADDEXTRAEXCVOL imagefile=/dev/exc/
extravol_exa01adm05_vol0054_b804ab74ccd7400ca5aef861f588be54 where
hostname=exa01adm05vm01
SAVE ACTION
MERGE ACTIONS

ALTER MACHINE imagefile=/dev/exc/
goldvol_exa01adm06_vol0055_1f7d044d6a93435d9bf79c788a811e3e where
hostname=exa01adm06vm01
SAVE ACTION
MERGE ACTIONS

ALTER MACHINE imagefileoption=ADDEXTRAEXCVOL imagefile=/dev/exc/
extravol_exa01adm06_vol0056_11b7c5feb9c74ca6b1e1e9b45578ea93 where
hostname=exa01adm06vm01
SAVE ACTION
MERGE ACTIONS

SAVE FILE

[root@exa01adm05 linux-x64]# ./oedacli -c /root/cluster01-es.xml -f ./
update.cmd
Version : 260504
SUCCESS - file loaded OK
Customer : Example - Example
processMerge
processMergeActions
Merging Action : ALTER MACHINE imagefile=/dev/exc/
goldvol_exa01adm05_vol0053_4bb5b73865774dcd8da0bec3ee73c176 where
hostname=exa01adm05vm01
Merging ALTER MACHINE
Action Validated and Merged OK
processMerge
processMergeActions
Merging Action : ALTER MACHINE imagefileoption=ADDEXTRAEXCVOL
imagefile=/dev/exc/
extravol_exa01adm05_vol0054_b804ab74ccd7400ca5aef861f588be54 where
hostname=exa01adm05vm01
Merging ALTER MACHINE
Action Validated and Merged OK
processMerge
processMergeActions
Merging Action : ALTER MACHINE imagefile=/dev/exc/
goldvol_exa01adm06_vol0055_1f7d044d6a93435d9bf79c788a811e3e where
hostname=exa01adm06vm01
Merging ALTER MACHINE
Action Validated and Merged OK
processMerge
processMergeActions
Merging Action : ALTER MACHINE imagefileoption=ADDEXTRAEXCVOL
imagefile=/dev/exc/
extravol_exa01adm06_vol0056_11b7c5feb9c74ca6b1e1e9b45578ea93 where
hostname=exa01adm06vm01
Merging ALTER MACHINE
Action Validated and Merged OK
```


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

```
File : /root/cluster01-es.xml saved OK
[root@exa01adm05 linux-x64]#
```

12. Deploy the Exadata guest VMs using the customized image.

    Use the OEDA deployment script (`install.sh`) to deploy the customized guest VMs. For example:

```
[root@exa01adm05 linux-x64]# ./install.sh -cf /root/cluster01-es.xml -s 4
Initializing
===== Executing Create Virtual Machine =====
Checking Secure Fabric setup in [exa01celadm09.example.com,
exa01celadm07.example.com, exa01celadm08.example.com]
Checking status of remote files...
Checking status of existing files on remote nodes...
Getting status of local files...
Creating KVM hosts vaults: sysvault_exa01adm06,sysvault_exa01adm05
Vault sysvault_exa01adm06 was already created
Vault sysvault_exa01adm05 was already created
Creating EDV volumes and attachments...
Setting up exascale on KVM hosts if not already setup
Setting up exascale with lock on host: exa01adm05.example.com
Setting up exascale with lock on host: exa01adm06.example.com
Created 10 EDV volumes in 3163 mS
Found 10 volume attachments
Created EDV volume attachments in 3875 mS
Successfully created EDV volumes and attachments.
Creating guests for cluster Cluster-cluster01
Creating KVM guests in cluster Cluster-cluster01
File name is /EXAVMIMAGES/upgrade/linux-x64/WorkDir/
exa01adm05vm01.example.com-vm.xml
File name is /EXAVMIMAGES/upgrade/linux-x64/WorkDir/
exa01adm06vm01.example.com-vm.xml
Creating KVM guest using /EXAVMIMAGES/upgrade/linux-x64/WorkDir/
exa01adm05vm01.example.com-vm.xml on node exa01adm05.example.com
Creating KVM guest using /EXAVMIMAGES/upgrade/linux-x64/WorkDir/
exa01adm06vm01.example.com-vm.xml on node exa01adm06.example.com
KVM Guests created, waiting for guests to start...
No config Keys in the configuration file..
Creating databasemachine.xml for EM discovery
Done Creating databasemachine.xml for EM discovery
===== Successfully completed execution of step Create Virtual Machine
[elapsed Time [Elapsed = 314491 mS [5.0 minutes] Thu May 07 11:49:36 PDT
2026]] =====

[root@exa01adm05 linux-x64]#
```

    Note that the deployment does not specifically identify when customized images are used.

    To confirm that the customized images were used, examine the newly created guest VMs and verify that they contain the expected customizations. For example, the following confirms that guest VM `exa01adm05vm01` has the new volume with the `/extra` file system:

```
[root@exa01adm05 ~]# ssh -l root exa01adm05vm01
root@exa01adm05vm01's password:
...
```


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

```
[root@exa01adm05vm01 ~]# df -h
Filesystem
                        Size  Used Avail Use% Mounted on
devtmpfs
                         40G     0   40G   0% /dev
tmpfs
                         41G  5.2M   41G   1% /run
tmpfs
                         41G     0   41G   0% /sys/fs/cgroup
/dev/mapper/VGExaDbDomU-
LVDbSys1
 15G  3.0G   12G  21% /
/dev/mapper/VGExaDbDomU-
LVDbTmp
3.0G   54M  2.9G   2% /tmp
/dev/
sda1
                  412M  170M  243M  42% /boot
/dev/mapper/VGExaDbDomU-
LVDbKdump
 20G  175M   20G   1% /crashfiles
/dev/mapper/vg_extra-
lv_extra
   200G  1.5G  199G   1% /extra
/dev/mapper/VGExaDbDomU-
LVDbVar1
2.0G  146M  1.8G   8% /var
/dev/mapper/VGExaDbDomU-
LVDbHome
4.0G   61M  3.9G   2% /home
/dev/mapper/VGExaDbDomU-
LVDbVarLog
 18G  191M   18G   2% /var/log
/dev/mapper/VGExaDbDomU-
LVDbVarLogAudit
924M   40M  885M   5% /var/log/audit
tmpfs
                         81G  4.0K   81G   1% /dev/shm
/dev/mapper/
VGExaDbDisk.scaqan08adm05vm04_u01_vol0042_f30f22cbc7c94b278b6426d726d9ce3c-
LVDBDisk     18G  161M   18G   1% /u01
/dev/mapper/
VGExaDbDisk.scaqan08adm05vm04_gih01_vol0031_cc03aa825ac64dfcb52f67a95c7d170
5-LVDBDisk   48G  3.1G   45G   7% /u01/app/23.0.0.0/grid
/dev/mapper/
VGExaDbDisk.scaqan08adm05vm04_dbh01_vol0041_e5f46b7535104263b1126520268d26f
7-LVDBDisk   48G  8.6G   40G  18% /u01/app/oracle/product/23.0.0.0/dbhome_4
tmpfs
                        8.1G     0  8.1G   0% /run/user/0
[root@exa01adm05vm01 ~]#
```

6.14 Imaging a New System

You have different options for preparing an Oracle Exadata system for deployment.

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

Using a `preconf.csv` file with the PXE 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 Elastic Configuration to Image a New System
  Elastic configuration is now the standard methodology for all new deployments, and applies to Oracle Exadata 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.
- Updates to Imaging Parameters
  Some of the command line options for imaging have changed with 12.1.2.1.0 and above.

6.14.1 Using Elastic Configuration to Image a New System

Elastic configuration is now the standard methodology for all new deployments, and applies to Oracle Exadata 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 `node`*`N`*, 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.

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

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 systen 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. On systems with InfiniBand Network Fabric only, run `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim` on each database server before installing any software.

   The command reclaims disk space reserved for the deployment type not selected, which is not required on RoCE-based Exadata database servers.

   The command typically takes less than 5 minutes.

   > **⚠ 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
```


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

```
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
```

6.14.2 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 `node`*N*, 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.

   Use one of the following connection options:

   - Attach a serial cable between a console and the SER MGT port on the server.

     Connections through the modular jack serial port must use shielded cables to comply with radio frequency emission limits.

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

Set the terminal device communication properties to these values: 9600 baud, 8 bit, no parity, 1 stop bit.

For more details, see Connect to Oracle ILOM.

- Attach an Ethernet cable to the NET MGT port on the server and then connect using a workstation on the network.

  For more details, see Connect to Oracle ILOM.

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. Prepare and locate the files required for system imaging.

   Use Exadata Database Machine and Exadata Storage Server Supported Versions (My Oracle Support Doc ID KB153930) to locate the relevant patch archive. Then, unzip the patch archive and follow the steps in the patch `README.txt` to copy the ISO file and optional `preconf.csv` file into place.

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

Note

- Using ISO with the `preconf.csv` file 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 (SAMBA 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*
- If a newly imaged system does not have access to DHCP, then the imaging procedure will not be able to access the `preconf.csv` file. In this case, the imaging procedure will attempt to locate the `preconf.csv` inside a dedicated ISO image located in the same directory and named `preconf.iso`.

  To use this method, you can create the `preconf.iso` file by using the following command:

  ```
  # mkisofs -l -input-charset ascii -o preconf.iso preconf.csv
  ```
- When a `preconf.csv` file is not being used, the ISO image can be made available as virtual media on the ILOM of the node being re-imaged.
- 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.

10. Reboot the database node or cell to start the imaging process from the ISO.
11. 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 systen 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
    ```


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

12. On systems with InfiniBand Network Fabric only, run `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim` on each database server before installing any software.

    The command reclaims disk space reserved for the deployment type not selected, which is not required on RoCE-based Exadata database servers.

    The command typically takes less than 5 minutes.

    > **⚠ 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
```

13. On systems using InfiniBand Network Fabric, 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.

14. 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
```


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

```
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 Database Machine and Exadata Storage Server Supported Versions (My Oracle Support Doc ID KB153930)

6.14.3 Using PXE to Image a New System

You can use PXE to prepare a new Exadata system 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 `node`*N*, 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.

   Use one of the following connection options:

   - Attach a serial cable between a console and the SER MGT port on the server.

     Connections through the modular jack serial port must use shielded cables to comply with radio frequency emission limits.

     Set the terminal device communication properties to these values: 9600 baud, 8 bit, no parity, 1 stop bit.

     For more details, see Connect to Oracle ILOM.
   - Attach an Ethernet cable to the NET MGT port on the server and then connect using a workstation on the network.

     For more details, see Connect to Oracle ILOM.
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).

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

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. Following are the recommended options for PXE imaging. 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
```

   On systems with Oracle Exadata System Software release 20.1.0 or later, you can enable or disable Exadata Secure RDMA Fabric Isolation by adding `qinq=yes` or `qinq=no` to the kernel options. For example:

```
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 qinq=no
```

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
```


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

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. 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
```

11. 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
ys1
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 systen 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
```

12. On systems with InfiniBand Network Fabric only, run `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim` on each database server before installing any software.

The command reclaims disk space reserved for the deployment type not selected, which is not required on RoCE-based Exadata database servers.

The command typically takes less than 5 minutes.

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

> **⚠ 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
```

13. 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.

14. 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
```


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

```
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 KB153930)

6.14.4 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:

- `updfrm`

  Do NOT use the `updfrm` 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.

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

6.15 Performing Initial Elastic Configuration of Oracle Exadata

This topic provides background information on elastic configuration and describes how to perform the initial configuration of Oracle Exadata.

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 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 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 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**
>
> - Disk space for unused system partitions must be reclaimed before installing software on all database servers with InfiniBand Network Fabric. Use the following command to reclaim the disk space: `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim`. This command is not required on RoCE-based Exadata database servers.
> - 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.

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.

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

> **Note**
>
> The network connection and configuration for Oracle Exadata 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. On systems with InfiniBand Network Fabric only, 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. If using virtualization on the rack, configure each database server as follows:

> **Note**
>
> 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 Server (OVM) is a Xen-based virtualization technology, which is used across Oracle Exadata Database Machine systems that use InfiniBand Network Fabric.
>
> This step applies to both virtualization technologies.

```
# /opt/oracle.SupportTools/switch_to_ovm.sh
```

The database servers will reboot during this step.

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

> **Note**
>
> It is not possible or supported to switch between physical and virtualized configurations.

5. On systems with InfiniBand Network Fabric only, 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 with InfiniBand Network Fabric. Use the following command to reclaim the disk space: `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim`. This command is not required on RoCE-based Exadata database servers.

   `reclaimdisks.sh` performs partition reconfiguration and a few other necessary configuration changes that take typically less than a minute to complete.

6. On systems with InfiniBand Network Fabric only, verify the disk configuration on the database servers.

   ```
   # /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.
   ```

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 -1 /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
        ```


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

> **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 `reclaimdidsks.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
```


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

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 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. 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.

9. 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.

10. Connect the corporate network cable for the management network to the Management Network Switch.

11. Connect the corporate network cables for the client access network to the database servers.

12. Restart all database servers and storage servers.

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

13. 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 to verify network connectivity to Oracle Exadata. 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.

14. Use OEDA to install and configure the software for Oracle Exadata.

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.

6.16 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.

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.

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

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. For a database server with InfiniBand Network Fabric only, log in to the server and reclaim the unused disk space.

   Disk space for unused system partitions must be reclaimed before installing software on all database servers with InfiniBand Network Fabric. Use the following command to reclaim the disk space: `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim`. This command is not required on RoCE-based Exadata database servers.

   `reclaimdisks.sh` performs partition reconfiguration and a few other necessary configuration changes that take typically less than a minute to complete.

8. 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.

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

6.17 Using the OEDA genPasswordHash Utility

You can configure OEDA to use customized OS user passwords by using the `genPasswordHash.sh` utility.

During initial deployment or while altering an existing system deployment, Oracle Exadata Deployment Assistant (OEDA) performs actions on Exadata system servers using various OS user accounts.

To do this, OEDA can use default user passwords or you must provide the user passwords to OEDA by using the `genPasswordHash.sh` command located in the OEDA software bundle.

When you supply a user password to OEDA using `genPasswordHash.sh`, the command generates a secure password hash enabling OEDA to later use the password as required. The `genPasswordHash.sh` command does not change the OS user password on the Exadata servers. It is your responsibility to ensure that the passwords provided to `genPasswordHash.sh` are the actual OS user passwords on all the servers in your Exadata system.

To specify an OS user password, run the `genPasswordHash.sh` command with one of the following options, then supply the password when prompted. The available options are:

- `-root` - Sets the `root` OS user password used by OEDA.
- `-non-root` - Sets the OS user password used by OEDA for the `grid` and `oracle` OS users.
- `-all` - Sets the OS user password used by OEDA for all OS users (`root`, `grid`, and `oracle`).
- `-celladmin` - Sets the `celladmin` OS user password used by OEDA on the Exadata storage servers.

> **Note**
>
> The `-celladmin` option is available only with Oracle Exadata System Software release 26.1.0 or later.

For example, to set the `root` password used by OEDA:

```
# /u01/oeda/linux-x86/genPasswordHash.sh -root
Enter root password::******
Re-enter root password:*****
```

Related Topics

- OEDA Deployment - Generating new or non-default hash password for root and non-root users (My Oracle Support Doc ID KB701699)

6.18 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.

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

- 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.

6.19 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.

- 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 OEDA to perform a deployment on an Oracle Exadata Rack that is already configured with a non-default `root` password.
- Deploying Oracle Exadata Using Enhanced Storage Server Security
  Use enhanced storage server security to adopt the most secure posture during initial system deployment.

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

> **See Also**
>
> Using Oracle Exadata Deployment Assistant for descriptions of the configuration files generated by OEDA.

6.19.1 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 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.

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

2. Disk space for unused system partitions must be reclaimed before installing software on all database servers with InfiniBand Network Fabric. Use the following command to reclaim the disk space: `/opt/oracle.SupportTools/reclaimdisks.sh -free -reclaim`. This command is not required on RoCE-based Exadata database servers.
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. If required, download and install a supported Java Development Kit (JDK).

   If you are using a version of OEDA released in conjunction with Oracle Exadata System Software release 26.1.0 or later, ensure that any system running OEDA programs contains Java Development Kit (JDK) version 17 or later. You can download a supported JDK from https://www.oracle.com/java/technologies/downloads/.

   Before running any OEDA program, check the following to ensure that the required JDK is installed and available:

   - Ensure that the `JAVA_HOME` environment variable references the intended JDK.
   - Ensure that `JAVA_HOME/bin/java` is executable.
   - Run `java --version` and confirm that the output reports Java version 17 or higher.
6. 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.
7. 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

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

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.

8. Log in to the database server or Oracle VM and go to the OEDA home directory.
9. 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 main 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 document KB153930. 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
```


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

Related Topics

- Exadata Database Machine and Exadata Storage Server Supported Versions (My Oracle Support Doc ID KB153930)

6.19.2 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.

6.19.3 Using Oracle Exadata Deployment Assistant on Systems with Non-Default root Passwords

You can use OEDA to perform a deployment on an Oracle Exadata Rack that is already configured with a non-default `root` password.

Oracle Exadata Deployment Assistant (OEDA) enables deployment on Oracle Exadata Racks that already have non-default `root` passwords. Using the following options, you can perform the deployment without having to change the `root` password back to the default:

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

- If non-default passwords are already set, you can add the `-p` option to the existing command line options when you run the OEDA deployment script (`install.sh`). For example:

  ```
  ./install.sh -cf <config.xml> -s 1 -p
  ```

  The `-p` option instructs OEDA to prompt for the `root` password when required. Using this feature, you can perform a deployment on a system that has a different `root` password on each server in the system. Ensure 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.

- Alternatively, you can use the `genPasswordHash.sh` command to store a password hash, which OEDA can use to perform secure login to Exadata servers with non-default `root` passwords. In this case, all of the servers must use the same `root` password. See Using the OEDA genPasswordHash Utility.

Related Topics

- Default Security Settings Enacted by OEDA

6.19.4 Deploying Oracle Exadata Using Enhanced Storage Server Security

Use enhanced storage server security to adopt the most secure posture during initial system deployment.

Starting with Oracle Exadata System Software release 26.1.0, you can implement enhanced storage server security features during initial system deployment on Oracle Exadata system hardware with RoCE Network Fabric (X8M-2 and later systems).

Use the following sequence of steps:

1. Immediately after imaging the Exadata storage servers, change the passwords for the `celladmin` and `cellmonitor` OS user accounts.

   For each user, ensure you use the same password across all of the Exadata storage servers.

2. Optionally, disable root SSH access on the Exadata storage servers.

   Run the following CellCLI command on each storage server:

   ```
   CellCLI> ALTER CELL accessLevelPerm=remoteLoginEnabledRootLoginDisabled
   ```

3. On the server running the OEDA deployment script (`install.sh`), run the `genPasswordHash.sh` command and set the password for `celladmin`.

   Run the following command and specify the `celladmin` password when prompted:

   ```
   # ./genPasswordHash.sh -celladmin
   Enter celladmin password::******
   Re-enter celladmin password:*****
   ```

   See also Using the OEDA genPasswordHash Utility.

4. Alter the Exadata XML configuration file (`es.xml`) to instruct OEDA to perform system deployment operations on the storage servers using the `celladmin` OS user account instead of `root`.

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

Use the following OEDA Command Line Utility (OEDACLI) commands to alter the XML configuration file:

```
oedacli> LOAD FILE NAME=es.xml
oedacli> ALTER PROPERTY name=USE_CELLADMIN_USER value=true
oedacli> SAVE ACTION
oedacli> MERGE ACTIONS
oedacli> SAVE FILE
```

5. Perform the initial deployment using the OEDA deployment script (`install.sh`).

   During system deployment, OEDA performs operations on the storage servers using the `celladmin` OS user account. The deployment result is exactly the same using `celladmin` instead of `root`. However, disabling `root` SSH access and using `celladmin` guards against malicious activities being performed during deployment.

   See also Configuring Oracle Exadata Database Machine Using OEDA.

6. After initial system deployment, perform routine storage server administration tasks using the `celladmin` and `cellmonitor` OS user accounts. You can optionally configure SSH equivalence for `celladmin` and `cellmonitor` to enable passwordless SSH access.

   If root access is absolutely required, you can temporarily allow remote login by setting the cell `accessLevelTemp` attribute. For example, the following command enables remote login for 30 minutes at the specified time:

   ```
   CellCLI> ALTER CELL accessLevelTemp=((accessLevel="remoteLoginEnabled", -
            startTime="2026-06-20T01:01:00-07:00",                 -
            duration="30m"))
   ```

   Root access should otherwise remain disabled to maintain a strong security posture.

6.20 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.

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

6.21 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 KB297213)
- MAA Best Practices website

6.22 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 XML system configuration file.

To add a VM cluster to an existing Oracle Exadata Database Machine deployment:

1. Start the OEDA Web-based configuration tool.

   See Getting Started with the OEDA Browser-Based User Interface.
2. Import the existing XML system configuration 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.

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

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. Unzip the zip file located in the directory containing the newly saved system configuration information.

    The directory containing the newly saved system configuration information contains multiple files, including the updated XML system configuration file and HTML installation template. However, the zip file contains additional cluster-specific XML configuration files.

15. Locate the cluster-specific XML configuration file for the new cluster.

    The new cluster-specific XML configuration file is located in the unzipped directory that you created in the previous step. The names of the cluster-specific XML configuration files

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

contain a numeric sequence that reflects the order of cluster creation, with the newest cluster having the highest number.

16. Use the cluster-specific XML configuration file to deploy the new cluster.

    Use the OEDA deployment utility (`install.sh`) in conjunction with the cluster-specific XML configuration file for the new cluster.

    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 they 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 cluster4.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 cluster4.xml -l

    1. Validate Configuration File
    2. Setup Required Files
    3. Create Users
    4. Setup Cell Connectivity
    5. Verify RDMA Network Connectivity
    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 cluster4.xml -s 1
    ```


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

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 cluster4.xml -r 12-16
```

Related Topics

- Configuring Oracle Exadata Database Machine Using OEDA
  Use Oracle Exadata Deployment Assistant (OEDA) to configure your engineered system.

6.23 Using IPv4 and IPv6 in an Existing Deployment

You can use OEDACLI to reconfigure an existing system with IPv4, IPv6, or both.

Starting with Oracle Exadata System Software release 26.1.0, you can use the Oracle Exadata Deployment Assistant (OEDA) Command Line Utility (OEDACLI) to easily and reliably reconfigure an existing Exadata deployment, using any combination of IPv4 and IPv6, without downtime or server reboots. With this capability, you can migrate Exadata database servers between IPv4 and IPv6 (in either direction), or implement a dual-stack configuration, allowing IPv4 and IPv6 to coexist together.

- Adding Network Interfaces to an Existing Guest
- Removing Network Interfaces from an Existing Guest

6.23.1 Adding Network Interfaces to an Existing Guest

You can use OEDACLI to add network interfaces to an existing virtual machine (VM) guest without downtime or server reboots.

> **Note**
>
> You can add interfaces using IPv6 to an existing guest using IPv4, or add IPv4 interfaces to a guest already using IPv6. The following examples focus on adding IPv6 interfaces. However, you can adapt the same commands to add IPv4 interfaces. Substitute your own hostname and IP address values as required.

Adding a Network Interface

To add a network interface to an existing guest, use the OEDACLI `ADD NETWORK` command.

For example, the following command sequence adds an IPv6 client network interface to an existing guest:

```
oedacli> ADD NETWORK NETWORKTYPE=CLIENT HOSTNAME=exa01vm01-client
IP='2001:db8::100' MASTER=bondeth0 NETMASK=64
         DOMAINNAME=example.com GATEWAY='2001:db8::' WHERE
HOSTNAME=exa01vm01.example.com
oedacli> save action
oedacli> merge action
oedacli> deploy actions
```

In the example:

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

- `NETWORKTYPE=CLIENT` specifies the addition a client network interface. Alternatively, you could use `NETWORKTYPE=ADMIN` to add a management network interface.
- `MASTER=bondeth0` specifies that the new interface is associated with the existing `bondeth0` network interface. This is the typical setting for client network interfaces on a guest. The typical setting for management network interfaces is `MASTER=eth0`.
- The other arguments specify attributes of the IPv6 interface. Note that for IPv6 interfaces, `NETMASK` specifies the prefix length, which is the IPv6 equivalent of the subnet mask in IPv4.

Adding a Cluster VIP

To add a cluster virtual IP (VIP) to an existing guest, use the OEDACLI `ADD VIP` command.

For example, the following command sequence adds an IPv6 VIP to an existing guest:

```
oedacli> ADD VIP NAME=exa01vm01-vip IP='2001:db8::101' NETMASK=64
         DOMAINNAME=example.com WHERE HOSTNAME=exa01vm01.example.com
oedacli> save action
oedacli> merge action
oedacli> deploy actions
```

Adding a Cluster SCAN

To add a cluster Single Client Access Name (SCAN) to an existing cluster, use the OEDACLI `ADD SCAN` command.

For example, the following command sequence adds an IPv6 SCAN to an existing cluster:

```
oedacli> ADD SCAN SCANNAME=exa01-scan
         SCANIPS="2001:db8::111,2001:db8::112,2606:2001:db8::113"
         SUBNET="2001:db8::/64" NETMASK=64
         SCANTYPE=ipv6 GATEWAY="2001:db8::" WHERE CLUSTERNUMBER=1
oedacli> save action
oedacli> merge action
oedacli> deploy actions
```

In the example:

- `SCANTYPE=ipv6` specifies that the SCAN uses IPv6. If `SCANTYPE` is not specified, IPv4 is assumed. Or, you could explicitly specify `SCANTYPE=ipv4` for an IPv4 SCAN.
- The `WHERE` clause identifies the existing cluster using the cluster number recorded in the Exadata XML configuration file (`es.xml`).
- The other arguments specify attributes of the IPv6 SCAN interface.

6.23.2 Removing Network Interfaces from an Existing Guest

For existing guests with IPv6 and IPv4 network interfaces, you can use OEDACLI to remove either the IPv4 or IPv6 interfaces network interfaces. This is typically done as the final phase of a migration between IPv4 and IPv6 (in either direction).

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

> **Note**
>
> The following examples focus on removing IPv4 interfaces. However, you can adapt the same commands to remove IPv6 interfaces. Substitute your own hostname and IP address values as required.

Removing a Network Interface

To remove a network interface from an existing guest, use the OEDACLI `DELETE NETWORK` command.

For example, the following command sequence removes the specified IPv4 client network interface from an existing guest:

```
oedacli> DELETE NETWORK WHERE IP=203.0.113.100 NETWORKTYPE=CLIENT
oedacli> save action
oedacli> merge action
oedacli> deploy actions
```

In the example, `NETWORKTYPE=CLIENT` specifies the removal of a client network interface. Alternatively, you could use `NETWORKTYPE=ADMIN` to remove a management network interface.

> **Note**
>
> If the interface you are removing is associated with the default gateway, you will get an error instructing you to move the default gateway to another network. For example, if you are trying to remove a client network interface, you can move the default gateway to the management network using the following command sequence:
>
> ```
> oedacli> ALTER MACHINE DEFAULTGATEWAYNETHOSTNAME=exa01vm01-admin WHERE
> HOSTNAME=exa01vm01-client
> oedacli> save action
> oedacli> merge action
> oedacli> deploy actions
> ```
>
> In the example, `exa01vm01-client` represents the hostname associated with the client network and `exa01vm01-admin` represents the hostname associated with the management network.

Removing a Cluster VIP

To remove a cluster virtual IP (VIP) from an existing guest, use the OEDACLI `DELETE VIP` command.

For example, the following command sequence removes an IPv4 VIP from an existing guest:

```
oedacli> DELETE VIP WHERE VIPIP=203.0.113.101
oedacli> save action
oedacli> merge action
oedacli> deploy actions
```


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

Remove a Cluster SCAN

To remove a cluster Single Client Access Name (SCAN) from an existing cluster, use the OEDACLI `DELETE SCAN` command.

For example, the following command sequence removes an IPv4 SCAN from an existing cluster:

```
oedacli> DELETE SCAN WHERE SCANNAME=exa01-scan SCANTYPE=ipv4 CLUSTERNUMBER=1
oedacli> save action
oedacli> merge action
oedacli> deploy actions
```

In the example:

- `SCANTYPE=ipv4` specifies that the SCAN being removed uses IPv4. Alternatively, you could use `SCANTYPE=ipv6` to remove an IPv6 SCAN.
- The other arguments identify the SCAN using the cluster number and SCAN name, as recorded in the Exadata XML configuration file (`es.xml`).

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

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.

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

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	Yes	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? • 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 Management Network Switch is being replaced, then have you arranged for the replacement switch installation?				

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

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: • 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?				

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

Data Center Room Checklist Items	Yes	No	N/A	Comment
If the data center has a raised floor: • 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
- Flooring Requirements
  Oracle Exadata Rack may be installed on raised floor or solid floor environments.
- Installing Oracle Exadata or Oracle Exadata Storage Expansion Rack at the Site

A.3 Data Center Environment Checklist

Use this checklist to ensure that the data center environment requirements have been addressed.

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

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: • 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.
- Ventilation and Cooling Requirements
  Always provide adequate space in front and behind the rack to allow for proper ventilation.
- General Environmental Requirements

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

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
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? Dimensions of the packaged rack: • Height: 2159 mm (85 inches) • Width: 1219 mm (48 inches) • Depth: 1575 mm (62 inches) Dimensions of the unpackaged rack: • Height: 1998 mm (78.66 inches) • Width: 600 mm with side panels (23.62 inches) • Depth 1200 mm (47.24 inches)				
Are there any stairs, ramps, or thresholds that are of concern? If yes, then provide details.				
Are all access route incline angles within the permitted range (6 degrees)?				
Is the access route free of any obstacles that would expose the equipment to shock?				

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

Access Route Considerations	Yes	No	N/A	Comment
Are all the surfaces acceptable for rolling the new unpacked and packed equipment?				
If a pallet jack is to be used: • Can the pallet jack support the weight of the rack? • Are the pallet jack tines compatible with the shipping pallet?				
If there are stairs, is a loading elevator available for the equipment?				
If an elevator is to be used: • 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?				

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

Related Topics

- Table: Access Route Requirements.
- Rack Weights for Oracle Exadata
  Review the weight of your Oracle Exadata Rack to ensure that it can be delivered and located 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 been reviewed? Power requirements listed in "Electrical Power Requirements".				
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?				

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

Facility Power Considerations	Yes	No	N/A	Comment
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?				
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?				

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

Safety Checklist Considerations	Yes	No	N/A	Comment
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.

> **See Also**
>
> "Powering On and Off Oracle Exadata Rack" in the *Oracle Exadata 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?				

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

Logistics Checklist Considerations	Yes	No	N/A	Comment
Are there any additional security access issues?				
Is computer room access available for installation personnel?				
Are laptops allowed in the data center?				
Are cell phones allowed in the data center?				
Are cameras allowed in the data center?				
Does the building have a delivery dock?				
Is there a delivery/ unpacking/staging area?				
Is inside delivery planned (direct to the final rack location in the data center room)?				
If the delivery is not inside, then is the site prepared for uncrating?				
Is the delivery/ unpacking/staging area protected from the elements?				
Does the building have adequate receiving space?				
Is the unpacking area air-conditioned to avoid thermal shock for various hardware components?				
Will sufficient moving personnel be available to transport the rack?				

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

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 *outside* 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.				
Is a tail or side lift required on the delivery carrier to unload the equipment at the delivery dock?				

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

Logistics Checklist Considerations	Yes	No	N/A	Comment
Will any special equipment be required to place the rack in the data center room? Examples: • 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

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 `checkip.sh` 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 Rack will be installed?				

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

Network Configuration Considerations	Yes	No	N/A	Comment
Will the network cables that will connect to the Oracle 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 Management Network Switch have IP routing disabled (recommended)?				

Related Topics

- Understanding the Network Requirements for Oracle Exadata
  Review the network requirements for Oracle Exadata 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?				

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

Auto Service Request Considerations	Yes	No	N/A	Comment
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 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?				

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

Re-racking Considerations	Yes	No	N/A	Comment
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?				
Do the target rack dimensions meet the following requirements? • Height: 42 RU • Width: 600 mm (23.62 inches) • Depth: 1112 mm (43.78 inches) without front and rear doors If the rack is less than 42 RU, then the customer must provide compatible PDUs to install in the target rack.				
Is the distance between the front and rear mounting planes between the minimum of 610 mm and the maximum 915 mm (24 inches to 36 inches)?				
Is the clearance depth in the front of the front mounting plane (distance to the front cabinet door) at least 25.4 mm (1 inch)?				
Does the target rack meet the following minimum load capacity? • 19 kg (41.89 lb) / RU				

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

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?				

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

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 two equivalent PDUs with equivalent power capacity and at least 17 x 10A C13 plugs per PDU?				

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

PDU Re-racking Considerations	Yes	No	N/A	Comment
Can the customer ensure that each PDU and its circuits can 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.				
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

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

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 maintenance overview and guidelines in My Oracle Support document KB649655?				
Have you reviewed the Oracle Exadata System Software Maximum Availability Architecture (MAA) Web site for post-deployment steps to achieve MAA?				

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

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 document KB74841.				
Will you follow Oracle Exadata Database Machine monitoring best practices available in My Oracle Support document KB145549?				
Have you reviewed the recommended software for Oracle Exadata Database Machine in My Oracle Support document KB153930?				
Have you reviewed Oracle Exadata Database Machine best practices in My Oracle Support document KB480584?				

Related Topics

- *Oracle AI Database High Availability Overview and Best Practices*
- Exadata Database Machine Software Maintenance Overview and Guidelines (My Oracle Support Doc ID KB649655)
- Oracle Maximum Availability Architecture website
- Oracle Exadata Database Machine Exachk (My Oracle Support Doc ID KB74841)
- Exadata Health and Resource Usage Monitoring Technical Brief (My Oracle Support Doc ID KB145549)
- Exadata Database Machine and Exadata Storage Server Supported Versions (My Oracle Support Doc ID KB153930)
- Oracle Exadata Best Practices (My Oracle Support Doc ID KB480584)

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