

Zero Data Loss Recovery Appliance Owner's Guide



Release 12.2

E88041-04

July 2018

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

ORACLE®

Zero Data Loss Recovery Appliance Owner's Guide, Release 12.2

E88041-04

Copyright © 2014, 2017, Oracle and/or its affiliates. All rights reserved.

Contributors: Glenn Maxey

Primary Author: Terence Buencamino

Contributors: Donna Carver, Doug Archambault, Andrew Babb, Tim Chien, Sam Corso, Donna Cooksey, Caroline Johnston, Leslie Keller, Lawrence To, Ericka Washington

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

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

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

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

Intel and Intel Xeon 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, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

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

Contents

Preface

Audience	xxii
Documentation Accessibility	xxii
Related Documentation	xxii
Conventions	xxiii

Part I Preinstallation

1 Introduction to Recovery Appliance

1.1	About Recovery Appliance	1-1
1.2	Recovery Appliance Rack Components	1-2
1.3	Compute Server Components	1-3
1.4	Storage Server Components	1-5
1.5	Spares Kit Components for Recovery Appliance	1-7
1.6	About Tape Backup Infrastructure	1-8
1.6.1	Oracle Recommended Stack	1-8
1.6.2	Third-Party Tape Systems	1-8
1.7	Restrictions on Modifying Recovery Appliance	1-9

2 Site Requirements for Recovery Appliance

2.1	General Environmental Requirements	2-1
2.2	Space Requirements	2-5
2.2.1	Receiving and Unpacking Requirements	2-6
2.2.2	Maintenance Access Requirements	2-7
2.3	Flooring Requirements	2-7
2.4	Electrical Power Requirements	2-8
2.4.1	PDUs for Recovery Appliance	2-8
2.4.1.1	Single-Phase PDUs for North America, South America, Japan, and Taiwan	2-8
2.4.1.2	Three-Phase PDUs for North America, South America, Japan, and Taiwan	2-9

2.4.1.3	Single-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific	2-10
2.4.1.4	Three-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific	2-11
2.4.2	Facility Power Requirements	2-12
2.4.3	Circuit Breaker Requirements	2-12
2.4.4	Grounding Guidelines	2-13
2.5	Temperature and Humidity Requirements	2-13
2.6	Ventilation and Cooling Requirements	2-14
2.7	Network Connection Requirements for Recovery Appliance	2-15
2.8	Ensuring That the Site Is Ready	2-16
2.8.1	Generate the Configuration Files	2-16
2.8.2	Configure the Network	2-16
2.8.3	Prepare the Site	2-16

3 Preparing Your Networks for Recovery Appliance

3.1	Overview of Network Requirements	3-1
3.1.1	About the Network Components and Interfaces	3-1
3.1.2	About the Networks	3-3
3.1.3	Example of Network Connections for Recovery Appliance	3-5
3.1.4	Connecting Recovery Appliance Rack Components to the Networks	3-6
3.1.5	Connecting Recovery Appliance to a Tape Library	3-7
3.1.5.1	Oracle Recommended Stack	3-7
3.1.5.2	Third-Party Tape Systems	3-8
3.1.6	Using Network VLAN Tagging with Recovery Appliance	3-9
3.2	Registering Recovery Appliance in the Domain Name System	3-10
3.3	Factory IP Address Settings	3-10
3.4	Port Assignments When Using a Firewall	3-11

4 Setting Up Auto Service Request

4.1	Understanding Auto Service Request	4-1
4.2	Getting Ready to Install ASR	4-2
4.3	Installing ASR Manager	4-3
4.4	Verifying the ASR Manager Installation	4-3
4.5	About the Trap Destinations on Recovery Appliance	4-4
4.6	Troubleshooting ASR	4-5

5 Using the Deployment Assistant

5.1	Overview of the Deployment Assistant	5-2
-----	--------------------------------------	-----

5.2	Generating the Configuration Files	5-2
5.3	About the Configuration Files	5-4
5.4	Customer Details Page	5-5
5.5	Hardware Selection Page	5-5
5.6	Define Customer Networks Page	5-6
5.6.1	Network Configuration Choices	5-7
5.6.2	Cabling Considerations	5-8
5.7	Administration Network Page	5-8
5.7.1	Modifying the Administration Host Name Format	5-9
5.8	Ingest Network Page	5-10
5.8.1	Modifying the Ingest Name Format	5-11
5.9	InfiniBand Network Configuration Page	5-12
5.9.1	Modifying the InfiniBand Host Name Format	5-12
5.10	Replication Network Page	5-13
5.10.1	Modifying the Replication Host Name Format	5-13
5.11	Identify Compute Node OS Page	5-14
5.12	Review and Edit Page	5-14
5.13	Define Clusters Page	5-14
5.14	Cluster Page	5-15
5.14.1	Cluster	5-15
5.14.2	Users and Groups	5-16
5.14.3	Software Locations	5-16
5.14.4	Disk Group Details	5-16
5.14.5	Initial Database	5-16
5.14.6	Ingest Network	5-17
5.14.7	Replication Network	5-17
5.15	Review and Edit Page	5-18
5.16	Replication Network SCAN Page	5-18
5.17	Cell Alerting Page	5-19
5.18	Platinum Configuration Page	5-20
5.18.1	Customer Details	5-20
5.18.2	Platinum Gateway Details	5-21
5.18.3	Platinum Agent Configuration Details	5-22
5.19	Oracle Configuration Manager Page	5-23
5.20	Auto Service Request Page	5-24
5.21	Grid Control Agent Page	5-25
5.22	Tape Library Connectivity Page	5-26
5.23	Comments Page	5-27
5.24	Generate Page	5-27
5.25	Finish Page	5-27

Part II Installation and Configuration

6 Installing the Recovery Appliance at the Site

6.1	Reviewing Safety Guidelines	6-1
6.2	Unpacking ZDLRA Rack	6-2
6.2.1	Contents of the Shipping Kit	6-3
6.2.2	Removing ZDLRA Rack from the Shipping Crate	6-4
6.3	Placing ZDLRA Rack in Its Allocated Space	6-5
6.3.1	Moving ZDLRA Rack	6-5
6.3.2	Stabilizing ZDLRA Rack	6-6
6.3.2.1	Stabilize ZDLRA Rack with Leveling Feet	6-7
6.3.2.2	Stabilize ZDLRA Rack X7 with Leveling Feet	6-8
6.3.3	Attaching a Ground Cable (Optional)	6-10
6.4	Installing the Fiber Channel Adapters	6-11
6.5	Acclimating the Recovery Appliance	6-12
6.6	Powering on the System for the First Time	6-12
6.6.1	Inspecting Recovery Appliance After It Is in Place	6-12
6.6.2	Connecting Power Cords	6-13
6.6.3	Powering Up Recovery Appliance	6-14

7 Expanding a Recovery Appliance Rack with Additional Storage Servers

7.1	Preparing to Extend Recovery Appliance	7-1
7.1.1	Estimating InfiniBand Cable Path Lengths	7-1
7.1.2	Obtaining the Current Configuration Information	7-2
7.1.3	Generating the Network Configuration Files	7-3
7.1.4	Moving the Audit and Diagnostic Files	7-4
7.1.5	Synchronizing Release and Patch Levels	7-4
7.2	Precautions and Guidelines	7-4
7.2.1	General Safety Precautions	7-4
7.2.2	InfiniBand Cable Precautions	7-4
7.2.3	InfiniBand Cable Bundling Guidelines	7-5
7.2.4	Cable Management Arm Guidelines	7-5
7.3	Adding Storage Servers to a Rack	7-6
7.3.1	Checking the Recovery Appliance Rack	7-6
7.3.2	Unpacking and Labeling the New Storage Servers	7-6
7.3.3	Recovery Appliance X7, X6, and X5 Population Rules	7-8
7.3.4	Removing the Doors	7-8
7.3.5	Adding New Servers	7-10

7.3.6	Cabling the Storage Servers	7-15
7.3.7	Closing the Rack	7-17
7.4	Making New Storage Available	7-18

8 Extending Recovery Appliance with Multiple Racks

8.1	About Multirack Configurations	8-1
8.2	Extending a Rack by Adding Another Rack	8-2
8.3	Adding the Spine Switches	8-2
8.4	Floor and Underfloor Delivery of InfiniBand Cables	8-4
8.5	Cabling Two Racks Together	8-4
8.6	Cabling Several Racks Together	8-5
8.7	Configuring Multiple Racks	8-7

9 Configuring Ingest Over InfiniBand

9.1	About Ingest Over InfiniBand	9-1
9.2	Cabling Recovery Appliance to Oracle Exadata Database Machine	9-2
9.3	Adding VIPs to the Recovery Appliance InfiniBand Network	9-4
9.4	Updating Recovery Appliance With the New VIPs for InfiniBand	9-5
9.5	Update the Database Initialization Parameters	9-6

10 Configuring a Recovery Appliance Rack

10.1	Supporting Auto Service Request	10-1
10.1.1	Prerequisites for Using Auto Service Request	10-1
10.1.2	Checking an Existing ASR Manager Installation	10-1
10.1.3	Installing ASR Manager	10-2
10.2	Installing the Tape Hardware	10-2
10.3	Connecting Recovery Appliance to Your Networks	10-3
10.4	Configuring the InfiniBand Switches	10-3
10.4.1	Configuring an InfiniBand Switch	10-4
10.4.2	Setting the Time on an InfiniBand Switch	10-5
10.4.3	Setting the Serial Number on a Spine Switch	10-6
10.4.4	Checking the Health of an InfiniBand Switch	10-7
10.4.5	Setting a Spine Switch as the Subnet Manager Master	10-8
10.5	Configuring the Cisco Ethernet Switch	10-10
10.5.1	Scope of the Configuration	10-10
10.5.2	Prerequisites for Configuring the Ethernet Switch	10-10
10.5.3	Configuring the Ethernet Switch on the Data Center Network	10-11
10.5.4	Disabling Telnet Connections	10-13
10.5.5	Configuring SSH on the Ethernet Switch	10-13

10.5.6	Setting the Clock and Time Zone on the Ethernet Switch	10-14
10.5.7	Disabling Spanning Tree on the Ethernet Switch	10-15
10.5.8	Verifying the Ethernet Configuration	10-16
10.6	Configuring the Power Distribution Units	10-16
10.6.1	Assigning Network Addresses to the PDUs	10-17
10.6.2	Configuring the PDU System Time Settings	10-17
10.6.3	Upgrading the PDU Firmware	10-18
10.7	Checking the Health of the Compute Servers	10-19
10.8	Checking the Health of the Storage Servers	10-22
10.9	Verifying the InfiniBand Network	10-24
10.10	Configuring the IP Addresses of the Servers	10-24

11 Installing the Recovery Appliance Software

11.1	About the Installation Software	11-1
11.2	Software Installation Prerequisites	11-2
11.2.1	About Adding a Rack to an Existing Rack	11-2
11.2.2	About Adding New Storage Servers to an Existing Rack	11-2
11.3	Reclaiming Disk Space on the Compute Servers	11-3
11.4	Installing the Software on Recovery Appliance	11-3
11.5	Install Utility Syntax	11-6
11.5.1	Install Utility Parameters	11-6
11.5.2	Install Utility Syntax Examples	11-6
11.5.3	About the Install Utility Steps	11-7
11.6	Setting Up Enterprise Manager to Administer Recovery Appliance	11-10
11.6.1	Deploying the Enterprise Manager Agents	11-10
11.6.2	Discovering the Cluster	11-11
11.6.3	Discovering the Cluster Database Targets	11-11
11.6.4	Discovering the Recovery Appliance	11-12

Part III Maintenance of Recovery Appliance

12 Maintaining the Recovery Appliance Hardware

12.1	Cautions and Warnings	12-1
12.2	Determining the Server Model	12-2
12.3	Powering On and Off a Recovery Appliance Rack	12-2
12.3.1	Emergency Power-Off Procedure	12-2
12.3.1.1	Powering Off in an Emergency	12-2
12.3.1.2	About the Emergency Power-Off Switch	12-2
12.3.2	Shutting Down Recovery Appliance	12-3

12.3.2.1	Stopping Recovery Appliance Services	12-3
12.3.2.2	Powering Down the Servers	12-4
12.3.2.3	Powering the Network Switches	12-5
12.3.3	Starting Up Recovery Appliance	12-5
12.3.3.1	Starting Up Recovery Appliance Components	12-5
12.3.3.2	Startup Sequence	12-6
12.3.3.3	Powering On Servers Remotely	12-6
12.3.3.4	Starting the Recovery Appliance Software	12-6
12.4	Replacing the Disk Controller Batteries	12-7
12.4.1	Monitoring Batteries in the Compute Servers	12-8
12.4.2	Replacing Batteries in Disk Controllers	12-8
12.5	Replacing a Power Distribution Unit	12-9
12.5.1	PDU Replacement Guidelines	12-9
12.5.2	Replacing a PDU	12-9
12.6	Resetting a Non-Responsive Oracle ILOM	12-11
12.6.1	Resetting Oracle ILOM Using SSH	12-12
12.6.2	Resetting Oracle ILOM Using the Remote Console	12-12
12.6.3	Resetting Oracle ILOM Using IPMItool	12-12
12.6.4	Resetting Oracle ILOM By Removing Power	12-12
12.7	Maintaining the Compute Servers	12-13
12.7.1	Verifying the RAID Status of a Compute Server	12-13
12.8	Reimaging a Compute Server	12-14
12.8.1	Contacting Oracle Support Services	12-14
12.8.2	Downloading the Latest Release of the Cluster Verification Utility	12-15
12.8.3	Removing the Failed Compute Server from the Cluster	12-15
12.8.4	Preparing the USB Flash Drive for Imaging	12-16
12.8.5	Copying the Image to the New Compute Server	12-17
12.8.6	Configuring the Replacement Compute Server	12-17
12.8.7	Preparing the Replacement Compute Server for the Cluster	12-18
12.8.8	Applying Patch Bundles to a Replacement Compute Server	12-20
12.8.9	Cloning the Oracle Grid Infrastructure	12-20
12.8.10	Clone Oracle Database Homes to the Replacement Compute Server	12-22
12.9	Maintaining the Storage Servers	12-24
12.9.1	Shutting Down a Storage Server	12-24
12.9.2	Enabling Network Connectivity Using the Diagnostics ISO	12-25
12.10	Maintaining the Physical Disks of Storage Servers	12-26
12.10.1	About System Disks and Data Disks	12-27
12.10.2	Monitoring the Status of Physical Disks	12-27
12.10.3	What Happens When Disk Errors Occur?	12-28
12.10.4	About Detecting Underperforming Disks	12-28
12.10.5	About Rebalancing the Data	12-29

12.10.6	Monitoring Hard Disk Controller Write-Through Cache Mode	12-30
12.10.7	Replacing a Failed Physical Disk	12-31
12.10.8	Replacing a Faulty Physical Disk	12-32
12.10.9	Removing an Underperforming Physical Disk	12-34
12.10.10	Moving All Drives from One Storage Server to Another	12-34
12.10.11	Removing and Replacing the Same Physical Disk	12-36
12.10.12	Reenabling a Rejected Physical Disk	12-36
12.11	Maintaining the Flash Disks of Storage Servers	12-37
12.11.1	About the Flash Disks	12-37
12.11.2	Faulty Status Indicators	12-38
12.11.3	Identifying Flash Disks in Poor Health	12-39
12.11.4	Identifying Underperforming Flash Disks	12-39
12.11.5	When Is It Safe to Replace a Faulty Flash Disk?	12-40
12.11.6	Replacing a Failed Flash Disk	12-40
12.11.7	Replacing a Faulty Flash Disk	12-41
12.11.8	Removing an Underperforming Flash Disk	12-42
12.11.9	About Write-Back Flash Cache	12-43
12.12	Replacing a Disk Controller Battery Backup Unit	12-43
12.12.1	Replacing a Disk Controller BBU on a Compute Server	12-44
12.12.2	Replacing a Disk Controller BBU on a Storage Server	12-45
12.13	Using the Storage Server Rescue Procedure	12-45
12.13.1	First Steps Before Rescuing the Storage Server	12-46
12.13.1.1	If the Server Has Normal Redundancy	12-46
12.13.1.2	If the Server Has High Redundancy	12-46
12.13.2	About the Rescue Procedure	12-46
12.13.3	Rescuing a Server Using the CELLBOOT USB Flash Drive	12-47
12.13.4	Reconfiguring the Rescued Storage Server	12-48
12.13.5	Recreating a Damaged CELLBOOT USB Flash Drive	12-49

13 Reconfiguring the Hardware

13.1	Changing Component Passwords	13-1
13.1.1	Changing the Compute Server Passwords	13-1
13.1.1.1	Changing the User Account Password	13-1
13.1.1.2	Changing the GRUB Account Password	13-2
13.1.2	Changing the Storage Server Passwords	13-2
13.1.3	Creating New Users and Changing Passwords on a PDU	13-2
13.1.4	Changing the Oracle ILOM Password	13-3
13.1.5	Changing an InfiniBand Switch Password	13-3
13.1.6	Changing the Cisco Ethernet Switch Password	13-4
13.1.7	Changing the Recovery Appliance Database User Password	13-4

13.1.8	Changing OSB Tape Backup Passwords	13-5
13.2	Changing the DNS Servers	13-5
13.2.1	Changing the DNS Server Address on an InfiniBand Switch	13-6
13.2.2	Changing the DNS Server Address on the Cisco Ethernet Switch	13-6
13.2.3	Changing the DNS Server Address on the Compute Server	13-7
13.2.4	Changing the DNS Server on Storage Servers	13-7
13.3	Changing the NTP Servers	13-8
13.3.1	Setting the NTP Server Address on the Compute Servers	13-8
13.3.2	Setting the NTP Server Address on the InfiniBand Switches	13-9
13.3.3	Set the NTP Server Address on the Cisco Ethernet Switch	13-9
13.3.4	Setting the NTP Server on the Storage Servers	13-10
13.3.5	Restart Recovery Appliance	13-10
13.4	Changing the Time Zone Settings	13-11
13.4.1	Changing the Time Zone Settings on the Storage Servers	13-11
13.4.2	Change Time Zone Settings on the Compute Servers	13-13
13.4.3	Change the Time Zone Settings on the InfiniBand Switches	13-13
13.4.4	Change the Time Zone Settings on the Cisco Switch	13-13

Part IV Appendixes

A RACLI Overview

A.1	racli add	A-4
A.2	racli add acfs	A-4
A.3	racli add cloud_storage	A-4
A.4	racli add credential	A-5
A.5	racli add credential_wallet	A-6
A.6	racli add keystore	A-7
A.7	racli add nfs	A-7
A.8	racli add storage_cells	A-8
A.9	racli add storage_disks	A-8
A.10	racli add vpc_user	A-9
A.11	racli alter	A-10
A.12	racli alter cloud_storage	A-10
A.13	racli alter credential_wallet	A-11
A.14	racli alter keystore	A-11
A.15	racli alter passphrase	A-12
A.16	racli alter rasys_user	A-12
A.17	racli alter vpc_user	A-13
A.18	racli disable	A-13
A.19	racli disable cloud_storage	A-13

A.20	racli disable keystore	A-14
A.21	racli disable nfs	A-14
A.22	racli enable	A-14
A.23	racli enable cloud_storage	A-14
A.24	racli enable keystore	A-15
A.25	racli enable nfs	A-15
A.26	racli export	A-16
A.27	racli export appliance	A-16
A.28	racli help	A-16
A.29	racli install	A-16
A.30	racli install appliance	A-16
A.31	racli install okv_endpoint	A-17
A.32	racli list	A-18
A.33	racli list cloud_storage	A-18
A.34	racli list nfs	A-18
A.35	racli list okv_endpoint	A-19
A.36	racli patch	A-19
A.37	racli patch appliance	A-19
A.38	racli remove	A-20
A.39	racli remove acfs	A-20
A.40	racli remove cloud_storage	A-20
A.41	racli remove keystore	A-21
A.42	racli remove nfs	A-21
A.43	racli remove vpc_user	A-21
A.44	racli start	A-22
A.45	racli start nfs	A-22
A.46	racli status	A-22
A.47	racli status appliance	A-22
A.48	racli status cloud_storage	A-22
A.49	racli status disk_group	A-23
A.50	racli status keystore	A-23
A.51	racli status nfs	A-23
A.52	racli status okv_endpoint	A-23
A.53	racli stop	A-24
A.54	racli stop nfs	A-24
A.55	racli uninstall	A-24
A.56	racli uninstall appliance	A-24
A.57	racli uninstall okv_endpoint	A-25
A.58	racli update	A-26
A.59	racli update storage_location	A-26
A.60	racli update support_tools	A-26

A.61	racli upgrade	A-26
A.62	racli upgrade appliance	A-26
A.63	racli version	A-27

B Default User Accounts

C Error Messages

C.1	Contacting Oracle Support Services	C-1
C.2	Generic Alert Messages: HALRT-00001 to HALRT-00500	C-2
C.3	Temperature Alert Messages: HALRT-00513 to HALRT-00600	C-5
C.4	Fan Alert Messages: HALRT-00616 to HALRT-00700	C-6
C.5	Power Supply/Converters Alert Messages: HALRT-00717 to HALRT-00800	C-6
C.6	Network Alert Messages: HALRT-00912 to HALRT-01000	C-7
C.7	Hard Disk and Flash Disk Alert Messages: HALRT-01015 to HALRT-01200	C-8
C.8	Accelerator/Controller Alert Messages: HALRT-01209 to HALRT-01300	C-28
C.9	Voltage Alert Messages: HALRT-01301 to HALRT-01400	C-32
C.10	Physical Disk Alert Messages: HALRT-01401 to HALRT-01500	C-32
C.11	Auto Service Request Alert Messages for Storage Servers: HALRT-02001 to HALRT-02013	C-34
C.12	HDD Controller Battery Alert Messages: HALRT-02101 to HALRT-02200	C-36
C.13	Generic Actions for Alert Messages: HALRT-10001 to HALRT-10500	C-39
C.14	Generic Actions for Alert Messages: HALRT-10503 to HALRT-11000	C-44
C.15	Generic Actions for Alert Messages: HALRT-11001 to HALRT-11100	C-45

D Replacement Units

D.1	About Warranty Coverage	D-1
D.2	Understanding Repair Categories	D-1
D.3	LED Status Descriptions	D-2
D.3.1	Oracle Server LEDs	D-2
D.3.2	Sun Datacenter InfiniBand Switch 36 LEDs	D-4
D.4	Parts for Compute Servers	D-6
D.4.1	Oracle Server X7 Oracle Database Server Parts	D-7
D.4.2	Oracle Server X6-2 and X5-2 Parts	D-13
D.4.3	Sun Server X4-2 Parts	D-18
D.5	Parts for Storage Servers	D-20
D.5.1	Oracle Exadata Storage Server X7-2 High Capacity Server Parts	D-21
D.5.2	Oracle Server X6-2L and X5-2L Parts	D-24
D.5.3	Sun Server X4-2L Parts	D-30
D.6	Parts for InfiniBand Switches	D-33

D.7	Parts for Cisco Switch	D-34
D.8	Parts for Power Distribution Units	D-34
D.9	Cables and Cable Accessories for Recovery Appliance	D-35

E In-Rack Cabling Tables for Recovery Appliance X7, X6, and X5

E.1	Cable Color Coding	E-1
E.2	Recovery Appliance X7, X6, and X5 Rack Layout	E-1
E.3	Oracle Integrated Lights Out Manager Cabling Connections	E-3
E.4	Administrative Gigabit Ethernet Cabling Connections	E-4
E.5	Single-Phase Power Distribution Unit Cabling Connections	E-5
E.6	Three-Phase Power Distribution Unit Cabling Connections	E-6
E.7	InfiniBand Switch-to-Switch Cabling Connections	E-7
E.8	InfiniBand Switch-to-Server Cabling Connections	E-8

F In-Rack Cabling Tables for Recovery Appliance X4

F.1	Recovery Appliance X4 Rack Layout	F-1
F.2	Oracle Integrated Lights Out Manager Cabling Connections	F-3
F.3	Administrative Gigabit Ethernet Cabling Connections	F-4
F.4	Single-Phase Power Distribution Unit Cabling Connections	F-5
F.5	Three-Phase Power Distribution Unit Cabling Connections	F-6
F.6	InfiniBand Switch-to-Switch Cabling Connections	F-7
F.7	InfiniBand Switch-to-Server Cabling Connections	F-7

G Multirack Cabling Tables

G.1	About Multirack Cabling	G-1
G.1.1	Cables Provided With a Spine Switch	G-2
G.1.2	Optical Cables Required for Four or More Racks	G-2
G.1.3	Cabling Three Switches Within a Rack	G-3
G.1.4	Cabling the Switches in Two Racks	G-4
G.1.5	Cabling the Switches in Up to Eight Racks	G-5
G.1.6	Cabling Switches in Up to 18 Racks	G-5
G.2	Key to Cabling Table Abbreviations	G-7
G.3	Two-Rack Cabling	G-7
G.4	Three-Rack Cabling	G-8
G.5	Four-Rack Cabling	G-10
G.6	Five-Rack Cabling	G-12
G.7	Six-Rack Cabling	G-15
G.8	Seven-Rack Cabling	G-18
G.9	Eight-Rack Cabling	G-22

G.10	Nine-Rack Cabling	G-27
G.11	10-Rack Cabling	G-27
G.12	11-Rack Cabling	G-28
G.13	12-Rack Cabling	G-29
G.14	13-Rack Cabling	G-29
G.15	14-Rack Cabling	G-30
G.16	15-Rack Cabling	G-31
G.17	16-Rack Cabling	G-31
G.18	17-Rack Cabling	G-32
G.19	18-Rack Cabling	G-33

Index

List of Figures

2-1	Low-Voltage Single Phase Power Connector	2-9
2-2	Low-Voltage Three-Phase Power Connector	2-10
2-3	High-Voltage Single Phase Power Connector	2-11
2-4	High-Voltage Three-Phase Power Connector	2-12
2-5	Typical Data Center Configuration for Perforated Floor Tiles	2-15
3-1	X7 Back—plane External Network Connectivity	3-2
3-2	Network Diagram for Recovery Appliance	3-6
3-3	Network Connections to the Recovery Appliance Rack Components	3-7
3-4	Recovery Appliance Connection to an Oracle Tape System	3-8
3-5	Recovery Appliance Connection to a Third-Party Tape System	3-9
4-1	Auto Service Request Network Connections	4-3
5-1	Oracle Exadata Deployment Assistant Welcome Page	5-4
6-1	ZDLRA Rack in Shipping Crate	6-3
6-2	Carefully Push ZDLRA Rack from Behind	6-6
6-3	Location of Leveling Feet on ZDLRA Rack	6-7
6-4	Securing ZDLRA Rack Using the Leveling Feet	6-8
6-5	Location of Leveling Feet on ZDLRA Rack X7	6-9
6-6	Securing ZDLRA Rack Using the Leveling Feet	6-10
6-7	Earth Ground Attachment Bolt Location	6-11
6-8	Power Cord Routing Example	6-14
6-9	PDU Switch Locations	6-14
7-1	Server Expansion Kit Components	7-7
7-2	Removing the Rack Doors	7-9
7-3	Removing the Filler Panels	7-10
7-4	Location of the Trough	7-11
7-5	Aligning the Side-Rail Assembly With the Rack	7-12
7-6	Inserting the Server with Mounting Brackets into the Slide Rails	7-14
9-1	Recovery Appliance Connected to an Exadata Quarter Rack	9-2
9-2	Recovery Appliance Connected to an Exadata Half or Full Rack	9-3
9-3	Recovery Appliance Connected to Multiple Exadata Racks	9-3
D-1	Front View of Oracle Server X7 Oracle Database Server	D-7
D-2	Rear View of X7 Oracle Database Server	D-7
D-3	Front View of Oracle Server X6-2 or X5-2	D-13
D-4	Rear View of Oracle Server X6-2 or X5-2	D-13
D-5	Front View of Sun Server X4-2	D-18

D-6	Rear View of Sun Server X4-2	D-18
D-7	Front View of Oracle Exadata Storage Server X7-2 High Capacity Server	D-21
D-8	Rear View of Oracle Exadata Storage Server X7-2 High Capacity Server	D-21
D-9	Front View of Oracle Server X6-2L or X5-2L	D-24
D-10	Rear View of Oracle Server X6-2L or X5-2L	D-25
D-11	Front View of Sun Server X4-2L	D-30
D-12	Rear View of Sun Server X4-2L	D-30
E-1	Rack Layout of the Recovery Appliance X7, X6, or X5 Full Rack	E-2
E-2	Rack Layout of the Recovery Appliance X7, X6, or X5 Minimum Configuration	E-3
F-1	Rack Layout of the Recovery Appliance X4 Full Rack	F-2
F-2	Rack Layout of the Recovery Appliance X4 Minimum Configuration	F-3
G-1	In-Rack Connections Between Spine Switch and Leaf Switches	G-3
G-2	In-Rack Connections Between the Servers and the Leaf Switches	G-4
G-3	Connecting Switches Across Two Racks	G-5
G-4	Connecting Switches Across Eight Racks	G-5
G-5	Connecting Switches Across 18 Racks	G-6

List of Tables

1-1	Components of a Recovery Appliance X7 Rack	1-2
1-2	Components of a Recovery Appliance X6 Rack	1-2
1-3	Components of a Recovery Appliance X5 Rack	1-2
1-4	Components of a Recovery Appliance X4 Rack	1-3
1-5	Oracle X7-2 Compute Server Components	1-3
1-6	Oracle Server X6-2 Components	1-4
1-7	Oracle Server X5-2 Components	1-4
1-8	Sun Server X4-2 Components	1-5
1-9	Oracle Server X7-2L Storage Components	1-5
1-10	Oracle X6-2L Storage Server Components	1-5
1-11	Oracle Server X5-2L Components	1-6
1-12	Sun Server X4-2L Components	1-6
1-13	Oracle Tape System Components	1-8
2-1	Environmental Requirements for Recovery Appliance X7 Racks	2-1
2-2	Environmental Requirements for Recovery Appliance X6 Racks	2-2
2-3	Environmental Requirements for Recovery Appliance X5 Racks	2-4
2-4	Environmental Requirements for Recovery Appliance X4 Racks	2-5
2-5	Shipping Measurements	2-6
2-6	Access Route Requirements	2-6
2-7	Maintenance Access Requirements for Recovery Appliance	2-7
2-8	Floor Load Requirements for Recovery Appliance	2-7
2-9	Low-Voltage 15 kVA Single-Phase PDUs	2-8
2-10	Low-Voltage 15 kVA Three-Phase PDUs	2-9
2-11	High-Voltage 15 kVA Single Phase PDUs	2-10
2-12	High-Voltage 15 kVA Three-Phase PDUs	2-11
2-13	Temperature, Humidity and Altitude Requirements	2-13
3-1	Factory IP Addresses for Recovery Appliance	3-11
3-2	Open Ports for the Firewall	3-12
5-1	Customer Details Page	5-5
5-2	Hardware Selection Page	5-6
5-3	Define Networks Page	5-7
5-4	Administration Network Page	5-8
5-5	Administration Network Format Masks Dialog Box	5-9
5-6	Ingest Network Page	5-10
5-7	Ingest Network Format Masks Dialog Box	5-11

5-8	InfiniBand Network Page	5-12
5-9	Private (IB) Network Format Masks Dialog Box	5-12
5-10	Replication Network Page	5-13
5-11	Replication Network Format Masks Page	5-14
5-12	Define Clusters Page	5-14
5-13	Cluster Details	5-15
5-14	Software Locations	5-16
5-15	Disk Group Details	5-16
5-16	Initial Database	5-17
5-17	Ingest Network Fields	5-17
5-18	Replication Network Fields	5-18
5-19	Replication Network SCAN Page	5-19
5-20	Cell Alerting Page	5-19
5-21	Platinum Customer Details	5-20
5-22	Platinum Gateway Details	5-21
5-23	Platinum Agent Configuration Details	5-22
5-24	Oracle Configuration Manager Page	5-23
5-25	Auto Service Request Page	5-25
5-26	Grid Control Agent Page	5-26
5-27	Tape Library Connectivity Page	5-26
7-1	Recovery Appliance X7, X6, and X5 Rack Population Rules	7-8
11-1	Options of the Install Utility	11-6
12-1	Battery Life Expectancy	12-8
12-2	Alerts Indicating Poor Disk Performance	12-29
B-1	Default Users and Passwords	B-1
D-1	Recovery Appliance Repair Categories	D-1
D-2	Oracle Server LED Status Descriptions	D-3
D-3	Sun InfiniBand Switch LED Status Descriptions	D-5
D-4	Replacement Parts for Oracle Server X7 Oracle Database Servers	D-8
D-5	Replacement Parts for Oracle Server X6-2 Compute Servers	D-14
D-6	Replacement Parts for Oracle Server X5-2 Compute Servers	D-16
D-7	Replacement Parts for Sun Server X4-2 Compute Servers	D-19
D-8	Replacement Parts for Oracle Exadata Storage Server X7-2 High Capacity Servers	D-22
D-9	Replacement Parts for Oracle Server X6-2L Storage Servers	D-25
D-10	Replacement Parts for Oracle Server X5-2L Storage Servers	D-28
D-11	Replacement Parts for Sun Server X4-2L Storage Servers	D-30
D-12	Replacement Parts for InfiniBand Switches	D-33

D-13	Replacement Parts for Cisco Switch	D-34
D-14	Replacement Parts for Power Distribution Units	D-35
D-15	Replacement Cables for Recovery Appliance	D-36
E-1	ILOM Cabling for the Recovery Appliance X7, X6, or X5 Full Racks	E-4
E-2	Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks	E-4
E-3	Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks	E-5
E-4	Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks	E-6
E-5	InfiniBand Switch-to-Switch Cabling	E-7
E-6	InfiniBand Switch-to-Server Cabling for the Recovery Appliance X6 or X5 Full Racks	E-8
F-1	ILOM Cabling for the Recovery Appliance X4 Full Racks	F-4
F-2	Gigabit Ethernet Cabling for the Recovery Appliance X4 Full Racks	F-4
F-3	Single-Phase PDU Cabling for the Recovery Appliance X4 Full Racks	F-5
F-4	Three-Phase PDU Cabling for the Recovery Appliance X4 Full Racks	F-6
F-5	InfiniBand Switch-to-Switch Cabling	F-7
F-6	InfiniBand Switch-to-Server Cabling for the Recovery Appliance X4 Full Racks	F-7
G-1	InfiniBand Cables Provided as Spares	G-2
G-2	Cable Requirements for Multirack Installations	G-2
G-3	Leaf Switch Connections for Rack 1 in a Two-Rack System	G-7
G-4	Leaf Switch Connections for Rack 2 in a Two-Rack System	G-8
G-5	Leaf Switch Connections for Rack 1 in a Three-Rack System	G-8
G-6	Leaf Switch Connections for Rack 2 in a Three-Rack System	G-9
G-7	Leaf Switch Connections for Rack 3 in a Three-Rack System	G-9
G-8	Leaf Switch Connections for Rack 1 in a Four-Rack System	G-10
G-9	Leaf Switch Connections for Rack 2 in a Four-Rack System	G-10
G-10	Leaf Switch Connections for Rack 3 in a Four-Rack System	G-11
G-11	Leaf Switch Connections for Rack 4 in a Four-Rack System	G-11
G-12	Leaf Switch Connections for Rack 1 in a Five-Rack System	G-12
G-13	Leaf Switch Connections for Rack 2 in a Five-Rack System	G-13
G-14	Leaf Switch Connections for Rack 3 in a Five-Rack System	G-13
G-15	Leaf Switch Connections for Rack 4 in a Five-Rack System	G-14
G-16	Leaf Switch Connections for Rack 5 in a Five-Rack System	G-14
G-17	Leaf Switch Connections for Rack 1 in a Six-Rack System	G-15
G-18	Leaf Switch Connections for Rack 2 in a Six-Rack System	G-15
G-19	Leaf Switch Connections for Rack 3 in a Six-Rack System	G-16
G-20	Leaf Switch Connections for Rack 4 in a Six-Rack System	G-17
G-21	Leaf Switch Connections for Rack 5 in a Six-Rack System	G-17
G-22	Leaf Switch Connections for Rack 6 in a Six-Rack System	G-18

G-23	Leaf Switch Connections for Rack 1 in a Seven-Rack System	G-18
G-24	Leaf Switch Connections for Rack 2 in a Seven-Rack System	G-19
G-25	Leaf Switch Connections for Rack 3 in a Seven-Rack System	G-19
G-26	Leaf Switch Connections for Rack 4 in a Seven-Rack System	G-20
G-27	Leaf Switch Connections for Rack 5 in a Seven-Rack System	G-21
G-28	Leaf Switch Connections for Rack 6 in a Seven-Rack System	G-21
G-29	Leaf Switch Connections for Rack 7 in a Seven-Rack System	G-22
G-30	Leaf Switch Connections for Rack 1 in an Eight-Rack System	G-22
G-31	Leaf Switch Connections for Rack 2 in an Eight-Rack System	G-23
G-32	Leaf Switch Connections for Rack 3 in an Eight-Rack System	G-23
G-33	Leaf Switch Connections for Rack 4 in an Eight-Rack System	G-24
G-34	Leaf Switch Connections for Rack 5 in a Eight-Rack System	G-25
G-35	Leaf Switch Connections for Rack 6 in a Eight-Rack System	G-25
G-36	Leaf Switch Connections for Rack 7 in a Eight-Rack System	G-26
G-37	Leaf Switch Connections for Rack 8 in a Eight-Rack System	G-26
G-38	Leaf Switch Connections for Rack 9	G-27
G-39	Leaf Switch Connections for Rack 10	G-28
G-40	Leaf Switch Connections for Rack 11	G-28
G-41	Leaf Switch Connections for Rack 12	G-29
G-42	Leaf Switch Connections for Rack 13	G-29
G-43	Leaf Switch Connections for Rack 14	G-30
G-44	Leaf Switch Connections for Rack 15	G-31
G-45	Leaf Switch Connections for Rack 16	G-31
G-46	Leaf Switch Connections for Rack 17	G-32
G-47	Leaf Switch Connections for Rack 18	G-33

Preface

Welcome to *Zero Data Loss Recovery Appliance Owner's Guide*.

This preface contains the following topics:

- [Audience](#) (page xxii)
- [Documentation Accessibility](#) (page xxii)
- [Related Documentation](#) (page xxii)
- [Conventions](#) (page xxiii)

Audience

This guide is intended for customers and those responsible for data center site planning, configuration, and maintenance of Zero Data Loss Recovery Appliance, commonly known as Recovery Appliance.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documentation

The following guides contain detailed information about the components of a Recovery Appliance rack:

- *Oracle Server X5-2 Service Manual* at http://docs.oracle.com/cd/E41059_01/index.html
- *Oracle Server X5-2L Service Manual* at http://docs.oracle.com/cd/E41033_01/index.html
- *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E36265_01/index.html

- *Sun Flash Accelerator F80 PCIe Card User's Guide* at http://docs.oracle.com/cd/E41278_01/html/E41251/index.html
- Oracle Integrated Lights Out Manager (ILOM) 3.0 documentation at <http://docs.oracle.com/cd/E19860-01/index.html>

The My Oracle Support website provides the following notes related to Recovery Appliance:

- Information Center: Zero Data Loss Recovery Appliance (Doc ID 1683791.2) at https://support.oracle.com/epmos/faces/DocumentDisplay?_afzLoop=235867602243456&id=1683791.2&_adf.ctrl-state=yhtsmq4wf_308
- Zero Data Loss Recovery Appliance Supported Versions (Doc ID 1927416.1) at https://support.oracle.com/epmos/faces/DocumentDisplay?_afzLoop=231507746454077&id=1927416.1&_adf.ctrl-state=15vsit3rr_102
- Oracle Exadata Database Machine exachk or HealthCheck (Doc ID 1070954.1) at https://support.oracle.com/epmos/faces/DocumentDisplay?id=1070954.1&_adf.ctrl-state=1ddvhqlhg_278&_afzLoop=3415586249981774
- Steps to shut down or reboot an Exadata storage cell without affecting ASM (Doc ID 1188080.1) at https://support.oracle.com/epmos/faces/DocumentDisplay?_afzLoop=3417278942550202&id=1188080.1&_adf.ctrl-state=1ddvhqlhg_1009
- Cluster Verification Utility (CLUVFY) FAQ (Doc ID 316817.1) at https://support.oracle.com/epmos/faces/DocumentDisplay?_afzLoop=3416334547680298&id=316817.1&_adf.ctrl-state=1ddvhqlhg_586
- Oracle Sun Database Machine Setup/Configuration Best Practices (Doc ID 1274318.1) at https://support.oracle.com/epmos/faces/DocumentDisplay?_afzLoop=3416846305484690&id=1274318.1&_adf.ctrl-state=1ddvhqlhg_956
- Oracle Auto Service Request (ASR) (Doc ID 1185493.1) at https://support.oracle.com/epmos/faces/DocumentDisplay?_afzLoop=3415885796905099&id=1185493.1&_adf.ctrl-state=1ddvhqlhg_455

Conventions

The following text conventions are used in this document:

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

Convention	Meaning
# prompt	The pound (#) prompt indicates a command that is run as the <code>root</code> user.

Part I

Preinstallation

Part I provides an introduction to Zero Data Loss Recovery Appliance. It describes the preinstallation procedures that you must complete before receiving shipment, to prevent delays with the installation.

Part I contains the following chapters:

- [Introduction to Recovery Appliance](#) (page 1-1)
- [Site Requirements for Recovery Appliance](#) (page 2-1)
- [Preparing Your Networks for Recovery Appliance](#) (page 3-1)
- [Setting Up Auto Service Request](#) (page 4-1)
- [Using the Deployment Assistant](#) (page 5-1)

1

Introduction to Recovery Appliance

This chapter describes the features and hardware components of Oracle's Zero Data Loss Recovery Appliance (Recovery Appliance).

This chapter contains these sections:

- [About Recovery Appliance](#) (page 1-1)
- [Recovery Appliance Rack Components](#) (page 1-2)
- [Compute Server Components](#) (page 1-3)
- [Storage Server Components](#) (page 1-5)
- [Spares Kit Components for Recovery Appliance](#) (page 1-7)
- [About Tape Backup Infrastructure](#) (page 1-8)
- [Restrictions on Modifying Recovery Appliance](#) (page 1-9)

1.1 About Recovery Appliance

Recovery Appliance is a backup and recovery solution that is specifically engineered to protect your Oracle Database data. It offers a massively scalable repository that records all delta changes (that is, differences) across hundreds to thousands of databases.

The Database Delta Push technology enables Oracle DBAs to run one full Recovery Manager (RMAN) backup, and thereafter run only incremental backups, from each database to Recovery Appliance. There are no recurring full backups, tape backups, or RMAN backup deletion, validation, and maintenance commands.

In the event of a disaster, DBAs can recover any database to any point in time or within subseconds of the current time. The Database Delta Store virtual full backup technology constructs physical backups from the resident incremental blocks, in response to RMAN database restore requests.

You can easily and quickly replicate backups on a local Recovery Appliance over a secure network to a remote Recovery Appliance for protection against server or site outage. Only the changed blocks are replicated. Virtual full backups on the remote Recovery Appliance are available for database restore operations.

Autonomous tape archives further protect against site disasters. Recovery Appliance automates the copying of virtual full or incremental backups to tape, for long-term retention and archival purposes. Recovery Appliance is preinstalled with tape backup software and support for an optional 16 Gb fiber channel card on each compute server, to connect directly to the tape hardware.

You can manage this entire infrastructure with end-to-end visibility into the status of data protection to ensure that data can be recovered successfully at any time. The primary administrative tools are Oracle Enterprise Manager Cloud Control and the RMAN command interface.

1.2 Recovery Appliance Rack Components

Recovery Appliance consists of compute servers, storage servers, and the network components to connect to your network. Oracle offers a flexible configuration for Recovery Appliance, so that you can acquire the resources that you need, and expand incrementally when you need more. You can configure up to 18 full racks into a single cluster.

[Table 1-2](#) (page 1-2) lists the components of a Recovery Appliance X7 rack.

Table 1-1 Components of a Recovery Appliance X7 Rack

Quantity	Component
2	Oracle Server X7-2 compute servers
3 to 18	Oracle Server X7-2L storage servers
2	Sun Datacenter InfiniBand Switch 36
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco Catalyst 4948E-F, model number WS-C4948E-F-S Ethernet switch
2	32 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports. Optional with OSB tape.

[Table 1-2](#) (page 1-2) lists the components of a Recovery Appliance X6 rack.

Table 1-2 Components of a Recovery Appliance X6 Rack

Quantity	Component
2	Oracle Server X6-2 compute servers
3 to 18	Oracle Server X6-2L storage servers
2	Sun Datacenter InfiniBand Switch 36
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco Catalyst 4948E-F, model number WS-C4948E-F-S Ethernet switch
2	16 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports

[Table 1-3](#) (page 1-2) lists the components of a Recovery Appliance X5 rack.

Table 1-3 Components of a Recovery Appliance X5 Rack

Quantity	Component
2	Oracle Server X5-2 compute servers
3 to 18	Oracle Server X5-2L storage servers
2	Sun Datacenter InfiniBand Switch 36
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)

Table 1-3 (Cont.) Components of a Recovery Appliance X5 Rack

Quantity	Component
1	48-port Cisco Catalyst 4948E-F, model number WS-C4948E-F-S Ethernet switch
2	16 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports

Table 1-4 (page 1-3) lists the components of a Recovery Appliance X4 rack.

Table 1-4 Components of a Recovery Appliance X4 Rack

Quantity	Component
2	Sun Server X4-2 compute servers
3 to 14	Sun Server X4-2L storage servers
2	Sun Datacenter InfiniBand Switch 36
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco Catalyst 4948E-F, model number WS-C4948E-F-S Ethernet switch
2	16 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports

1.3 Compute Server Components

Table 1-5 (page 1-3) lists the components of an individual Oracle X7-2 compute server.

Table 1-5 Oracle X7-2 Compute Server Components

Quantity	Description
2	24-core Xeon 8160 processors (2.1 GHz)
8	384 GB RAM
2	QDR InfiniBand (40 Gb/s) ports
1	Oracle Integrated Lights Out Manager (ILOM) for remote management. Ethernet Port
1	1 GB copper Ethernet Port (mgmt)
2	On-board 10 Gigabit copper Base-T Ethernet ports
2	On-board 10/25 Gigabit optical Ethernet ports
2	Optional PCIe card 10/25 Gigabit optical Ethernet ports Sun Dual 10/25 Gb Ethernet SFP+/28 PCIe 2.0 Low Profile Adapter incorporating Intel 82599 10/25 Gb Ethernet controller and supporting pluggable SFP+ Transceivers. ROHS-5
2	32 Gb Fibre Channel Ports Sun Storage Dual 32 Gb Fibre Channel PCIe Universal HBA, QLogic for tape connectivity (optional)
2	Redundant Hot-Swappable Power Supplies

Table 1-5 (Cont.) Oracle X7-2 Compute Server Components

Quantity	Description
2	Redundant Hot-Swappable Fans
1	Disk Controller HBA with 1 GB Supercap-backed Write Cache

[Table 1-6](#) (page 1-4) lists the components of an Oracle X6-2 compute server.

Table 1-6 Oracle Server X6-2 Components

Quantity	Description
2	22-core Intel Xeon E5-2699 v4 processors (2.2 GHz)
8	32 GB RAM (256 GB total)
4	600 GB 10K RPM SAS disks
1	Disk controller HBA with 1 GB supercap-backed write cache
2	InfiniBand 4X QDR (40 Gb/s) ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	1 GbE/10GbE Base-T Ethernet ports
2	10 GbE Ethernet SFP+ ports (1 dual-port 10GbE PCIe 2.0 network card based on the Intel 82599 10 GbE controller technology)
1	Ethernet port for Oracle Integrated Lights Out Manager (ILOM) for remote management
1	16 Gb dual-port QLogic PCIe 3.0 HBA card (optional)

[Table 1-7](#) (page 1-4) lists the components of an Oracle X5-2 compute server.

Table 1-7 Oracle Server X5-2 Components

Quantity	Description
2	18-core Intel Xeon E5-2699 v3 processors (2.3 GHz)
8	32 GB RAM (256 GB total)
4	600 GB 10K RPM SAS disks
1	Disk controller HBA with 1 GB supercap-backed write cache
2	InfiniBand 4X QDR (40 Gb/s) ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	1 GbE/10GbE Base-T Ethernet ports
2	10 GbE Ethernet SFP+ ports (1 dual-port 10GbE PCIe 2.0 network card based on the Intel 82599 10 GbE controller technology)
1	Ethernet port for Oracle Integrated Lights Out Manager (ILOM) for remote management
1	16 Gb dual-port QLogic PCIe 3.0 HBA card (optional)

[Table 1-8](#) (page 1-5) lists the components of a Sun X4-2 compute server.

Table 1-8 Sun Server X4-2 Components

Quantity	Description
2	12-Core Intel Xeon E5-2697 v2 processors (2.7 GHz)
16	16 GB RAM (256 total)
4	600 GB 10K RPM SAS disks
1	Disk controller HBA with 512 MB battery-backed write cache, and swappable battery backup unit (BBU)
2	InfiniBand 4X QDR (40 Gb/s) ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	1 GbE/10GbE Base-T Ethernet ports
2	10 GbE Ethernet SFP+ ports (1 dual-port 10 GbE PCIe 2.0 network card based on the Intel 82599 10 GbE controller technology)
1	Ethernet port for Oracle Integrated Lights Out Manager (ILOM) for remote management
1	16 Gb dual-port QLogic PCIe 3.0 HBA card (optional)

1.4 Storage Server Components

[Table 1-9](#) (page 1-5) lists the components of an Oracle X7–2L storage server.

Table 1-9 Oracle Server X7–2L Storage Components

Quantity	Description
2	10-core Xeon 4114 processors (2.0 GHz)
8	192 GB RAM
12	10 TB 7.2 K RPM High Capacity SAS disks
2	3.2 TB non-volatile memory express (NVMe) PCIe 3.0 flash accelerator cards
1	Disk controller HBA with 1 GB supercap-backed write cache
2	36 port QDR InfiniBand (40 Gb/s) ports
2	Embedded Gigabit Ethernet ports
1	Ethernet port for ILOM for remote management
2	Redundant Hot-Swappable Power Supplies
2	Redundant Hot-Swappable Fans

[Table 1-10](#) (page 1-5) lists the components of an Oracle X6-2L storage server.

Table 1-10 Oracle X6-2L Storage Server Components

Quantity	Description
2	10-core Intel Xeon CPU E5-2630 v4 (2.2 GHz)
8	16 GB RAM (128 GB total RAM)

Table 1-10 (Cont.) Oracle X6-2L Storage Server Components

Quantity	Description
12	8 TB 7.2 K RPM High Capacity SAS disks
4	3.2 TB non-volatile memory express (NVMe) PCIe 3.0 flash accelerator cards
1	Disk controller HBA with 1 GB supercap-backed write cache
2	InfiniBand 4X QDR (40 Gb/s) ports (PCIe 3.0), both ports active
4	Embedded Gigabit Ethernet ports
1	Ethernet port for ILOM for remote management

[Table 1-11](#) (page 1-6) lists the components of an Oracle Server X5-2L storage server.

Table 1-11 Oracle Server X5-2L Components

Quantity	Description
2	8-core Intel Xeon CPU E5-2630 v3 (2.4 GHz)
4	8 GB RAM
4	16 GB RAM (96 GB total RAM)
12	4 TB 7.2 K RPM High Capacity SAS disks
4	1.6 TB non-volatile memory express (NVMe) PCIe 3.0 flash accelerator cards
1	Disk controller HBA with 1 GB supercap-backed write cache
2	InfiniBand 4X QDR (40 Gb/s) ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	Embedded Gigabit Ethernet ports
1	Ethernet port for ILOM for remote management

[Table 1-12](#) (page 1-6) lists the components of a Sun Server X4-2L storage server.

Table 1-12 Sun Server X4-2L Components

Quantity	Description
2	Six-core Intel Xeon E5-2630 v2 processors (2.6 GHz)
4	8 GB RAM
4	16 GB RAM (96 GB total RAM)
12	4 TB 7.2K RPM High Capacity SAS disks
4	800 GB Sun Flash Accelerator F80 PCIe Cards
1	Disk controller HBA with 512 MB battery-backed write cache and swappable BBU
2	InfiniBand 4 X QDR (40 Gb/s) InfiniBand ports (1 dual-port PCIe 3.0 Host Channel Adapter (HCA))
4	Embedded Gigabit Ethernet ports

Table 1-12 (Cont.) Sun Server X4-2L Components

Quantity	Description
1	Ethernet port for ILOM for remote management

1.5 Spares Kit Components for Recovery Appliance

The spares kit for Recovery Appliance X7 has the following components:

- 1 x 10 TB 7.2K RPM high capacity SAS disk
- 1 x flash accelerator F320 PCIe card
- Tools for the rack
- Cisco accessory adapter and documentation ship kit
- Server documentation

The spares kit for Recovery Appliance X6 has the following components:

- 1 x 8 TB 7.2K RPM high capacity SAS disk
- 1 x flash accelerator F320 PCIe card for high capacity storage servers only
- Tools for the rack
- Cisco accessory adapter and documentation ship kit
- Server documentation

The spares kit for Recovery Appliance X5 has the following components:

- 1 x 4 TB 7.2K RPM high capacity SAS disk
- 1 x flash accelerator F160 PCIe card for high capacity storage servers only
- Tools for the rack
- Cisco accessory adapter and documentation ship kit
- Server documentation

The spares kit for Recovery Appliance X4 has the following components:

- 1 x 4 TB 7.2K RPM high capacity SAS disk
- 1 x 800 GB Exadata Smart Flash Cache card
- Tools for the rack
- Cisco accessory adapter and documentation ship kit
- Server documentation

These spare cables are tied inside the Recovery Appliance rack:

- 1 or 2 (depending on rack size) 3m InfiniBand cables
- 1 blue Ethernet cable
- 1 red Ethernet cable
- 1 black Ethernet cable

Two sets of two keys are provided for the rack door and side panel locks.

1.6 About Tape Backup Infrastructure

Tape backup is an optional feature. You can choose between Oracle and third-party tape systems.

Oracle's Recovery Appliance supports the installation and configuration of Oracle tape systems. You can connect Recovery Appliance to the storage area network (SAN) in your data center for backups to an Oracle tape system.

- [Oracle Recommended Stack](#) (page 1-8)
- [Third-Party Tape Systems](#) (page 1-8)

1.6.1 Oracle Recommended Stack

Recovery Appliance is engineered to work with industry leading tape infrastructure components, including those described in [Table 1-13](#) (page 1-8).

Table 1-13 Oracle Tape System Components

Component	Models	Notes
Oracle Tape Libraries	SL150, SL3000, SL8500	ACSLs, ACSLS Logical Libraries, FC-SCSI
Oracle Tape Drives	T10000D, T10000C, LTO	All interfaces supported by the recommended Oracle libraries
Automated Cartridge System Library Software (ACSLs)	Automated Cartridge System Library Software (ACSLs) 8.x	API, Logical Libraries
Tape Encryption	Oracle Key Manager (OKM) 2.3	Recommended Key Management for enterprise key management. Operates seamlessly with Recovery Appliance

Recovery Appliance supports Oracle Secure Backup tape management software out of the box. It provides the following benefits:

- Oracle Secure Backup is deployed and configured as part of the Recovery Appliance software installation.
- Upgrades are included with Recovery Appliance software upgrades.
- Tape backups automatically protect the Recovery Appliance metadata.
- Oracle Secure Backup provides disaster recovery automation from tape.
- Oracle Secure Backup copies data from its storage location directly to tape.
- Oracle Secure Backup provides superior local tape performance by using only the InfiniBand and fiber channel networks. Tape backups do not impact your IP network.

1.6.2 Third-Party Tape Systems

Out of the box, Recovery Appliance supports the use of any third-party (non Oracle) media management software that is compatible with the Recovery Appliance.

You can install the *client* software component of a third-party media management product, such as Symantec NetBackup or HP Data Protector, on the compute servers. You can then configure Recovery Appliance as a client over the 10 GbE network.

Under the terms of the Recovery Appliance license, you may not install the media manager *server* component because only Oracle Secure Backup can back up Recovery Appliance to the tape devices directly (or via SAN) attached to the fibre channel interfaces of the compute nodes. Any non Oracle media manager software must be configured only as a network client.

Note the following:

- Recovery Appliance software upgrades might not preserve the client software or its configuration.
- Third-party media managers do not provide Recovery Appliance metadata protection support.
- If you configure third-party media managers with Recovery Appliance, carefully validate your configuration to ensure that backups and restores work as expected and leverage all nodes in the Recovery Appliance RAC cluster.

1.7 Restrictions on Modifying Recovery Appliance

The following restrictions apply to software and hardware modifications to Recovery Appliance. Violating these restrictions can result in loss of warranty and support.

- Customers may not modify any internal software component or configuration on a Recovery Appliance compute server except as indicated in a My Oracle Support note or as explicitly directed by an Oracle support analyst. This includes actions such as modifying or adding listeners to the internal database, modifying the operating system or network configuration files, installing custom scripts, creating NFS mounts for any purpose except for polling a backup location, and creating additional database instances.
- Although Oracle discourages installation of third-party software products on the appliance, customers may install them under the conditions stated in My Oracle Support note 2014361.1.
- Customers may install the client software component of a third-party media management product in "LAN backup agent" mode on the compute servers. Customers may not install the media manager server software component, because only Oracle Secure Backup supports Recovery Appliance backups to tape over fibre channel connections.
- Customers may not load additional software or alter the installed software on a storage server.
- Customers may not install Oracle Enterprise Manager Ops Center agents. However, Oracle Enterprise Manager Cloud Control Management agents can be installed on the compute servers to monitor the system, as part of the Recovery Appliance software installation process.
- Recovery Appliance hardware cannot be modified or customized. There are two exceptions. The allowed hardware modifications to Recovery Appliance are as follows:
 - Modification to the administrative 48-port Cisco Gigabit Ethernet switch included with Recovery Appliance. Customers may choose to do the following:

- * Replace the Gigabit Ethernet switch, at their own expense, with an equivalent 1U 48-port Gigabit Ethernet switch that conforms to their internal data center network standards. This replacement must be performed by the customer, at their expense and labor, after delivery of Recovery Appliance. If the customer chooses to make this change, then Oracle cannot make or assist with this change given the numerous possible scenarios involved, and it is not included as part of the standard installation. The customer must supply the replacement hardware, and make or arrange for this change through other means.
- * Remove the CAT5 cables connected to the Cisco Ethernet switch, and connect them to the customer's network through an external switch or patch panel. The customer must perform these changes at their expense and labor. In this case, the Cisco Ethernet switch in the rack can be turned off and disconnected from the data center network.
- Addition of client access network switches in Recovery Appliance. The following restrictions apply for the network switches:
 - * A maximum of two client access network switches may be installed in the rack.
 - * The switch must be 1 rack unit (RU) in height.
 - * Air flow must go from the front of the rack to the back of the rack.
 - * The switch must be installed in rack slot U41 or U42.
 - * Power consumption must be less than 400 watts for each switch.
 - * Power over Ethernet (PoE) functionality, if applicable, should not be used on these switches.
- Customers can update the firmware of the other components.
 - Customers can update the IOS and firmware versions on the Cisco Gigabit Ethernet switch to meet their data center requirements.
 - Customers can update the firmware of the components of the compute servers only by applying a patch bundle provided by Oracle.
 - Customers can update the firmware of the InfiniBand switches provided they comply with the validated versions documented in My Oracle Support note 1927416.1, and its related notes.
- Customers cannot connect USB devices that draw more than 100 mA of power to the compute servers.

2

Site Requirements for Recovery Appliance

This chapter describes the site requirements for Recovery Appliance. It contains these sections:

- [General Environmental Requirements](#) (page 2-1)
- [Space Requirements](#) (page 2-5)
- [Flooring Requirements](#) (page 2-7)
- [Electrical Power Requirements](#) (page 2-8)
- [Temperature and Humidity Requirements](#) (page 2-13)
- [Ventilation and Cooling Requirements](#) (page 2-14)
- [Network Connection Requirements for Recovery Appliance](#) (page 2-15)
- [Ensuring That the Site Is Ready](#) (page 2-16)

 **See Also:**

Zero Data Loss Recovery Appliance Site Checklists

2.1 General Environmental Requirements

The environmental requirements for Recovery Appliance depend on the type and size of the system. The following tables provide an overview of the rack and disk requirements. Measurements are approximate.

[Table 2-1](#) (page 2-1) identifies the general environmental requirements for a Recovery Appliance X7 rack.

Table 2-1 Environmental Requirements for Recovery Appliance X7 Racks

Environmental Component	Full Rack	Minimum Configuration ¹
Weight See Also: " Flooring Requirements (page 2-7)"	899.9 kg (1983.9 lbs)	385.7 kg (850.4 lbs)
Power See Also: " Electrical Power Requirements (page 2-8)"	Maximum: 14.1 kVA (13.8 kW) Typical: 9.8 kVA (9.6 kW)	Maximum: 4.4 kVA (4.3 kW) Typical: 3.1 kVA (3.0 kW)

Table 2-1 (Cont.) Environmental Requirements for Recovery Appliance X7 Racks

Environmental Component	Full Rack	Minimum Configuration ¹
Cooling See Also: " Temperature and Humidity Requirements (page 2-13)", and " Ventilation and Cooling Requirements (page 2-14)"	Maximum: 47,016 BTU/hour (49,602 kJ/hour) Typical: 32,911 BTU/hour (34,7217 kJ/hour)	Maximum: 14,822 BTU/hour (15,638 kJ/hour) Typical: 10,376 BTU/hour (10,946 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: " Temperature and Humidity Requirements (page 2-13)", and " Ventilation and Cooling Requirements (page 2-14)"	Maximum: 2,177 CFM Typical: 1,524 CFM	Maximum: 686 CFM Typical: 480 CFM
Acoustic levels	9.5 Bel	8.9 Bel
IP addresses See Also: " Factory IP Address Settings (page 3-10)"	69 for Ethernet network (single cluster) 36 maximum for InfiniBand network	22 for Ethernet network (single cluster) 8 maximum for InfiniBand network
Network drops See Also: " Overview of Network Requirements (page 3-1)"	11 network drops minimum	5 network drops minimum
Fiber san drops (optional) See Also: " Overview of Network Requirements (page 3-1)"	4 x OM3 fiber cables	4 x OM3 fiber cables
External connectivity See Also: " Network Connection Requirements for Recovery Appliance (page 2-15)"	6 x 1 Gbps Ethernet ports 4 x 10 Gbps Ethernet ports 12 or more InfiniBand ports 4 x 32 Gb fiber channel SAN connections (optional)	6 x 1 Gbps Ethernet ports 4 x 10 Gbps Ethernet ports 12 or more InfiniBand ports 4 x 32 Gb fiber channel SAN connections (optional)

¹ Recovery Appliance rack with the minimum number (3) of storage servers

[Table 2-2](#) (page 2-2) identifies the general environmental requirements for a Recovery Appliance X6 rack.

Table 2-2 Environmental Requirements for Recovery Appliance X6 Racks

Environmental Component	Full Rack	Minimum Configuration ¹
Weight See Also: " Flooring Requirements (page 2-7)"	875.7 kg (1930.5 lbs)	381.2 kg (840.5 lbs)

Table 2-2 (Cont.) Environmental Requirements for Recovery Appliance X6 Racks

Environmental Component	Full Rack	Minimum Configuration ¹
Power See Also: "Electrical Power Requirements (page 2-8)"	Maximum: 13.1 kVA (12.8 kW) Typical: 9.2 kVA (9.0 kW)	Maximum: 4.1 kVA (4.0 kW) Typical: 2.9 kVA (2.8 kW)
Cooling See Also: "Temperature and Humidity Requirements (page 2-13)", and "Ventilation and Cooling Requirements (page 2-14)"	Maximum: 43,774 BTU/hour (46,182 kJ/hour) Typical: 30,642 BTU/hour (32,327 kJ/hour)	Maximum: 13,679 BTU/hour (14,432 kJ/hour) Typical: 9,575 BTU/hour (10,102 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements (page 2-13)", and "Ventilation and Cooling Requirements (page 2-14)"	Maximum: 2,027 CFM Typical: 1,419 CFM	Maximum: 633 CFM Typical: 443 CFM
Acoustic levels	8.5 Bel	8.3 Bel
IP addresses See Also: "Factory IP Address Settings (page 3-10)"	69 for Ethernet network (single cluster) 36 maximum for InfiniBand network	22 for Ethernet network (single cluster) 8 maximum for InfiniBand network
Network drops See Also: "Overview of Network Requirements (page 3-1)"	11 network drops minimum	5 network drops minimum
Fiber san drops (optional) See Also: "Overview of Network Requirements (page 3-1)"	4 x OM3 fiber cables	4 x OM3 fiber cables
External connectivity See Also: "Network Connection Requirements for Recovery Appliance (page 2-15)"	6 x 1 Gbps Ethernet ports 4 x 10 Gbps Ethernet ports 12 or more InfiniBand ports 4 x 16 Gb fiber channel SAN connections (optional)	6 x 1 Gbps Ethernet ports 4 x 10 Gbps Ethernet ports 12 or more InfiniBand ports 4 x 16 Gb fiber channel SAN connections (optional)

¹ Recovery Appliance rack with the minimum number (3) of storage servers

Table 2-3 (page 2-4) identifies the general environmental requirements for a Recovery Appliance X5 rack.

Table 2-3 Environmental Requirements for Recovery Appliance X5 Racks

Environmental Component	Full Rack	Minimum Configuration ¹
Weight See Also: "Flooring Requirements (page 2-7)"	875.7 kg (1930.5 lbs)	381.2 kg (840.5 lbs)
Power See Also: "Electrical Power Requirements (page 2-8)"	Maximum: 12.0 kW Typical: 8.4 kW	Maximum: 4.0 kW Typical: 2.8 kW
Cooling See Also: "Temperature and Humidity Requirements (page 2-13)", and "Ventilation and Cooling Requirements (page 2-14)"	Maximum: 40,942 BTU/hour (43,194 kJ/hour) Typical: 28,660 BTU/hour (30,236 kJ/hour)	Maximum: 13,560 BTU/hour (14,306 kJ/hour) Typical: 9,492 BTU/hour (10,014 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements (page 2-13)", and "Ventilation and Cooling Requirements (page 2-14)"	Maximum: 1,895 CFM Typical: 1,327 CFM	Maximum: 628 CFM Typical: 439 CFM
Acoustic levels	8.5 Bel	8.3 Bel
IP addresses See Also: "Factory IP Address Settings (page 3-10)"	69 for Ethernet network (single cluster) 36 maximum for InfiniBand network	22 for Ethernet network (single cluster) 8 maximum for InfiniBand network
Network drops See Also: "Overview of Network Requirements (page 3-1)"	11 network drops minimum	5 network drops minimum
Fiber san drops (optional) See Also: "Overview of Network Requirements (page 3-1)"	4 x OM3 fiber cables	4 x OM3 fiber cables
External connectivity See Also: "Network Connection Requirements for Recovery Appliance (page 2-15)"	6 x 1 Gbps Ethernet ports 4 x 10 Gbps Ethernet ports 12 or more InfiniBand ports 4 x 16 Gb fiber channel SAN connections (optional)	6 x 1 Gbps Ethernet ports 4 x 10 Gbps Ethernet ports 12 or more InfiniBand ports 4 x 16 Gb fiber channel SAN connections (optional)

¹ Recovery Appliance rack with the minimum number (3) of storage servers

Table 2-4 (page 2-5) identifies the general environmental requirements for a Recovery Appliance X4 rack.

Table 2-4 Environmental Requirements for Recovery Appliance X4 Racks

Environmental Component	Full Rack	Minimum Configuration ¹
Weight See Also: "Flooring Requirements (page 2-7)"	724.9 kg (1598 lb)	393.3 kg (867 lb)
Power See Also: "Electrical Power Requirements (page 2-8)"	Maximum: 8.1 kVA (8.0 kW) Typical: 5.7 kVA ² (5.6 kW)	Maximum: 3.2 kVA (3.1 kW) Typical: 2.3 kVA ² (2.2 kW)
Cooling See Also: "Temperature and Humidity Requirements (page 2-13)", and "Ventilation and Cooling Requirements (page 2-14)"	Maximum: 27,200 BTU/hour (28,700 kJ/hour) Typical: 19,000 BTU/hour (20,100 kJ/hour)	Maximum: 10,600 BTU/hour (11,200 kJ/hour) Typical: 7,500 BTU/hour (7,900 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements (page 2-13)", and "Ventilation and Cooling Requirements (page 2-14)"	Maximum: 1,260 CFM Typical: 880 CFM	Maximum: 490 CFM Typical: 350 CFM
Acoustic levels	9.1 Bel	8.5 Bel
IP addresses See Also: "Factory IP Address Settings (page 3-10)"	69 for Ethernet network (single cluster) 36 maximum for InfiniBand network	22 for Ethernet network (single cluster) 8 maximum for InfiniBand network
Network drops See Also: "Overview of Network Requirements (page 3-1)"	11 network drops minimum	5 network drops minimum
Fiber san drops (optional) See Also: "Overview of Network Requirements (page 3-1)"	4 x OM3 fiber cables	4 x OM3 fiber cables
External connectivity See Also: "Network Connection Requirements for Recovery Appliance (page 2-15)"	6 x 1 Gbps Ethernet ports 4 x 10 Gbps Ethernet ports 12 or more InfiniBand ports 4 x 16 Gb fiber channel SAN connections (optional)	6 x 1 Gbps Ethernet ports 4 x 10 Gbps Ethernet ports 12 or more InfiniBand ports 4 x 16 Gb fiber channel SAN connections (optional)

¹ Recovery Appliance rack with the minimum number (3) of storage servers
² Varies based on application load

2.2 Space Requirements

All Recovery Appliance hardware configurations use the same 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 2300 mm (90 inches), measured from the true floor or raised floor, whichever is higher. An additional 914 mm (36 inches) is required above the rack height in the front and rear aisle space for maintenance access. 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.

2.2.1 Receiving and Unpacking Requirements

Before Recovery Appliance arrives, ensure that the receiving area is large enough for the package.

[Table 2-5](#) (page 2-6) lists the dimensions and weight of the shipment.

Table 2-5 Shipping Measurements

Measurement	Full Rack	Minimum Configuration
Height	2159 mm (85 inches)	2159 mm (85 inches)
Width	1219 mm (48 inches)	1219 mm (48 inches)
Depth	1575 mm (62 inches)	1575 mm (62 inches)
X7 Weight	899.9 kg (1983.9 lbs)	385.7 kg (850.4 lbs)
X6 and X5 Weight	875.7 kg (1930.5 lbs)	381.2 kg (840.5 lbs)
X4 Weight	866 kg (1909 lb)	469 kg (1034 lb)

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.

Leave Recovery Appliance in its shipping container until it arrives at the installation site. Ensure sufficient clearance and clear pathways for moving it from the unpacking location to the installation location. The entire access route to the installation site should be free of raised-pattern flooring that can cause vibration.

Use a conditioned space when removing the packaging material, to reduce particles before entering the data center. Allow enough space for unpacking Recovery Appliance from its shipping cartons. [Table 2-6](#) (page 2-6) lists the access route requirements.

Table 2-6 Access Route Requirements

Access Route Item	With Shipping Pallet	Without Shipping Pallet
Minimum door height	218.4 cm (86 inches)	204 cm (80.5 inches)
Minimum door width	127 cm (50 inches)	64 cm (25.19 inches)
Minimum elevator depth	162.6 cm (64 inches)	124 cm (49 inches)

Table 2-6 (Cont.) Access Route Requirements

Access Route Item	With Shipping Pallet	Without Shipping Pallet
Maximum incline	6 degrees	6 degrees
Minimum elevator, pallet jack, and floor loading capacity	1134 kg (2500 lb)	1134 kg (2500 lb)

2.2.2 Maintenance Access Requirements

The maintenance area must be large enough for Recovery Appliance and have the required access space. For example, the required space to remove the side panels is 67.6 cm (26.6 inches). [Table 2-7](#) (page 2-7) lists the maintenance access requirements.

Table 2-7 Maintenance Access Requirements for Recovery Appliance

Location	Maintenance Access Requirement
Rear maintenance	91.4 cm (36 inches)
Front maintenance	123 cm (48.5 inches)
Top maintenance	91.4 cm (36 inches)

Note:

- Open tiles are required for electrical access.
- The Recovery Appliance racks can be placed side-by-side when interconnecting them.

2.3 Flooring Requirements

Oracle recommends that the Recovery Appliance racks be installed on raised flooring. The site floor and the raised flooring must be able to support the total weight of the rack.

[Table 2-8](#) (page 2-7) lists the floor load requirements.

Table 2-8 Floor Load Requirements for Recovery Appliance

Description	Requirement
Maximum allowable weight of installed rack equipment	952.5 kg (2100 lb)
Maximum allowable weight of installed power distribution units	52 kg (115 lb)
Maximum dynamic load (maximum allowable weight of installed equipment including PDUs)	1005 kg (2215 lb)

2.4 Electrical Power Requirements

Recovery Appliance can operate effectively over a wide range of voltages and frequencies. However, they must have a reliable power source. Damage may occur if the ranges are exceeded. Electrical disturbances such as the following might damage Recovery Appliance:

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

To protect Recovery Appliance from such disturbances, you should have a dedicated power distribution system, power-conditioning equipment, and either lightning arresters or power cables to protect from electrical storms.

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

2.4.1 PDUs for Recovery Appliance

The type of PDUs depends on the geographic location of the data center. The following tables lists the requirements for Recovery Appliance:

- [Single-Phase PDUs for North America, South America, Japan, and Taiwan](#) (page 2-8)
- [Three-Phase PDUs for North America, South America, Japan, and Taiwan](#) (page 2-9)
- [Single-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific](#) (page 2-10)
- [Three-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific](#) (page 2-11)

2.4.1.1 Single-Phase PDUs for North America, South America, Japan, and Taiwan

[Table 2-9](#) (page 2-8) lists the requirements for a low-voltage single-phase PDU for North America, South America, Japan, and Taiwan. There are two PDUs per rack.

Table 2-9 Low-Voltage 15 kVA Single-Phase PDUs

Option	Requirement per PDU
Voltage	200 to 240 VAC
Frequency	50/60 Hz
Current	24A maximum per input
Power rating	15 kVA
Output current	72A (3 x 24A)

Table 2-9 (Cont.) Low-Voltage 15 kVA Single-Phase PDUs

Option	Requirement per PDU
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection ¹	20A
Data center receptacle	15 kVA, with three 30A/250V 2-pole/3-wire NEMA L6-30P plugs

¹ UL489 2-pole circuit breaker

You need the following items to connect Recovery Appliance to a low-voltage single-phase power source:

- 6 power cords for two PDUs, 30A at 200 to 240 VAC
- 6 receptacles to connect the PDUs to 6 NEMA L6-30 data center receptacles

[Figure 2-1](#) (page 2-9) shows the low-voltage, single-phase PDU power connector for North America, South America, Japan and Taiwan.

Figure 2-1 Low-Voltage Single Phase Power Connector



2.4.1.2 Three-Phase PDUs for North America, South America, Japan, and Taiwan

[Table 2-10](#) (page 2-9) lists the requirements for a low-voltage three-phase PDU for North America, South America, Japan, and Taiwan. There are two PDUs per rack.

Table 2-10 Low-Voltage 15 kVA Three-Phase PDUs

Option	Requirement per PDU
Voltage	190 to 220 VAC
Frequency	50/60 Hz
Current	40A maximum per phase
Power rating	14.4 kVA
Output current	69.3A (3 x 23.1A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection ¹	20A

Table 2-10 (Cont.) Low-Voltage 15 kVA Three-Phase PDUs

Option	Requirement per PDU
Data center receptacle	15 kVA IEC 60309 60A 4-pin 250 VAC three-phase IP67

¹ UL489 2-pole circuit breaker

You need the following items to connect Recovery Appliance to a low-voltage three-phase power source:

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

Figure 2-2 (page 2-10) shows the low-voltage three-phase PDU power connector for North America, South America, Japan and Taiwan.

Figure 2-2 Low-Voltage Three-Phase Power Connector



2.4.1.3 Single-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific

Table 2-11 (page 2-10) lists the requirements for a high-voltage single-phase PDU for Europe, the Middle East, Africa, and Asia Pacific, except Japan and Taiwan. There are two PDUs per rack.

Table 2-11 High-Voltage 15 kVA Single Phase PDUs

Option	Requirement per PDU
Voltage	220 to 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 ¹	20A
Data center receptacle	15 kVA, with three blue 32A/240V splash-proof 2-pole/3-wire IEC 60309 plugs

¹ UL489 1-pole circuit breaker

You need the following items to connect Recovery Appliance to a high-voltage single-phase power source:

- 6 power cords for two PDUs, 25A at 220/380 to 240/415 VAC 1-phase voltage
- 6 receptacles to connect the PDUs to 2 IEC 60309 32A 3-pin 250 VAC IP44 data center receptacles, and that require 22 kVA, 1-phase

Figure 2-3 (page 2-11) shows the high-voltage single-phase PDU power connector for EMEA, and APAC, except for Japan and Taiwan.

Figure 2-3 High-Voltage Single Phase Power Connector



2.4.1.4 Three-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific

Table 2-12 (page 2-11) lists the requirements for a high-voltage, three-phase PDU for Europe, the Middle East, Africa, and Asia Pacific, except Japan and Taiwan. There are two PDUs per rack.

Table 2-12 High-Voltage 15 kVA Three-Phase PDUs

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

¹ UL489 1-pole circuit breaker

You need the following items to connect Recovery Appliance to a high-voltage three-phase power source:

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

Figure 2-4 (page 2-12) shows the high-voltage three-phase PDU power connector for EMEA and APAC, except Japan and Taiwan.

Figure 2-4 High-Voltage Three-Phase Power Connector



2.4.2 Facility Power Requirements

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

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 and Canada, ensure that the overall system AC input current load does not exceed 80 percent of the branch circuit AC current rating.

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.

2.4.3 Circuit Breaker Requirements

To prevent catastrophic failures, the design of your power system must ensure that adequate power is provided to the servers. Use dedicated AC breaker panels for all power circuits that supply power to the server. Electrical work and installations must obey applicable local, state, and national electrical codes. Servers require that electrical circuits be grounded to the earth.

In addition to the circuit breakers, you must provide a stable power source, such as an uninterruptible power supply (UPS) to reduce the possibility of component failures. If computer equipment is subjected to repeated power interruptions and fluctuations, then it is susceptible to a higher rate of component failure.



Note:

The owner supplies the circuit breakers. Each power cord requires a circuit breaker.

2.4.4 Grounding Guidelines

The Recovery Appliance cabinets are shipped with grounding-type power cords. Always connect the cords to grounded power outlets. Because different grounding methods are used depending on location, check the grounding type and refer to documentation, such as IEC documents, for the correct grounding method. Ensure that the facility administrator or a qualified electrical engineer verifies the grounding method for the building and performs the grounding work.

2.5 Temperature and Humidity Requirements

Airflow through the Recovery Appliance rack is from front to back. Refer to [Table 2-4](#) (page 2-5) for information on cooling and airflow.

 **Note:**

Studies show 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.

Excessive internal temperatures may result in a full or partial shut down of Recovery Appliance.

[Table 2-13](#) (page 2-13) lists the temperature, humidity, and altitude requirements for operating and nonoperating Recovery Appliances in the data center.

Table 2-13 Temperature, Humidity and Altitude Requirements

Condition	Operating Requirement	Nonoperating Requirement	Optimum
Temperature	5 to 32 degrees Celsius (41 to 89.6 degrees Fahrenheit)	-40 to 70 degrees Celsius (-40 to 158 degrees Fahrenheit)	21 to 23 degrees Celsius (70 to 74 degrees Fahrenheit)
Relative humidity	10 to 90 percent relative humidity, noncondensing	Up to 93 percent relative humidity	45 to 50 percent, noncondensing
Altitude	3,048 meters (10,000 feet) maximum	12000 meters (40,000 feet).	Ambient temperature is reduced by 1 degree Celsius for each 300 meters above 900 meters altitude above sea level

Set conditions to the optimal temperature and humidity ranges to minimize the chance of downtime because of component failure. Operating Recovery Appliance for extended periods at or near the operating range limits, or installing it in an environment where it remains at or near nonoperating range limits, can 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 if the air conditioning system fails.

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:

- Helps protect computer systems from corrosion problems associated with high humidity levels.
- Provides the greatest operating time buffer if the air conditioner control fails
- Helps prevent failures or temporary malfunctions caused by intermittent interference from static discharges, which may occur when the relative humidity is too low.

 **Note:**

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

2.6 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 through the rear of the rack. There is no air flow requirement for the left and right sides because of 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 it is inversely proportional to the air temperature. Humidity goes down when the temperature rises, and it 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 form.

Air conditioning facilities usually do not precisely monitor or control temperature and humidity throughout an entire computer room. Generally, you should monitor individual points corresponding to multiple exhaust vents in the main unit and other units in the room, because the distribution of temperature and humidity is uneven across the room. Pay special attention to humidity when using underfloor ventilation.

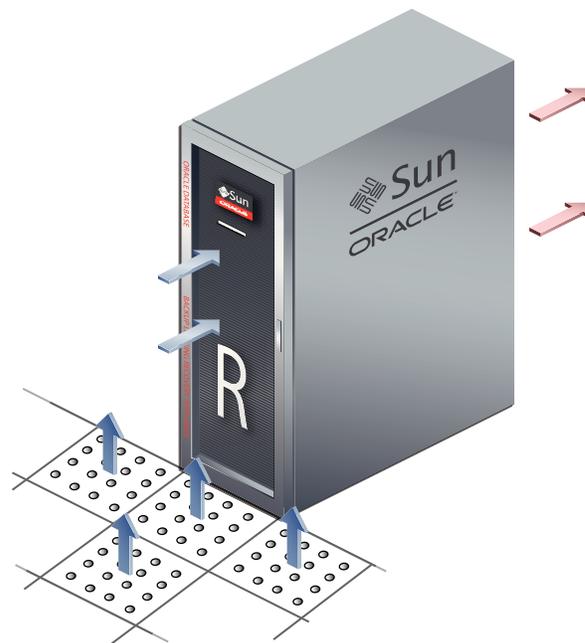
The Recovery Appliance racks are designed to function while mounted in a natural convection air flow. Follow these requirements to meet the environmental specification:

- Ensure that the rack has adequate air flow
- Ensure that the rack has front-to-back cooling. The air inlet is at the front of the servers, and the air is let out the rear.
- Allow a minimum clearance of 152.5 cm (60 inches) at the front of the rack and 91.4 cm (36 inches) at the rear for ventilation.

Use perforated tiles, rated for 400 cubic feet per minute (CFM) per tile air flow, in front of the rack for cold air intake. The tiles can be arranged in any order in front of the rack, if 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. Oracle recommends four perforated floor tiles for a full Recovery Appliance rack, and one tile for a minimum configuration.

Figure 2-5 (page 2-15) shows a typical installation of the floor tiles for a full rack.

Figure 2-5 Typical Data Center Configuration for Perforated Floor Tiles



2.7 Network Connection Requirements for Recovery Appliance

Before installation, network cables must be run from your existing network infrastructure to the installation site. When you run Oracle Exadata Deployment Assistant, it calculates the number of network connections that you need, based on your rack configuration and other information you provide. Use this information when preparing your site for the installation.

 **See Also:**

- [Preparing Your Networks for Recovery Appliance](#) (page 3-1)
- [Using the Deployment Assistant](#) (page 5-1)

2.8 Ensuring That the Site Is Ready

Before Recovery Appliance is delivered to the site, ensure that it is ready by completing these tasks:

- "[Generate the Configuration Files](#) (page 2-16)"
- "[Configure the Network](#) (page 2-16)"
- "[Prepare the Site](#) (page 2-16)"

2.8.1 Generate the Configuration Files

Run the Oracle Exadata Deployment Assistant and generate the configuration files.

 **See Also:**

[Using the Deployment Assistant](#) (page 5-1) for information about the configuration utility

2.8.2 Configure the Network

Configure the network using the information generated by Oracle Exadata Deployment Assistant. Configuration includes the following tasks:

- Registering the networks in the Domain Name System (DNS)
- Assigning IP addresses
- Configuring data center switches and firewalls

 **See Also:**

- [Preparing Your Networks for Recovery Appliance](#) (page 3-1)
- [Using the Deployment Assistant](#) (page 5-1)

2.8.3 Prepare the Site

Review the site requirements in this chapter and the *Zero Data Loss Recovery Appliance Site Checklists* when preparing the site for installation.

Before the Recovery Appliance rack arrives at your site:

1. Review the safety guidelines.

 **See Also:**

"[Reviewing Safety Guidelines](#) (page 6-1)"

2. Note problems or peculiarities at the site that require special equipment. For example, ensure the doors are tall enough and wide enough for the Recovery Appliance rack.

 **See Also:**

"[Space Requirements](#) (page 2-5)"

3. Verify that the installation site flooring has a strength rating to withstand the combined weight of Recovery Appliance and any other installed equipment.

 **See Also:**

"[Flooring Requirements](#) (page 2-7)"

4. Install network cables for Recovery Appliance.

 **See Also:**

"[Network Connection Requirements for Recovery Appliance](#) (page 2-15)"

5. Install all necessary electrical equipment, and ensure that sufficient power is provided for Recovery Appliance.

 **See Also:**

- "[Electrical Power Requirements](#) (page 2-8)"
- The *Sun Rack II Power Distribution Units User's Guide* for the Sun Rack II Power Distribution Unit (PDU) power requirements

6. Ensure that the installation site provides adequate air conditioning.

 **See Also:**

"[Ventilation and Cooling Requirements](#) (page 2-14)"

7. Operate the air conditioning system for 48 hours to bring the room temperature to the appropriate level.
8. Clean and vacuum the area thoroughly in preparation for installation.

3

Preparing Your Networks for Recovery Appliance

This chapter describes the network requirements for Recovery Appliance, so that you can prepare your data center for installation.

This chapter contains the following sections:

- [Overview of Network Requirements](#) (page 3-1)
- [Registering Recovery Appliance in the Domain Name System](#) (page 3-10)
- [Factory IP Address Settings](#) (page 3-10)
- [Port Assignments When Using a Firewall](#) (page 3-11)

3.1 Overview of Network Requirements

Recovery Appliance includes two compute servers, 3 to 18 storage servers (depending on your hardware version), and the equipment to connect these servers to your network. The network connections allow the servers to be administered remotely, and for data to move from the protected databases to Recovery Appliance. Use the information in this section to prepare your data center for the addition of Recovery Appliance.

- [About the Network Components and Interfaces](#) (page 3-1)
- [About the Networks](#) (page 3-3)
- [Example of Network Connections for Recovery Appliance](#) (page 3-5)
- [Connecting Recovery Appliance Rack Components to the Networks](#) (page 3-6)
- [Connecting Recovery Appliance to a Tape Library](#) (page 3-7)
- [Using Network VLAN Tagging with Recovery Appliance](#) (page 3-9)

See Also:

"[About Ingest Over InfiniBand](#) (page 9-1)" for information about configuring local Oracle Exadata Database Machine racks to backup to Recovery Appliance over the InfiniBand network

3.1.1 About the Network Components and Interfaces

Each compute server in the X7 configuration consists of the following network components and interfaces:

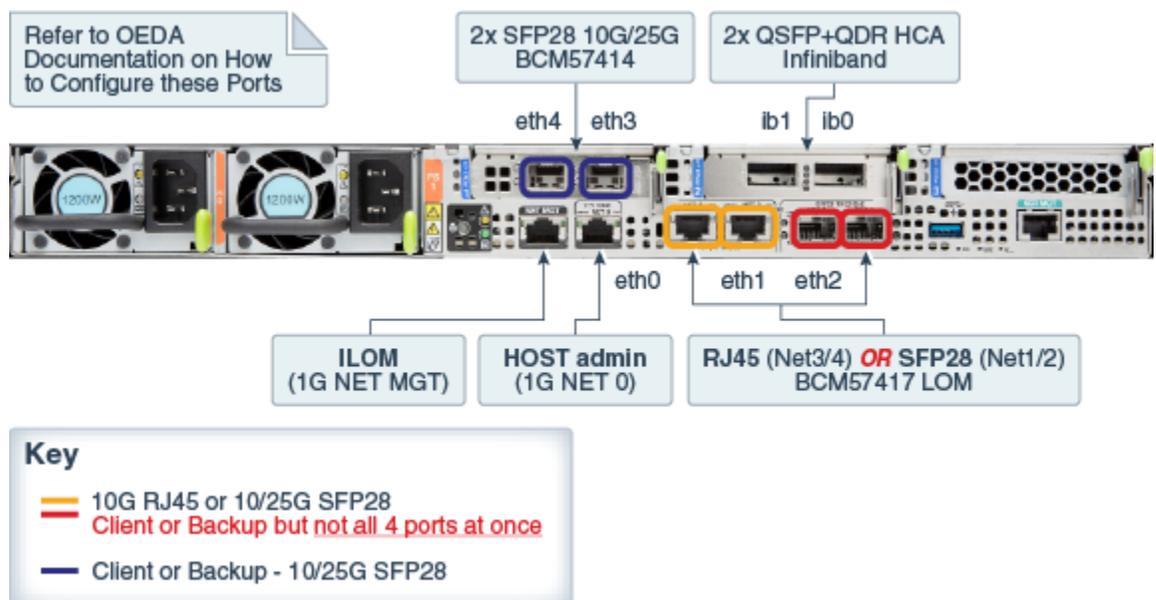
- Ethernet ingest and replication network connectivity

- On-board: 2 x 10 Gb copper Ethernet (eth1)
- On-board: 2 x 10/25 Gb optical Ethernet Ports (eth2)
- PCIe card: 2 x 10/25 Gb optical Ethernet Ports (eth3 and eth4)
- 1 dual-port 4X QDR (40 Gbps) InfiniBand Host Channel Adapter (HCA) (IB0 and IB1)
- 1 Ethernet port for Oracle Integrated Lights Out Manager (ILOM) remote management
- 1 dual-port 32 GB FC Converged Network Adapter (CNA) FC ports 0 and 1.

 **Note:**

The corresponding SFP modules that work with the 10/25 GbE PCIe 2.0 network cards are purchased separately.

Figure 3-1 X7 Back—plane External Network Connectivity



Given that the base rack has two (2) compute servers, the maximum for ingest is 2 x 10 Gb or 2 x 25 Gb Ethernet ports, while the maximum for replication is 2 x 10 Gb or 2 x 25 Gb Ethernet ports. The following are valid combinations of the options.

- 2 x 10Gb on-board copper (ingest) + 2 x 10/25Gb PCIe card optical (replication)
- 2 x 10/25Gb PCIe card optical (ingest) + 2 x 10Gb on-board copper (replication)
- 2 x 10/25Gb PCIe card optical (ingest) + 2 x 10/25Gb on-board optical (replication)
- 2 x 10/25Gb on-board optical (ingest) + 2 x 10/25Gb PCIe card optical (replication)

 **Note:**

If ingest and replication traffic is desired to be configured on the **same network**, define the required network interface in OEDA in the ingest network section, and leave the replication network section blank. With this setup, Recovery Appliance will use the ingest network for replication traffic.

Each storage server consists of the following network components and interfaces:

- 1 embedded Gigabit Ethernet port (NET0)
- 1 dual-port 4X QDR (40 Gbps) InfiniBand Host Channel Adapter (HCA) (IB0 and IB1)
- 1 Ethernet port for Oracle Integrated Lights Out Manager remote management (Oracle ILOM)

Additional configuration, such as defining multiple virtual local area networks (VLANs) for the management (NET0 and/or ILOM) interfaces or enabling routing, might be required for the switch to operate properly in your environment and is beyond the scope of the installation service. If additional configuration is needed, then your network administrator must perform the necessary configuration steps during installation of Recovery Appliance.

3.1.2 About the Networks

To deploy Recovery Appliance, ensure that you meet the minimum network requirements. There are up to five networks.

 **Note:**

Each network must be on a distinct and separate subnet from the others. The only exception is that the replication network defaults to using physical ingest network, if not assigned in OEDA.

The network descriptions are as follows:

- **Management network:** This required, 1 gigabit, Ethernet (GbE) network connects to your existing management network. You use it to do administrative work on all components of Recovery Appliance. The management network connects to the servers, Oracle ILOM, and InfiniBand switches through the Cisco Ethernet switch in the rack.

There are three uplinks to the management network:

- One from the Ethernet switch
- One from each of the two power distribution units (PDUs). Network connectivity to the PDUs is required only if the electric current is monitored remotely.

Each server has two network interfaces for management. One provides management access to the operating system through the NET0 Ethernet interface, and the other provides access to Oracle ILOM through the ILOM Ethernet

interface. Recovery Appliance is delivered with the ILOM and NET0 interfaces connected to the Ethernet switch on the rack.

Cabling or configuration changes to these interfaces is not permitted. Do not use the management network interfaces for client or application network traffic. Instead, you can use NET1, NET2, NET3, or the two interfaces on the PCI slot.

- **Ingest network:** This required network connects the protected Oracle Database servers to Recovery Appliance for backup within the same data center. Also known as a backup network, this high-speed, private Ethernet network must be designed to support the transfer of large volumes of data. Recovery Appliance connects to this network using two 10/25 GB connections to each of the two compute servers in the rack. You can configure the two connections as active/passive (redundant) or active/active. The compute servers support channel bonding to provide higher bandwidth and availability.

Single client access name (SCAN) supports failover between the two compute servers in Recovery Appliance. In an installation with multiple Recovery Appliance racks configured as a cluster, virtual IP (VIP) addresses support failover among the racks. The protected database systems can resolve the host names to dynamically assigned addresses.

Third-party tape hardware and software also uses the ingest network.

- **Replication network:** The optional replication network connects the local Recovery Appliance (the **upstream** appliance) with a remote Recovery Appliance (the **downstream** appliance). Oracle recommends a broadband, encrypted network, instead of an insecure public network, wherever possible.

Recovery Appliance supports the following configurations between the upstream and downstream appliances:

- **One way:** Data flows from the upstream appliance to the downstream appliance.
- **Bidirectional:** Data flows in both directions between the upstream appliance and the downstream appliance.
- **Hub and spoke:** Data flows from multiple upstream appliances to one downstream appliance.

 **Note:**

A downstream Recovery Appliance or a tape library can reside in the local data center. The replication network is not used in a local configuration.

- **Fibre Channel SAN network:** If you are using Oracle Secure Backup, then you can back up Recovery Appliance to the storage area network (SAN) in your data center for backups to tape. The network connections depend on whether you have an Oracle tape solution or use third-party hardware.
- **InfiniBand private network:** This network connects the compute servers and storage servers using the InfiniBand switches. Oracle Database uses this network for Oracle RAC cluster interconnect traffic and for accessing data on the storage servers. This nonroutable network is fully contained in Recovery Appliance.

This network does not connect to your data center networks. It is automatically configured during installation.

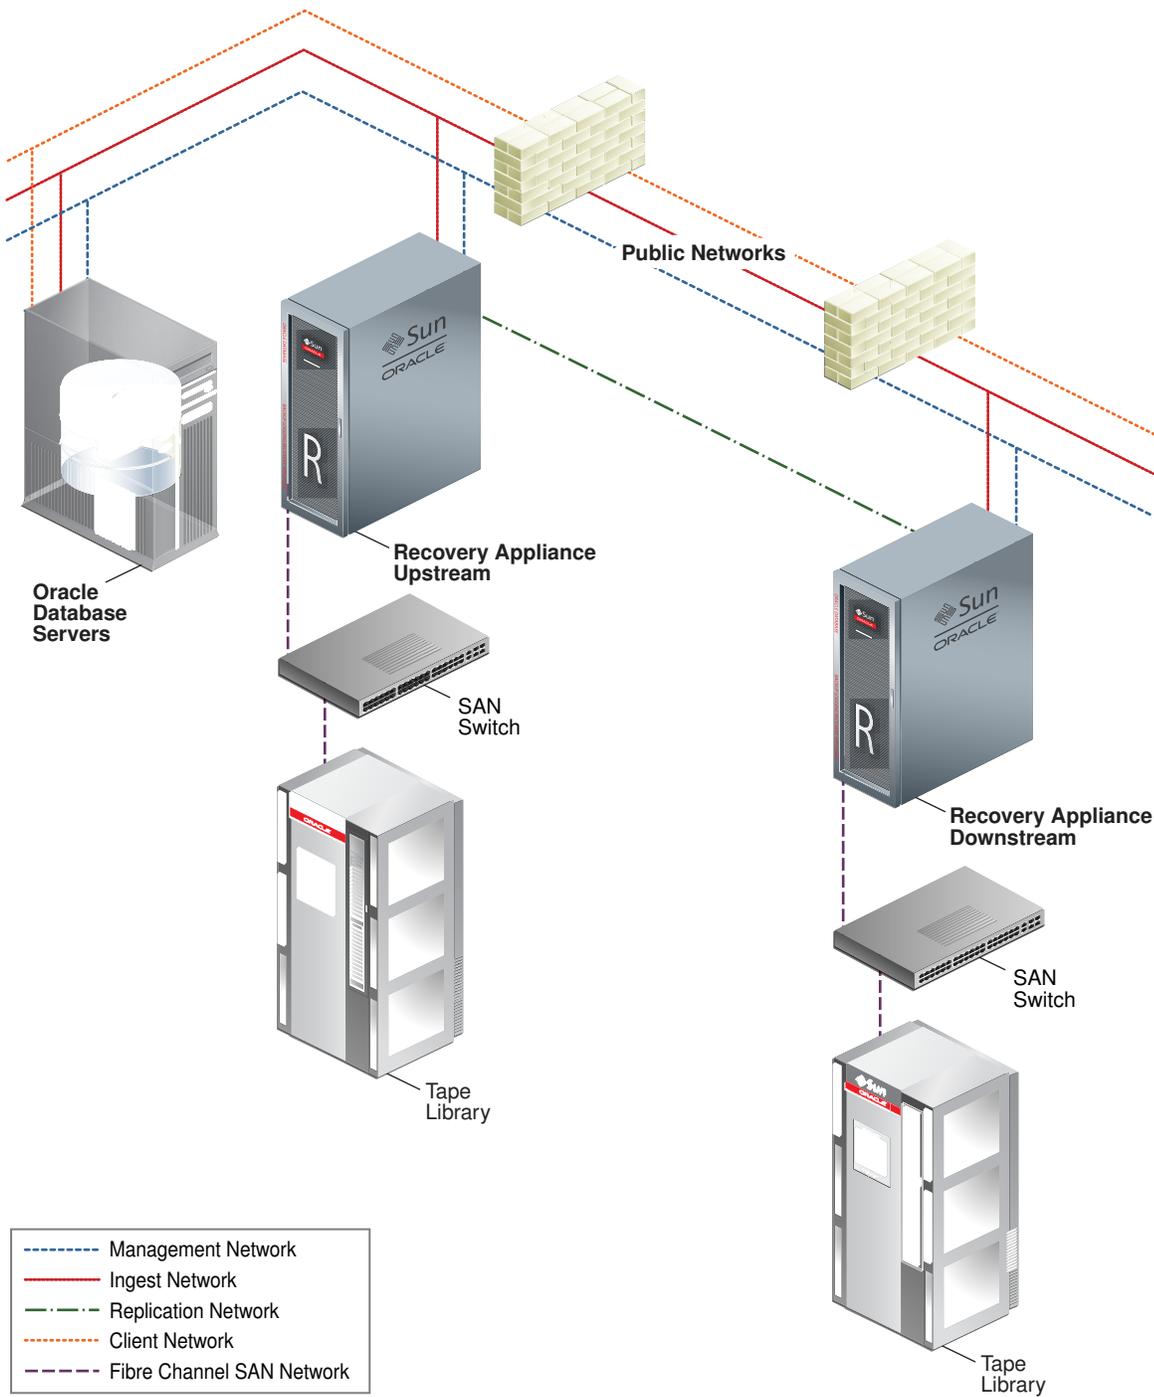
 **See Also:**

- *Oracle Clusterware Administration and Deployment Guide* for a discussion of SCANs and VIPs in network configurations for Oracle Database.
- "[Connecting Recovery Appliance to a Tape Library](#) (page 3-7)" for information about how a fibre channel SAN network is configured for backups to tape in a Recovery Appliance environment.

3.1.3 Example of Network Connections for Recovery Appliance

[Figure 3-2](#) (page 3-6) shows the network cabling of a sample configuration. Two Recovery Appliance racks are installed in separate data centers. The protected Oracle databases are connected to the upstream Recovery Appliance over the ingest network. The upstream Recovery Appliance is connected to the downstream Recovery Appliance over the replication network. Both racks are configured to use an Oracle tape solution.

Figure 3-2 Network Diagram for Recovery Appliance



3.1.4 Connecting Recovery Appliance Rack Components to the Networks

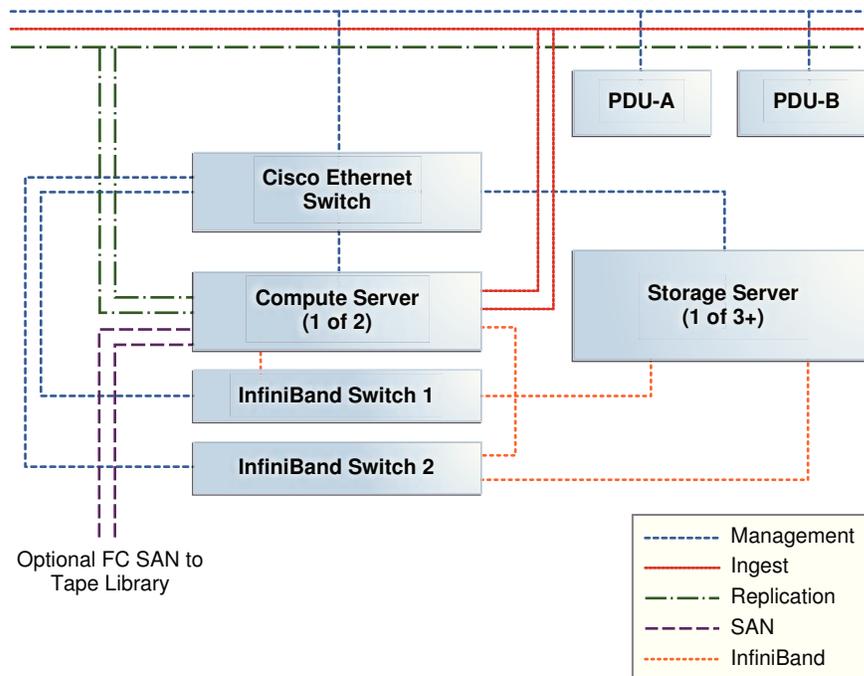
Figure 3-3 (page 3-7) shows the network connections to components of Recovery Appliance rack.

The management network connects through the Ethernet switch to the compute servers, the storage servers, and the InfiniBand switches. The management network connects directly to the PDUs.

The ingest network, the optional replication network, and the optional fiber channel SAN network connect to the two compute servers.

The InfiniBand network connects the InfiniBand switches to the compute servers and the storage servers.

Figure 3-3 Network Connections to the Recovery Appliance Rack Components



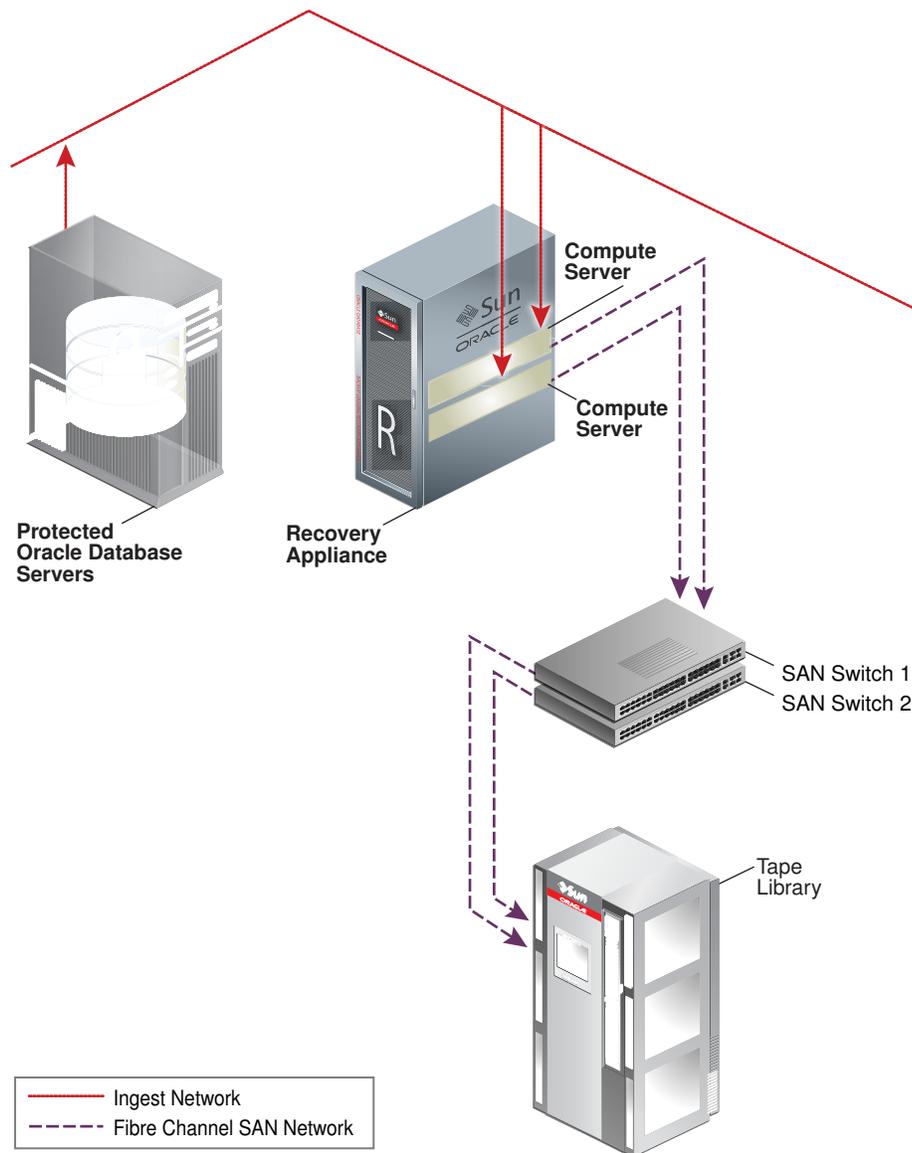
3.1.5 Connecting Recovery Appliance to a Tape Library

The network connections between Recovery Appliance and an optional tape library depend on whether you are using Oracle or third-party tape components. See "[About Tape Backup Infrastructure](#) (page 1-8)" for the differences in support provided by Recovery Appliance.

3.1.5.1 Oracle Recommended Stack

When you use the Oracle recommended tape solution, a fiber channel adapter is installed in each compute server to provide a connection to the fiber channel storage area network (SAN). Tape backups are isolated on this network, and thus do not interfere with the performance of the other networks. [Figure 3-4](#) (page 3-8) provides an overview of the network connections when using an Oracle tape system.

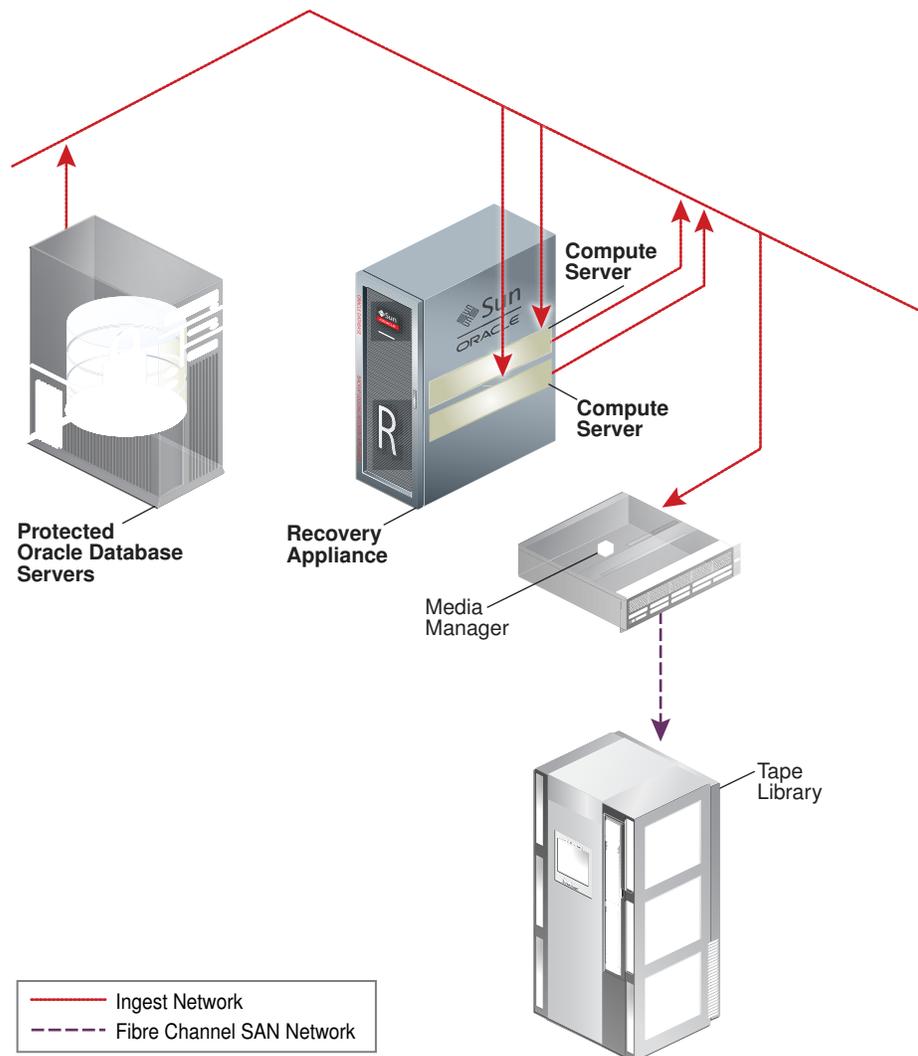
Figure 3-4 Recovery Appliance Connection to an Oracle Tape System



3.1.5.2 Third-Party Tape Systems

When you use a third-party tape system, the backups to tape use the 10 Gb ingest network. This is the same network that the local protected databases use to backup to Recovery Appliance. [Figure 3-5](#) (page 3-9) provides an overview of the network connections when using a third-party tape system.

Figure 3-5 Recovery Appliance Connection to a Third-Party Tape System



3.1.6 Using Network VLAN Tagging with Recovery Appliance

The Recovery Appliance supports VLAN port tagging only on the ingest network. You configure VLAN port tagging after you complete the Recovery Appliance installation.

If applicable, ensure that you also set the Access VLAN on the network switches, including on the Cisco switch that is included in the Recovery Appliance rack for the management network.

 **See Also:**

"[Installing the Software on Recovery Appliance](#) (page 11-3)" for instructions on when and how to configure VLAN tagging

3.2 Registering Recovery Appliance in the Domain Name System

Before receiving your Recovery Appliance rack, use Oracle Exadata Deployment Assistant. The assistant generates a file to be used when setting up the system. The host names and IP addresses specified in the assistant-generated file must be registered in Domain Name System (DNS) before the initial configuration. In addition, all public addresses, single client access name (SCAN) addresses, and VIP addresses must be registered in DNS before 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 Recovery Appliance. Configure DNS for round robin resolution for the SCAN name to these three SCAN addresses.

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

See Also:

- [Using the Deployment Assistant](#) (page 5-1)
- *Oracle Grid Infrastructure Installation Guide for Linux* for additional information about SCAN addresses
- Your DNS vendor documentation for additional information about configuring round-robin name resolution

3.3 Factory IP Address Settings

Recovery Appliance has default IP addresses set at the factory:

- Gateway: 192.168.1.254 in all devices as required
- Subnet Mask: 255.255.252.0 in all devices as required
- IP Address Range: 192.168.1.1 to 192.168.1.203

Before connecting Recovery Appliance to the network, ensure that these IP addresses do not conflict with other addresses on the network. The `checkip.sh` script checks for conflicts. Oracle recommends running the script before connecting the network to avoid problems, even when a check was performed before Recovery Appliance was delivered. See "[Installing the Software on Recovery Appliance](#) (page 11-3)" for additional information about the `checkip.sh` script.

Table 3-1 (page 3-11) lists the factory IP addresses for a Recovery Appliance full rack.

Table 3-1 Factory IP Addresses for Recovery Appliance

Rack Unit	Component	Management Network Addresses	InfiniBand Active Bonded IP Addresses	Oracle ILOM IP Addresses
U41	Storage server	192.168.1.23	192.168.10.45	192.168.1.123
U39	Storage server	192.168.1.22	192.168.1.43	192.168.1.122
U37	Storage server	192.168.1.21	192.168.10.41	192.168.1.121
U35	Storage server	192.168.1.20	192.168.10.39	192.168.1.120
U33	Storage server	192.168.1.19	192.168.10.37	192.168.1.119
U31	Storage server	192.168.1.18	192.168.10.35	192.168.1.118
U29	Storage server	192.168.1.17	192.168.10.33	192.168.1.117
U27	Storage server	192.168.1.16	192.168.10.31	192.168.1.116
U25	Storage server	192.168.1.14	192.168.10.27	192.168.1.114
U23	Storage server	192.168.1.12	192.168.10.23	192.168.1.112
U22	InfiniBand switch	Not applicable	Not applicable	192.168.1.203
U21	Ethernet switch	Not applicable	Not applicable	192.168.1.200
U20	InfiniBand switch	Not applicable	Not applicable	192.168.1.202
U18	Storage server	192.168.1.10	192.168.10.19	192.168.1.110
U17	Compute server	192.168.1.9	192.168.10.17	192.168.1.109
U16	Compute server	192.168.1.8	192.168.10.15	192.168.1.108
U14	Storage server	192.168.1.7	192.168.10.13	192.168.1.107
U12	Storage server	192.168.1.6	192.168.10.11	192.168.1.106
U10	Storage server	192.168.1.5	192.168.10.9	192.168.1.105
U08	Storage server	192.168.1.4	192.168.10.7	192.168.1.104
U06	Storage server	192.168.1.3	192.168.10.5	192.168.1.103
U04	Storage server	192.168.1.2	192.168.10.3	192.168.1.102
U02	Storage server	192.168.1.1	192.168.10.1	192.168.1.101

3.4 Port Assignments When Using a Firewall

When network communication between Recovery Appliance and other components requires access through a firewall, you must open ports used by the Recovery Appliance services.

 **Note:**

A firewall may not be used between components of the Recovery Appliance.

[Table 3-2](#) (page 3-12) lists the ports used by services on Recovery Appliance. Review the list and open the necessary ports. All ports are on the management network, unless otherwise noted.

Table 3-2 Open Ports for the Firewall

Source	Target	Protocol	Port	Application
NA	Database management	SSH over TCP	22	SSH
NA	Compute servers, storage servers, and InfiniBand ILOMs	SSH over TCP	22	SSH
NA	Storage management	SSH over TCP	22	SSH
Storage servers	email server	SMTP	25 465 if using SSL	SMTP (Simple Mail Transfer Protocol)
Compute servers, storage servers, and InfiniBand ILOMs	NA	TFTP over UDP	69	Outgoing TFTP (Trivial File Transfer Protocol)
NA	Compute servers, storage servers, and InfiniBand ILOMs	HTTP over TCP	80	Web (user configurable)
NA	PDU	HTTP over TCP	80	Browser interface
Database management	NA	NTP over UDP	123	Outgoing Network Time Protocol (NTP)
Compute servers, storage servers, and InfiniBand ILOMs	NA	NTP over UDP	123	Outgoing NTP
Storage management	NA	NTP over UDP	123	Outgoing NTP
ASR Manager	ASR asset	SNMP (get)	161	FMA enrichment for additional diagnostic information
NA	Compute servers, storage servers, and InfiniBand ILOMs	SNMP over UDP	161	SNMP (Simple Network Management Protocol) (user configurable)
NA	PDU	SNMP over UDP	161	SNMP (user configurable)
Storage servers	SNMP subscriber such as Oracle Enterprise Manager Cloud Control or an SNMP manager	SNMP	162	SNMP version 1 (SNMPv1) outgoing traps (user-configurable)
Compute servers and storage server ILOMs	ASR Manager	SNMP	162	Telemetry messages sent to ASR Manager
Compute servers, storage servers, and InfiniBand ILOMs	NA	IPMI over UDP	162	Outgoing IPMI (Intelligent Platform Management Interface) Platform Event Trap (PET)
PDU	NA	SNMP over UDP	162	Outgoing SNMPv2 traps

Table 3-2 (Cont.) Open Ports for the Firewall

Source	Target	Protocol	Port	Application
NA	Compute servers, storage servers, and InfiniBand ILOMs	LDAP over UDP/TCP	389	Outgoing LDAP (Lightweight Directory Access Protocol) (user configurable)
ASR Manager	ASR back end	HTTPS	443	Telemetry messages sent to ASR back end
NA	Compute servers, storage servers, and InfiniBand ILOMs	HTTPS over TCP	443	Web (user configurable)
NA	PDU	HTTPS over TCP	443	Browser interface
Compute servers, storage servers, and InfiniBand ILOMs	NA	Syslog over UDP	514	Outgoing Syslog
PDU	NA	Syslog over UDP	514	Outgoing Syslog
Compute servers, storage servers, and InfiniBand ILOMs	NA	DHCP over UDP	546	client DHCP (Dynamic Host Configuration Protocol)
PDU	NA	DHCP over UDP	546	DHCP (Dynamic Host Configuration Protocol) client
NA	Compute servers, storage servers, and InfiniBand ILOMs	IPMI over UDP	623	IPMI (Intelligent Platform Management Interface)
Oracle Enterprise Manager Cloud Control	NA	TCP	1159	Oracle Enterprise Manager Cloud Control HTTPS upload port
Oracle Enterprise Manager Cloud Control	NA	TCP	1159	Oracle Enterprise Manager Cloud Control HTTPS upload port
NA	Database data	SQL*Net over TCP	1521	Database listener
Protected database	Recovery Appliance	SQL*Net over TCP	1521 (ingest network)	RMAN backup and restore
Upstream Recovery Appliance	Downstream Recovery Appliance	SQL*Net over TCP	1522 (replication network)	Recovery Appliance Replication
Compute servers, storage servers, and InfiniBand ILOMs	NA	RADIUS over UDP	1812	Outgoing RADIUS (Remote Authentication Dial In User Service) (user configurable)

Table 3-2 (Cont.) Open Ports for the Firewall

Source	Target	Protocol	Port	Application
Oracle Enterprise Manager Grid Control	NA	TCP	4889	Oracle Enterprise Manager Cloud Control HTTP upload port
Oracle Enterprise Manager Grid Control	NA	TCP	4889	Oracle Enterprise Manager Cloud Control HTTP upload port
NA	Compute server and storage server ILOMs	TCP	5120	ILOM remote console: CD
NA	Compute server and storage server ILOMs	TCP	5121	ILOM remote console: keyboard and mouse
NA	Compute server and storage server ILOMs	TCP	5123	ILOM remote console: diskette
NA	Compute server and storage server ILOMs	TCP	5555	ILOM remote console: encryption
NA	Compute server and storage server ILOMs	TCP	5556	ILOM remote console: authentication
ASR Manager	Compute server and storage server ILOMs	HTTP	6481	Service tags listener for asset activation
NA	Compute server and storage server ILOMs	TCP	6481	ILOM remote console: servicetag daemon
NA	Compute server and storage server ILOMs	TCP	7578	ILOM remote console: video
NA	Compute server and storage server ILOMs	TCP	7579	ILOM remote console: serial
NA	Compute servers	TCP	7777	Oracle Enterprise Manager Grid Control HTTP console port
NA	Storage servers	TCP	7777	Oracle Enterprise Manager Grid Control HTTP console port
NA	Compute servers	TCP	7799	Oracle Enterprise Manager Grid Control HTTPS console port
NA	Storage servers	TCP	7799	Oracle Enterprise Manager Grid Control HTTPS console port
Protected database	Recovery Appliance	HTTP	8001 (ingest network)	RMAN backup and restore

Table 3-2 (Cont.) Open Ports for the Firewall

Source	Target	Protocol	Port	Application
Upstream Recovery Appliance	Downstream Recovery Appliance	HTTP	8001 (replication network)	Recovery Appliance Replication

4

Setting Up Auto Service Request

This chapter explains how to install and configure Auto Service Request for Recovery Appliance. It contains these sections:

- [Understanding Auto Service Request](#) (page 4-1)
- [Getting Ready to Install ASR](#) (page 4-2)
- [Installing ASR Manager](#) (page 4-3)
- [Verifying the ASR Manager Installation](#) (page 4-3)
- [About the Trap Destinations on Recovery Appliance](#) (page 4-4)
- [Troubleshooting ASR](#) (page 4-5)

4.1 Understanding Auto Service Request

Auto Service Request (ASR) is designed to automatically open service requests when specific Recovery Appliance hardware faults occur. ASR detects faults in the most common server components, such as disks, fans, and power supplies, and automatically opens a service request when a fault occurs. ASR monitors only server components and does not detect all possible faults.

ASR is not a replacement for other monitoring mechanisms, such as SMTP and SNMP alerts, within the customer data center. It is a complementary mechanism that expedites and simplifies the delivery of replacement hardware. ASR should not be used for downtime events in high-priority systems. For high-priority events, contact Oracle Support Services directly.

When ASR detects a hardware problem, ASR Manager submits a service request to Oracle Support Services. In many cases, Oracle Support Services can begin work on resolving the issue before the administrator is even aware the problem exists.

An email message is sent to both the My Oracle Support email account and the technical contact for Recovery Appliance to notify them of the creation of the service request.

A service request might not be filed automatically on some occasions. This can happen because of the unreliable nature of the SNMP protocol or a loss of connectivity to ASR Manager. Oracle recommends that you continue to monitor your systems for faults and call Oracle Support Services if you do not receive notice that a service request was filed automatically.

 **See Also:**

- Oracle Auto Service Request web page at <http://www.oracle.com/technetwork/systems/asr/overview/index.html>
- Oracle Auto Service Request user documentation at <http://www.oracle.com/technetwork/systems/asr/documentation/index.html>

4.2 Getting Ready to Install ASR

Before installing ASR, complete these prerequisites:

1. Create a My Oracle Support account at <http://support.oracle.com>.
2. Ensure that the following are set up correctly in My Oracle Support:
 - Oracle Premier Support for Systems, Oracle Premier Support for Operating Systems, or Oracle Limited Warranty
 - The technical contact person at the customer site who is responsible for Recovery Appliance
 - A valid shipping address at the customer site for the Recovery Appliance parts
3. Identify and designate a system to host ASR Manager.

ASR Manager must be installed on a server that has connectivity to Recovery Appliance and an outbound Internet connection using HTTPS or an HTTPS proxy. To submit a service request (SR), the server must be able to access the Internet.

4. ASR Manager uses the following ports for normal operations. Ensure that Recovery Appliance and ASR Manager server ports are open:
 - All Recovery Appliance servers must have port 6481 open for HTTP-based asset activation requests that they receive from ASR Manager.
 - The server running ASR Manager must have port 162 open for SNMP-based telemetry messages that the Recovery Appliance servers send to it.
 - The Oracle `transport.oracle.com` web server has port 443 open for HTTPS-based telemetry messages that ASR Manager sends to it.

See [Figure 4-1](#) (page 4-3) for the port locations.

5. Ensure that the designated system conforms to the "Hardware and Network Configuration Recommendations for ASR." See the Oracle ASR website:
http://docs.oracle.com/cd/E37710_01/nav/products.htm
6. Confirm that Java Development Kit 6 (JDK 1.6.0_04 or later) is running on the designated ASR Manager system:

```
java -version
```

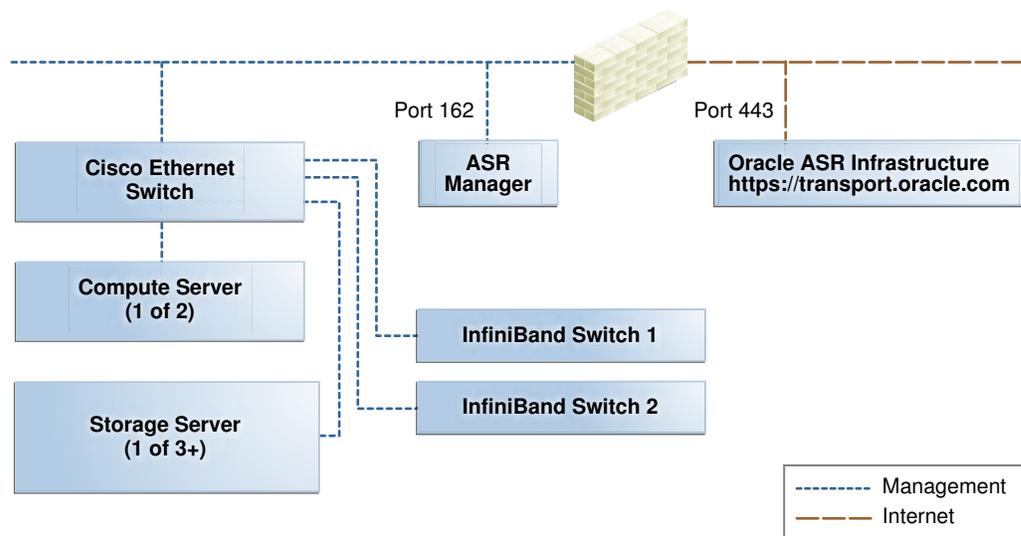
If necessary, download and install the latest version of JDK from the Java SE Downloads website:

<http://www.oracle.com/technetwork/java/javase/downloads/index.html>

7. Obtain `root` access to the designated ASR Manager host.
8. Identify and verify connectivity to Recovery Appliance.
9. Verify connectivity to the Internet using HTTPS.

Figure 4-1 (page 4-3) shows the network connections between ASR and Recovery Appliance.

Figure 4-1 Auto Service Request Network Connections



4.3 Installing ASR Manager

To install ASR Manager, download the current version from My Oracle Support Doc ID 1185493.1. Then follow the instructions in the *Oracle Auto Service Request Installation and Operations Guide* at

<http://www.oracle.com/technetwork/systems/asr/documentation/index.html>

4.4 Verifying the ASR Manager Installation

Perform these checks as `root` on ASR Manager to ensure that it is installed properly.

- Ensure that the `asr` command is in the search path of the `root` user.
- Verify that ASR Manager 3.5 or later is running:


```
asr show_rules_version
```
- Check the registration status:


```
asr show_reg_status
```
- Test the connection by sending a test message to the transport server.


```
asr test_connection
```
- Verify the ASR assets:

1. In ASR Manager, verify that ASR is activated:

```
asr list_asset -i asset_ip
```

In the preceding command, *asset_ip* is the IP address of a server or an Oracle ILOM. To list all assets, enter this command:

```
asr list_asset
```

2. If no assets are listed, then verify that ASR is configured on Recovery Appliance.

 **Note:**

If an IP address or host name changes, then you must deactivate and reactivate the asset.

4.5 About the Trap Destinations on Recovery Appliance

ASR is an optional software component of Recovery Appliance. When configuring the software in Oracle Exadata Deployment Assistant, you must complete the page about ASR. Then during the software installation, Recovery Appliance configures the servers to trap the component fault indicators. See "[Auto Service Request Page](#) (page 5-24)".

The following are examples of Simple Network Management Protocol (SNMP) traps sent to ASR Manager for disk failures.

[Example 4-1](#) (page 4-4) shows the SNMP trap for an storage server disk failure. The hardware alert code (HALRT-02001) is underlined.

[Example 4-2](#) (page 4-4) shows the SNMP trap from a compute server disk failure. The hardware alert code (HALRT-02007) is underlined.

Example 4-1 Example of a Storage Server SNMP Trap

```
2011-09-07 10:59:54 server1.example.com [UDP: [192.85.884.156]:61945]:
RFC1213-MIB::sysUpTime.0 = Timeticks: (52455631) 6 days, 1:42:36.31
SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: SUN-HW-TRAP-MIB::sunHwTrapHardDriveFault
SUN-HW-TRAP-MIB::sunHwTrapSystemIdentifier = STRING: Sun Oracle Database Machine
1007AK215C
SUN-HW-TRAP-MIB::sunHwTrapChassisId = STRING: 0921XFG004
SUN-HW-TRAP-MIB::sunHwTrapProductName = STRING: SUN FIRE X4270 M2 SERVER
SUN-HW-TRAP-MIB::sunHwTrapSuspectComponentName = STRING: SEAGATE ST32000SSSUN2.0T;
Slot: 0SUN-HW-TRAP-MIB::sunHwTrapFaultClass = STRING: NULL
SUN-HW-TRAP-MIB::sunHwTrapFaultCertainty = INTEGER: 0
SUN-HW-TRAP-MIB::sunHwTrapFaultMessageID = STRING: HALRT-02001
SUN-HW-TRAP-MIB::sunHwTrapFaultUUID = STRING: acb0a175-70b8-435f-9622-38a9a55ee8d3
SUN-HW-TRAP-MIB::sunHwTrapAssocObjectId = OID: SNMPv2-SMI::zeroDotZero
SUN-HW-TRAP-MIB::sunHwTrapAdditionalInfo = STRING: Exadata Storage Server:
cellname Disk Serial Number: E06S8K
server1.example.com failure trap.
```

Example 4-2 Example of Oracle Database Server SNMP Trap

```
2011-09-09 10:59:54 dbserv01.example.com [UDP: [192.22.645.342]:61945]:
RFC1213-MIB::sysUpTime.0 = Timeticks: (52455631) 6 days, 1:42:36.31
SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: SUN-HW-TRAP-MIB::sunHwTrapHardDriveFault
SUN-HW-TRAP-MIB::sunHwTrapSystemIdentifier = STRING: Sun Oracle Database Machine
```

```
1007AK215C
SUN-HW-TRAP-MIB::sunHwTrapChassisId = STRING: 0921XFG004
SUN-HW-TRAP-MIB::sunHwTrapProductName = STRING: SUN FIRE X4170 M2 SERVER
SUN-HW-TRAP-MIB::sunHwTrapSuspectComponentName = STRING: HITACHI H103030SCSUN300G
Slot: 0SUN-HW-TRAP-MIB::sunHwTrapFaultClass = STRING: NULL
SUN-HW-TRAP-MIB::sunHwTrapFaultCertainty = INTEGER: 0
SUN-HW-TRAP-MIB::sunHwTrapFaultMessageID = STRING: HALRT-02007
SUN-HW-TRAP-MIB::sunHwTrapFaultUUID = STRING: acb0a175-70b8-435f-9622-38a9a55ee8d3
SUN-HW-TRAP-MIB::sunHwTrapAssocObjectId = OID: SNMPv2-SMI::zeroDotZero
SUN-HW-TRAP-MIB::sunHwTrapAdditionalInfo = STRING: Exadata Database Server: db03
Disk Serial Number: HITACHI H103030SCSUN300GA2A81019GGDE5E
dbserv01.example.com failure trap.
```

4.6 Troubleshooting ASR

For troubleshooting procedures for the ASR software, see Chapter 5 of the *Oracle ASR Installation and Operations Guide* at:

<http://www.oracle.com/technetwork/systems/asr/documentation/index.html>

If you continue to have issues, contact ASR support.

5

Using the Deployment Assistant

This chapter describes how to use the Deployment Assistant. The information that you enter in it is used to create the Recovery Appliance configuration files. These files work with installation utility to automate the installation process.

This chapter contains the following topics:

- [Overview of the Deployment Assistant](#) (page 5-2)
- [Generating the Configuration Files](#) (page 5-2)
- [About the Configuration Files](#) (page 5-4)
- [Customer Details Page](#) (page 5-5)
- [Hardware Selection Page](#) (page 5-5)
- [Define Customer Networks Page](#) (page 5-6)
- [Administration Network Page](#) (page 5-8)
- [Ingest Network Page](#) (page 5-10)
- [InfiniBand Network Configuration Page](#) (page 5-12)
- [Replication Network Page](#) (page 5-13)
- [Identify Compute Node OS Page](#) (page 5-14)
- [Review and Edit Page](#) (page 5-14)
- [Define Clusters Page](#) (page 5-14)
- [Cluster Page](#) (page 5-15)
- [Review and Edit Page](#) (page 5-18)
- [Replication Network SCAN Page](#) (page 5-18)
- [Cell Alerting Page](#) (page 5-19)
- [Platinum Configuration Page](#) (page 5-20)
- [Oracle Configuration Manager Page](#) (page 5-23)
- [Auto Service Request Page](#) (page 5-24)
- [Grid Control Agent Page](#) (page 5-25)
- [Tape Library Connectivity Page](#) (page 5-26)
- [Comments Page](#) (page 5-27)
- [Generate Page](#) (page 5-27)
- [Finish Page](#) (page 5-27)

5.1 Overview of the Deployment Assistant

Oracle Exadata Deployment Assistant can configure Oracle Exadata Database Machine, Recovery Appliance, or Oracle SuperCluster. You choose which system to configure on the Welcome page. After guiding you through a series of pages, the assistant generates a set of configuration files. These files are used with the Onecommand installation utility to help automate the deployment process and ensure that the system is configured to your specifications.

Before using Oracle Exadata Deployment Assistant, work with your network and database administrators to evaluate the current network settings, such as the current IP address use and network configuration. Then identify the settings to use for Recovery Appliance.

Ensure that you run Oracle Exadata Deployment Assistant, generate the configuration files, and validate the planned configuration before the delivery of Recovery Appliance. All information is required unless otherwise indicated.

The following buttons are available on pages throughout the assistant:

- **Help:** Displays information about completing the current page.
- **Import:** Loads the configuration settings previously saved in an XML file. Only select configurations created explicitly for Recovery Appliance. Do not import Oracle Exadata or Oracle SuperCluster configurations.

You can import a file to complete or modify the configuration settings, and when upgrading to a new software version.

This button is available only on the Welcome page.

- **Save:** Enables you to save your configuration in an XML file. You can import the file later to resume the configuration process. The file is saved by default in the `ExadataConfigurations` subdirectory.
- **Back:** Displays the previous page.
- **Next:** Displays the next page.
- **Cancel:** Cancel the configuration process and close the Deployment Assistant.

5.2 Generating the Configuration Files

Before running Oracle Exadata Deployment Assistant, ensure that you have all of the information needed to deploy Recovery Appliance at your site, as described in this chapter.

To generate the configuration files:

1. Log in to the computer where you plan to run Oracle Exadata Deployment Assistant. Ensure that the computer is on the network where Recovery Appliance will be deployed.

Oracle Exadata Deployment Assistant supports the following operating systems: Linux, Solaris 64-bit, SPARC Solaris 9, Apple Mac OS, and Microsoft Windows.

2. From My Oracle Support Doc ID 1927416.1, download the specified OneCommand ZIP file for your operating system.

The ZIP file contains all the components needed to configure and deploy the software. A separate ZIP file is provided for each supported operating system.

3. Extract the downloaded file, using an operating system utility. The following example extracts a Linux file:

```
$ unzip p19363943_121111_Linux-x86-64.zip
Archive:  p19363943_121111_Linux-x86-64.zip
  inflating: linux-x64/properties/patchname.properties
  inflating: linux-x64/properties/s_Sparc.properties
  inflating: linux-x64/properties/s_Solaris.properties
  inflating: linux-x64/properties/commonQuestions.txt
  inflating: linux-x64/properties/s_LinuxXen.properties
  inflating: linux-x64/properties/jaxb.properties
  inflating: linux-x64/properties/es.properties
  inflating: linux-x64/properties/s_Windows.properties
  inflating: linux-x64/properties/s_Linux.properties
  inflating: linux-x64/README.txt
  inflating: linux-x64/README
  inflating: linux-x64/config.sh
  inflating: linux-x64/install.sh
  inflating: linux-x64/jre1.7.0_25/man/man1/javaws.1
  inflating: linux-x64/jre1.7.0_25/man/man1/rmiregistry.1
  inflating: linux-x64/jre1.7.0_25/man/man1/policytool.1
  inflating: linux-x64/jre1.7.0_25/man/man1/pack200.1
  inflating: linux-x64/jre1.7.0_25/man/man1/java.1
  inflating: linux-x64/jre1.7.0_25/man/man1/orbd.1
  .
  .
  .
```

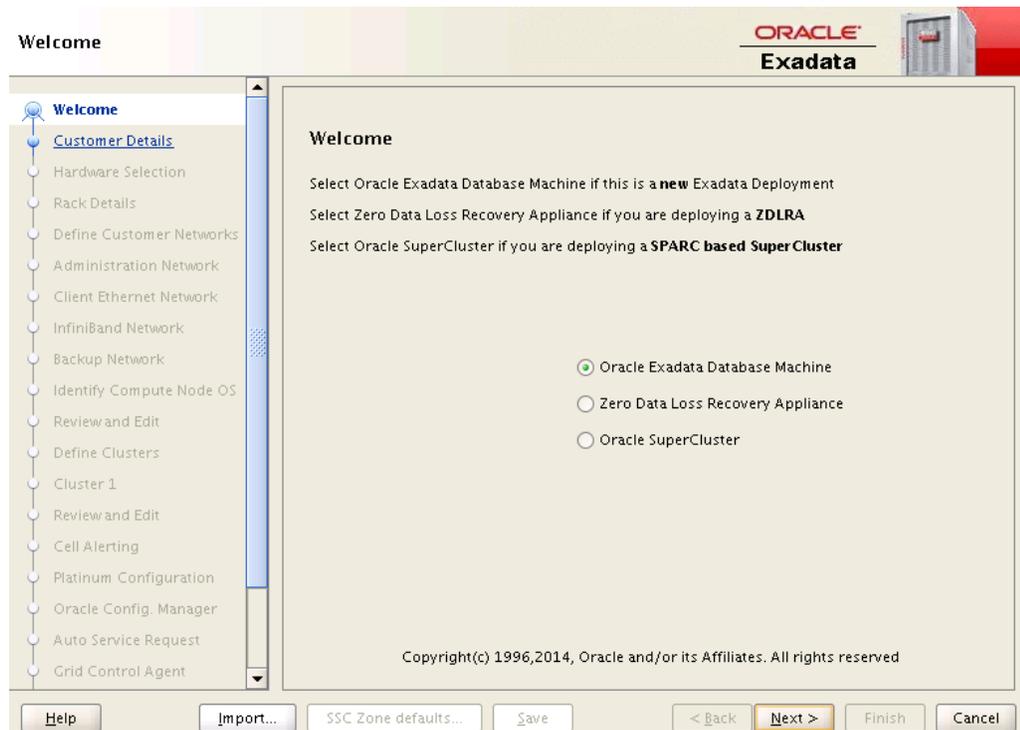
4. Change to the installation directory, such as `linux-x64`.
5. To start Oracle Exadata Deployment Assistant, run either `config.sh` or `config.cmd`, depending on your operating system.

The Welcome page is displayed. See [Figure 5-1](#) (page 5-4).

6. Select **Zero Data Loss Recovery Appliance**.
7. To use a previously saved configuration file, click **Import** and select the XML file. Configuration files are saved by default in the `ExadataConfigurations` subdirectory.
8. Follow the steps of the wizard. Click **Help** for details about your choices on each page.
9. On the Generate page, click **Create Files**.
10. Send the generated zip file to your Oracle representative. The file has a name in the form `customer_name-rack_name.zip`
11. Run the `checkip` script to validate the network settings. See "[About the Configuration Files](#) (page 5-4)".

[Figure 5-1](#) (page 5-4) shows the Oracle Exadata Deployment Assistant Welcome page.

Figure 5-1 Oracle Exadata Deployment Assistant Welcome Page



5.3 About the Configuration Files

Oracle Exadata Deployment Assistant generates the following files to use to configure the system. You can select the directory where they are saved.

customer_name-rack_name.xml

Contains all the settings for a completed configuration. To alter the configuration of a Recovery Appliance deployment, you can load this file, enter the changes, and regenerate the configuration files.

This file is only used by Oracle Exadata Deployment Assistant. It is not used for the actual configuration of Recovery Appliance.

customer_name-rack_name_temporary_config.xml

Contains the settings of an incomplete configuration. Use this file to continue from where you stopped working on the configuration.

customer_name-rack_name.zip

Contains copies of the configuration files. If an Oracle representative will perform the installation, then send this file to Oracle before the installation date. Otherwise, transfer the file to a USB drive for copying to Recovery Appliance.

customer_name-rack_name-checkip.sh or -checkip.cmd

Runs a series of tests. You use it at various points in the installation process:

- Before you connect the Cisco switch in Recovery Appliance to the data center network, to ensure that the factory IP addresses are unique.

- After you connect the Cisco switch to the data center network, to ensure that the custom names and IP addresses were added correctly to the DNS.
- After you install the Recovery Appliance software, to ensure that the deployment onto the data center network was successful.

customer_name-rack_name-InstallationTemplate.html

Provides a report of all the configuration details. You can view this report in a browser. Check it carefully to ensure that all of the settings are correct.

customer_name-rack_name-platinum.csv

Identifies the compute servers and storage servers covered by Oracle Platinum Services.

customer_name-rack_name-preconf_rack_n.csv

Contains the component names and IP addresses used by the `checkip` script.

5.4 Customer Details Page

The following table describes the customer details fields.

Table 5-1 Customer Details Page

Customer Details Field	Description
Customer Name	The customer name, which can contain alphanumeric characters and spaces. Required.
Application	The application used on Oracle Exadata Database Machine or another protected Oracle Database system.
Network Domain Name	The name of the domain in which Recovery Appliance will operate. Required.
Name Prefix	The prefix used to generate host names for all rack components. Required. The prefix can contain up to 20 alphanumeric characters. White space is invalid. For example, a value of <code>ra01</code> results in a compute server name of <code>ra01db01</code> , and a storage server name of <code>ra01cel01</code> .
Region	The geographic area where Recovery Appliance will be installed. Select a region from the list.
Time Zone	The time zone for the installation site. Select a time zone from the list. The displayed time zones are within the selected region.
DNS	Up to three IP addresses for the Domain Name System (DNS) servers.
NTP	Up to three IP addresses for the Network Time Protocol (NTP) servers.

5.5 Hardware Selection Page

The hardware selection page identifies the number of racks that you want to deploy at the same time, and the number of storage servers within each rack. Recovery Appliance supports a flexible configuration. From the a minimum configuration X5

through X7, you can add up to 15 storage servers, and up to 11 storage servers for X4. Multiple racks must be cabled together.

The following table describes the hardware choices.

Table 5-2 Hardware Selection Page

Hardware Selection Field	Description
Select interconnected hardware to deploy	Lists the available hardware configurations for a rack. Each rack can have 3 to 18 storage servers (cells), depending on your hardware version. Choose one or more racks. You can choose the same type of rack multiple times.
This is your deployment	Lists the hardware selected for your site.

For example, with X7 hardware, to deploy two full racks (18 storage servers) and one minimally configured rack (3 storage servers) at the same time:

1. Expand the respective X7 ZDLRA folder.
2. Add X7–2 ZDLRA Elastic Rack HC 10 TB three times (once for each rack), and click **Next**.

The Rack Details page appears.

3. For the first 2 racks, enter 18 in the Storage Cell count field.
4. For the third rack, enter 3 in the Storage Cell count field.

With X6 hardware, to deploy two full racks (18 storage servers) and one minimally configured rack (3 storage servers) at the same time:

1. Expand the respective X6 ZDLRA folder.
2. Add X6–2 ZDLRA Elastic Rack HC 8TB three times (once for each rack), and click **Next**.

The Rack Details page appears.

3. For the first 2 racks, enter 18 in the Storage Cell count field.
4. For the third rack, enter 3 in the Storage Cell count field.

With X4 hardware, to deploy two full racks (14 storage servers) and one minimally configured rack (3 storage servers) at the same time:

1. Expand the X4 ZDLRA folder.
2. Add X4-2 ZDLRA 14 Cell HC 4TB twice.
3. Add X4-2 ZDLRA 3 Cell HC 4TB once.

5.6 Define Customer Networks Page

The deployment process configures the following networks for Recovery Appliance. Each network must be on a separate subnet from the others:

- Subnet 1: The **administration network** provides access to the servers and switches in Recovery Appliance, either directly or through Oracle ILOM. This network uses copper cable. Required.

- Subnet 2: The **ingest network** connects Oracle Exadata Database Machine to a local Recovery Appliance. This network can be bonded, and it uses either a 1/10 GbE copper cable or a 10 GbE optical fiber cable. Required.

The default setting is Bonded, which supports active/passive bonding, but Oracle recommends LACP, which supports active/active bonding using an Link Aggregation Control Protocol (LACP) layers 3 and 4 (host and port). This configuration is the best for using both network interfaces in a High Availability installation. The data center must provide LACP support across two separate network switches using a technology like Cisco vPC, Arista MLAG, or Juniper MC-LAG.

- Subnet 3: The private **InfiniBand network** connects components within the rack and, optionally, multiple racks together. This network uses InfiniBand cables. Required.
- Subnet 4: The **replication network** connects a Recovery Appliance to an optional, remote Recovery Appliance. This network can be bonded, and it uses either a 1/10 GbE copper cable or a 10 GbE optical fiber cable, whichever type the ingest network does not use. Optional.

A replication network is also called a backup or DR network.

See Also:

- ["Overview of Network Requirements \(page 3-1\)"](#)
- ["Cabling Considerations \(page 5-8\)"](#)

5.6.1 Network Configuration Choices

The following table describes the configuration choices for the networks. The shaded fields and options are read only, and appear only to provide you with additional network details. Also see the previous network descriptions.

Table 5-3 Define Networks Page

Network Field	Description
Subnet Mask	The subnet mask to use for the network.
Bonded or Non Bonded	Identifies whether the network is bonded. The Admin and Private InfiniBand networks are always non bonded.
Enable LACP	Select this option for the ingest network if LACP is enabled in the data center.
Gateway	The IP address of the gateway for the network.
Network Format	Identifies the appropriate Ethernet card for the cable type: <ul style="list-style-type: none"> • 1/10 Gb Copper Base-T for copper cables • 10 GbE Optical for fiber optic cables
Available Network	Select this option to configure the optional Replication network.

5.6.2 Cabling Considerations

Each compute server has five ports, two optical and three copper. Most data centers support 10 GbE optical only, thus the ingest network is typically configured for 10 GbE optical.

Oracle recommends bonding of the ingest and replication networks for high availability and throughput. The default bonding is active/passive, which limits the network to 10 GbE or 1.2 GB/second. In contrast, the optional LACP configuration has a potential of 20 GbE or 2.4 GB/second for each compute server. For this throughput, you must configure the network switches appropriately for layer 3 (hashing over the source and destination IP addresses) and layer 4 (hashing over the TCP ports used by applications). Further, the Inter Switch Links (ISL) must have bandwidth equal to or greater than the total number of recovery appliance compute nodes.

The replication network is optional. If you choose the optical ports for the ingest network, then you must choose the copper ports for the replication network. High availability is less important for the replication network than the ingest network. Thus, you can configure the replication network with no bonding, and thus no high availability capabilities.

5.7 Administration Network Page

Each compute server and storage server has two network interfaces for the administration network:

- An Ethernet interface provides management access to the operating system.
- An ILOM Ethernet interface provides access to Oracle Integrated Lights Out Manager (Oracle ILOM).

The following table describes the administration network fields. The shaded fields are read only, and appear only to provide you with additional configuration details.

Table 5-4 Administration Network Page

Administration Network Field	Description
Starting IP Address for Pool	The first IP address in a pool of consecutive addresses assigned to the administration network port of the compute servers, storage servers, and InfiniBand switches. Addresses in this pool are also assigned to the ILOM port of the database and storage servers.
Pool Size	This value is calculated from the hardware selection. The pool should consist of consecutive IP addresses. If consecutive IP addresses are not available, then you can modify individual IP addresses on the review page.
Ending IP Address for Pool	This value is calculated from the starting IP address and the pool size.
Is the default gateway for database servers	Select this option if the gateway IP address is the default gateway for the compute servers. Oracle recommends using the ingest network gateway instead, because it has a bonded interface.

Table 5-4 (Cont.) Administration Network Page

Administration Network Field	Description
Defines the host name for the database servers	Select this option if the administration network names define the host names for the compute servers. If selected, the Oracle Management Server must have SQL*Net connectivity to the administration network.
Sample host names	Displays sample names on the administration network, using the prefix entered on the Customer Details page. To modify the name format, click Modify .

5.7.1 Modifying the Administration Host Name Format

You can change the default administration host names by using the Administration Network Format Masks dialog box. Names are typically changed to meet corporate naming standards or for compatibility with existing systems.

See "[Overview of Network Requirements](#) (page 3-1)" for a description of the administration network.

The following table describes the administration network format mask fields. Enter percent signs (%) in the names where you want the ID numbers to appear. The device ID numbering begins with the values in the Starting ID fields. For example, if you enter a 1 for the starting ID, and two percent signs (%%) in the name, then the devices are numbered 01, 02, 03, and so forth.

 **Note:**

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

Table 5-5 Administration Network Format Masks Dialog Box

Administration Network Format Masks Field	Description
Compute Node Name	The alphanumeric administration name of the compute servers. Enter percent signs for the ID number. Required.
Compute Node Starting ID	The number of the first compute server. This number replaces the percent signs in the name. Required.
Sample Name	Displays a compute server name created using the format mask.
Compute ILOM Name	The alphanumeric administration name of the compute server ILOMs. Enter percent signs for the ID number. Required.
Compute ILOM Starting ID	The number of the first compute server ILOM. This number replaces the percent signs in the name. Required.
Sample Name	Displays a compute server ILOM name created using the format mask.

Table 5-5 (Cont.) Administration Network Format Masks Dialog Box

Administration Network Format Masks Field	Description
Cell Node Name	The alphanumeric administration name for the storage servers. Required.
Cell Node Starting ID	The number of the first storage server. This number replaces the percent signs in the name. Required. The other storage servers are numbered sequentially.
Sample Name	Displays a storage server name created using the format mask.
Cell ILOM Name	The alphanumeric administration name for the storage server ILOMs. Required.
Cell ILOM Starting ID	The number of the first storage server ILOM. This number replaces the percent signs in the name. Required. The other storage servers are numbered sequentially.
Sample Name	Displays a storage server ILOM name created using the format mask.
Cisco Switch	The alphanumeric administration name for the Cisco Ethernet switch. Required.
KVM Switch	Not applicable. Recovery Appliance has no KVM switch.
PDU-A	The alphanumeric administration name for the first power distribution unit (PDU). Required.
PDU-B	The alphanumeric administration name for the second PDU. Required.
InfiniBand Spine	The alphanumeric administration name for the InfiniBand spine switch.
InfiniBand Leaf	The alphanumeric administration name for the first InfiniBand leaf switch. Required.
InfiniBand Leaf	The alphanumeric administration name for the second InfiniBand leaf switch. Required.

5.8 Ingest Network Page

The ingest network provides the connection between the protected databases and Recovery Appliance for bulk data transfers. It is the backup network.

See "[Overview of Network Requirements](#) (page 3-1)" for a description of the ingest network.

The following table describes the ingest network fields. The shaded fields are read only, and appear only to provide you with additional configuration details.

Table 5-6 Ingest Network Page

Ingest Network Field	Description
Starting IP Address for Pool	The first IP address in a pool of consecutive addresses assigned to the ingest network port of the compute servers.

Table 5-6 (Cont.) Ingest Network Page

Ingest Network Field	Description
Pool Size	This value is calculated from the hardware selection. The pool should consist of consecutive IP addresses. If consecutive IP addresses are not available, then you can modify individual IP addresses later in the configuration process.
Ending IP Address for Pool	This value is calculated from the starting IP address and the pool size.
Is the default gateway for compute servers	Select this option if the gateway IP address is the default gateway for the compute servers. Oracle recommends this configuration.
Defines the host name for the compute servers	Select this option if the ingest network names define the host names for the compute servers. If selected, the Oracle Management Server must have SQL*Net connectivity to the ingest network.
Sample ingest names	Displays sample names on the ingest network, using the prefix entered on the Customer Details page. To modify the format, click Modify .

5.8.1 Modifying the Ingest Name Format

You can change the default ingest network host names on the Ingest Network Format Masks dialog box.

The following table describes the ingest network format mask fields. Enter percent signs (%) in the names where you want the ID numbers to appear. The device ID numbering begins with the values in the Starting ID fields. For example, if you enter a 1 for the starting ID, and two percent signs (%%) in the name, then the devices are numbered 01, 02, 03, and so forth.

Table 5-7 Ingest Network Format Masks Dialog Box

Ingest Network Format Masks Field	Description
Ingest Network Name	The alphanumeric ingest name for the rack. Enter percent signs where the ID number appears. Required.
Ingest Network Starting ID	The first numeric identifier for an ingest name in the rack. Required.
VIP Name	The alphanumeric virtual IP name for the rack. Enter percent signs where the ID number appears. Required.
VIP Starting ID	The first numeric identifier for a virtual IP name. Required.
Sample Name	Displays a name created using the format mask.
SCAN Name	The alphanumeric single client access network (SCAN) name for the rack. Required.

5.9 InfiniBand Network Configuration Page

The private InfiniBand network connects the compute servers and storage servers through the InfiniBand gateway switches in the rack. The InfiniBand network can also connect multiple Recovery Appliance racks.

See "[Overview of Network Requirements](#) (page 3-1)" for a description of the InfiniBand network.

The following table describes the InfiniBand network fields. The shaded fields are read only, and appear only to provide you with additional configuration details.

Table 5-8 InfiniBand Network Page

InfiniBand Network Field	Description
Starting IP Address for Pool	The first IP address in a pool of consecutive addresses assigned to the InfiniBand network port of the compute servers, storage servers, and InfiniBand switches.
Pool Size	This value is calculated from the hardware selection. The pool should consist of consecutive IP addresses. If consecutive IP addresses are not available, then you can modify individual IP addresses on the review page.
Ending IP Address for Pool	This value is calculated from the starting IP address and the pool size.
Enable active bonding on compute node network	Select this option to support failover between the two compute servers on the InfiniBand network.
Sample InfiniBand names	Displays sample compute server and storage server names on the InfiniBand network, using the prefix entered on the Customer Details page. To modify the name format, click Modify .

5.9.1 Modifying the InfiniBand Host Name Format

You can change the format of the InfiniBand network host names by using the Private (IB) Network Format Masks dialog box.

The following table describes the InfiniBand network format mask fields. Enter percent signs (%) in the names where you want the ID numbers to appear. The device ID numbering begins with the values in the Starting ID fields. For example, if you enter a 1 for the starting ID, and two percent signs (%%) in the name, then the devices are numbered 01, 02, 03, and so forth.

Table 5-9 Private (IB) Network Format Masks Dialog Box

Private (IB) Network Format Masks Field	Description
Compute Private Name	The alphanumeric InfiniBand access name for the compute servers. Enter percent signs where the ID number appears. Required.

Table 5-9 (Cont.) Private (IB) Network Format Masks Dialog Box

Private (IB) Network Format Masks Field	Description
Compute Private Starting ID	The number of the first compute server in the rack. This number replaces the percent signs in the host name. Required. The other compute servers are numbered sequentially.
Cell Private Name	The alphanumeric InfiniBand name for storage servers. Enter percent signs where the ID number appears. Required.
Cell Private Starting ID	The number of the first storage server in the rack. This number replaces the percent signs in the host name. Required. The other storage servers are numbered sequentially.
Sample Name	Displays a name created using the format mask.

5.10 Replication Network Page

The optional replication network connects Recovery Appliance to another (downstream) Recovery Appliance. It provides dedicated bandwidth for this task. Complete this page if you have a replication network.

If a replication network is unavailable, then you can use the ingest network to connect to a downstream appliance. However, this configuration creates competition for bandwidth on the ingest network.

See "[Overview of Network Requirements](#) (page 3-1)" for a description of the replication network.

The following table describes the replication network fields. The shaded fields are read only, and appear only to provide you with additional configuration details.

Table 5-10 Replication Network Page

Replication Network Field	Description
Starting IP Address for Pool	The first IP address in a pool of consecutive addresses assigned to the replication network port of the storage servers. The pool size identifies the required number of IP addresses. If consecutive IP addresses are not available, then you can modify individual IP addresses on the review page.
Pool Size	This value is calculated from the hardware selection.
Ending IP Address for Pool	This value is calculated from the starting IP address and the pool size.
Sample replication name	Displays a sample compute server name on the replication network, using the prefix entered on the Customer Details page. To modify the format, click Modify .

5.10.1 Modifying the Replication Host Name Format

You can change the replication host name format by using the Replication Network Masks dialog box.

The following table describes the replication network format mask fields. Enter percent signs (%) in the names where you want the ID numbers to appear. The device ID numbering begins with the values in the Starting ID fields. For example, if you enter a 1 for the starting ID, and two percent signs (%%) in the name, then the devices are numbered 01, 02, 03, and so forth.

Table 5-11 Replication Network Format Masks Page

Replication Network Format Masks Field	Description
Name	The alphanumeric host name of the compute servers on the replication network. Enter percent signs where the ID number appears. Required.
Starting ID	The number of the first compute server in the rack. This number replaces the percent signs in the host name. Required.
Sample Name	Displays a name created using the format mask.

5.11 Identify Compute Node OS Page

Recovery Appliance runs the Linux operating system. There are no choices on this page.

5.12 Review and Edit Page

Use this page to review the network names and IP addresses assigned to each device. If you identified a pool of IP addresses that are not all available for use, then replace the duplicate addresses on this page. You can edit the host names and IP addresses just by clicking the field. All host names and IP addresses must be unique.

Regenerate Data: Click this button to see the current list of generated names, after modifying the names or format masks.

5.13 Define Clusters Page

You can create one cluster on Recovery Appliance.

The following table describes the fields on this page.

Table 5-12 Define Clusters Page

Define Clusters Field	Description
Cluster Name	A unique name for the cluster. Required. The name must start with an alphanumeric character. It can contain additional letters, numbers, periods (.), and hyphens (-). The default name is cluster-clu1.
<i>Available</i>	Lists all compute servers and storage servers included in this installation.
<i>Selected</i>	Lists all compute servers and storage servers that are assigned to the cluster. Click All to assign all servers to the cluster.

5.14 Cluster Page

Many fields on this page have values that you provided on earlier configuration pages. You can review these values and revise them if necessary. The shaded fields are read only, and appear only to provide you with additional details about the installation.

- [Cluster](#) (page 5-15)
- [Users and Groups](#) (page 5-16)
- [Software Locations](#) (page 5-16)
- [Disk Group Details](#) (page 5-16)
- [Initial Database](#) (page 5-16)
- [Ingest Network](#) (page 5-17)
- [Replication Network](#) (page 5-17)

5.14.1 Cluster

The values at the top of the page are derived from the rack configuration settings on the [Customer Details Page](#) (page 5-5).

The following table describes these general cluster details.

Table 5-13 Cluster Details

Cluster Field	Description
Cluster Name	A unique name for the cluster. Required. The name must start with an alphanumeric character. It can contain additional letters, numbers, periods (.), and hyphens (-). The default name is cluster-clu1.
Prefix	The prefix used to generate host names for all rack components. Required. The prefix can contain up to 20 alphanumeric characters. White space is invalid. The default prefix is ra01. For example, a value of ra01 results in a compute server name of ra01db01, and a storage server name of ra01cel01.
DNS	Up to three IP addresses for the Domain Name System (DNS) servers.
NTP	Up to three IP addresses for the Network Time Protocol (NTP) servers.
Domain Name	The name of the domain in which Recovery Appliance will operate. Required.
Region	The geographic area where Recovery Appliance will be installed. Select a region from the list.
Time Zone	The time zone for the installation site. Select a time zone from the list. The displayed time zones are within the selected region.

5.14.2 Users and Groups

This section is read only. It displays the users and groups defined for Oracle Database on Recovery Appliance. You do not access this database directly, because it exclusively supports the backup and recovery of protected databases.

5.14.3 Software Locations

This section identifies the location of the Oracle Database software on Recovery Appliance. Most fields are read only.

The following table describes the editable field.

Table 5-14 Software Locations

Software Locations Field	Description
Software Languages	The two-letter abbreviations for the languages supported in the database. The default language is English (en).

5.14.4 Disk Group Details

This section defines the disk groups for Oracle Database on Recovery Appliance. Most fields are read only.

The following table describes the editable field.

Table 5-15 Disk Group Details

Disk Group Field	Description
DELTA Disk Group	The type of redundancy for the DELTA disk group. <code>NORMAL</code> is the correct setting for backups. (Oracle recommends <code>HIGH</code> only for mission-critical applications.)

5.14.5 Initial Database

The following table describes the editable field for the initial database.

Table 5-16 Initial Database

Initial Database Field	Description
Database Name	<p>The name of the database on Recovery Appliance.</p> <div style="border: 1px solid #0070C0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <p>This user-specified database name (the <code>DB_NAME</code>) must be different from any other Recovery Appliance or general database within your enterprise.</p> </div>

5.14.6 Ingest Network

Review the values for the ingest network, and correct them as necessary. The values are primarily derived from the [Ingest Network Page](#) (page 5-10).

The following table describes the ingest network fields.

Table 5-17 Ingest Network Fields

Ingest Network Field	Description
Domain	The name of the domain in which Recovery Appliance will operate.
Start IP	The first IP address in a pool of consecutive addresses assigned to the ingest network port of the compute servers.
Pool Size	The number of IP addresses required for the servers on the ingest network.
Gateway IP	The IP address of the gateway for the ingest network.
Name Mask	The name mask used to generate server names on the ingest network. Percent signs (%) in the name indicate where the ID numbers for the individual servers appear.
Name Mask Start ID	The ID number for the first server on the ingest network. This number replaces the percent signs in the name mask.
VIP Name Mask	The alphanumeric virtual IP name for the rack. Enter percent signs where the ID number appears.
VIP Name Mask Start ID	The first numeric identifier for a virtual IP name.
SCAN Name	The alphanumeric single client access network (SCAN) name for the rack.

5.14.7 Replication Network

Review the values for the replication network, and correct them as necessary. The values are primarily derived from the [Replication Network Page](#) (page 5-13).

The following table describes the fields.

Table 5-18 Replication Network Fields

Replication Network Field	Description
Base Adapter	Indicates whether a replication network is being configured for Recovery Appliance. Choose Replication if it is, or Not in use otherwise.
Domain	The name of the domain in which Recovery Appliance will operate.
Start IP	The first IP address in a pool of consecutive addresses assigned to the replication network port of the storage servers.
Pool Size	The number of IP addresses required for the servers on the replication network.
Name Mask	The name mask used to generate server names on the replication network. Percent signs (%) in the name indicate where the ID numbers for the individual servers appear.
Start ID	The ID number assigned to the first server on the replication network. This number replaces the percent signs in the name mask

5.15 Review and Edit Page

Use this page to review the network names and IP addresses on the single client access network (SCAN), ingest network, virtual IP (VIP) network, and replication network. If you identified a pool of IP addresses that are not all available for use, then replace the duplicate addresses on this page. You can edit the host names and IP addresses just by clicking the field. All host names and IP addresses must be unique.

Regenerate Data: Click this button to update the list of generated names, if you modified any names or format masks.

Lookup IP: Checks that the IP address matches the device name, and corrects it if they do not match.

5.16 Replication Network SCAN Page

This page appears only if you previously configured a replication network. Recovery Appliance defines a single client access network (SCAN) for the rack, and a virtual IP (VIP) for the compute servers. All names and IP addresses must be unique.

Regenerate Data: Click this button to update the list of generated names, if you modified any names or format masks.

Lookup IP: Checks that the IP address matches the device name, and corrects it if they do not match.

The following table describes the replication network SCAN fields.

Table 5-19 Replication Network SCAN Page

Replication Network SCAN Field	Description
Replication SCAN Name	The name of the replication single client access network (SCAN) for Recovery Appliance.
Replication SCAN IP	The IP addresses of the replication network.
Nodex Replication VIP Name	The VIP name of Recovery Appliance.
VIP IP	The IP address of the replication VIP.

5.17 Cell Alerting Page

Cell alerts are triggered when a problem is detected in a storage server. The alert can be delivered using simple mail transfer protocol (SMTP), simple network management protocol (SNMP), or both.

Recovery Appliance does not require cell alerting, but you must configure one of the following alerting options:

- Auto Service Request
- Cell alerting
- Oracle Configuration Manager

The following table describes the cell alert configuration options.

Table 5-20 Cell Alerting Page

Cell Alerting Field	Description
Enable Email Alerting	Select this option to send cell alerts automatically. Then you can enter the configuration details.
Recipients Addresses	The email addresses for the alert recipients. Click this field to display a dialog box, where you can enter multiple addresses, one per line. When you are done, verify that the number of email addresses is correct in the adjacent display.
Name	The SMTP email user name that appears in the alert notifications, such as the Oracle backup administrator.
Email Address	The SMTP email address that sends the alert notifications, such as admin@example.com.
SMTP Server	The SMTP email server used to send alert notifications, such as mail.example.com.
Uses SSL	Select this option to use Secure Socket Layer (SSL) security when sending alert notifications.
Port	The SMTP email server port used to send alert notifications, such as 25 or 465.
Enable SNMP Alerting	Select this option if alerts will be delivered using SNMP.
SNMP Server	The host name of the SNMP server, such as snmp.example.com.
Port	The port for the SNMP server. The default port is 162.

Table 5-20 (Cont.) Cell Alerting Page

Cell Alerting Field	Description
Community	The SNMP community name. The default is <code>public</code> .

5.18 Platinum Configuration Page

Oracle Platinum Services provide enhanced support at no additional cost for qualified engineered systems. This is a connected service and requires the installation of Oracle Advanced Support Gateway software within the system.

The Oracle Advanced Support Gateway is configured with the user information that is required by the Platinum Service monitoring agent.

The following sections describe the Oracle Platinum Service configuration options:

- [Customer Details](#) (page 5-20)
- [Platinum Gateway Details](#) (page 5-21)
- [Platinum Agent Configuration Details](#) (page 5-22)



See Also:

- Oracle Platinum Services information at <http://www.oracle.com/us/support/premier/engineered-systems-solutions/platinum-services/overview/index.html>
- Oracle Advanced Support Gateway Hardware Minimum Requirements at <http://www.oracle.com/us/support/library/advanced-support-gateway-host-reqs-1896462.pdf>

5.18.1 Customer Details

The following table describes the fields for the customer details section of the Platinum Configuration page.

Table 5-21 Platinum Customer Details

Platinum Field	Description
Capture data for Platinum configuration	Select this option to enter the Oracle Platinum Services configuration details. Clear the option if you do not want these services.
Customer Name	The customer name must match the name in the purchase order for Recovery Appliance. Check the Printed Purchase Orders Report (Portrait).
CSI	The hardware customer support identifier (CSI) for the Recovery Appliance rack.

Table 5-21 (Cont.) Platinum Customer Details

Platinum Field	Description
My Oracle Support Email	The My Oracle Support email address that corresponds to the CSI.

5.18.2 Platinum Gateway Details

Recovery Appliance can use an existing Oracle Advanced Support Gateway in the network. Otherwise, you can create a gateway on a server or a supported virtual machine, which is connected to the same management network or ingest network as Recovery Appliance.

Table 5-22 Platinum Gateway Details

Platinum Field	Description
Use Existing Gateway	Select this option to use an existing Oracle Advanced Support Gateway. Then you only need to identify the host name and primary IP address. Clear the box to create a gateway and provide the full configuration details.
Gateway Machine Type	Identifies the gateway server as a physical device or a virtual machine.
Gateway Machine Description	The hardware model where the gateway system will be installed, such as Oracle Advanced Support Gateway Server X4-2.
Platinum Gateway Host Name	The Oracle Advanced Support Gateway host name. The name format is <code>ct-customer_name-nn</code> .
Primary IP Address	The primary IP address for Oracle Advanced Support Gateway.
Subnet Mask	The subnet mask for the Oracle Advanced Support Gateway primary network.
Gateway IP Address	The Oracle ILOM primary subnet gateway server.
VPN	The VPN connection type for access to the gateway from Oracle Continuously Connected Network. Choose between Secure Socket Layer (SSL) and Internet Protocol Security (IPSEC). Oracle recommends using Secure Socket Layer (SSL)
Gateway to Exadata Link	The Oracle Advanced Support Gateway location in the network. Choose between DMZ (demilitarized zone or perimeter network) and Direct. Oracle recommends DMZ.
Static Routes	The number of defined static routes. Click Modify to change the routes for the gateway.

Table 5-22 (Cont.) Platinum Gateway Details

Platinum Field	Description
Modify	Displays the Create Static Routes dialog box, where you can create or delete static routes. In the dialog box, click the plus sign (+) to create a route. Click the minus sign (-) next to a route to delete it. Destination Network: The remote network for the router. Use classless inter-domain routing (CIDR) format, such as 192.0.2.1/255. Gateway IP: The remote network gateway IP address. It must be in the Oracle Advanced Support Gateway primary IP subnet.
Gateway machine has an ILOM	Select this option to install Oracle Advanced Support Gateway on Oracle hardware that has an ILOM service processor. Then you can edit the configuration fields for Oracle ILOM. Clear this option to install Oracle Advanced Support Gateway on a third-party server.
ILOM IP Address	The Oracle ILOM IP address of the gateway server.
Subnet Mask	The Oracle ILOM subnet mask of the gateway server. Select a value from the list.
Gateway IP Address	The IP address of the gateway server.
HTTP Proxy Required	Select this option if Oracle Advanced Support Gateway must use a proxy to connect to the required Oracle sites over the Internet. This choice enables the other proxy fields.
Proxy Host Name	The proxy host name.
Port	The proxy port number.
HTTP Proxy Requires Authentication	Select this option if the proxy server requires authentication. Then you can provide the user name.
Proxy User Name	A user name recognized by the proxy server.

5.18.3 Platinum Agent Configuration Details

Oracle Platinum Services uses an agent to monitor the health of Recovery Appliance. Use the default configuration settings for the agent, unless you have a specific reason to change them. For example, the default port number might be assigned already to another device.

The following table describes the Platinum agent fields.

Table 5-23 Platinum Agent Configuration Details

Platinum Field	Description
Agent OS User Name	The operating system user for the Platinum agent. The default name is <code>oracom</code> .
Agent OS User ID	The user ID. The default ID is 1010.

Table 5-23 (Cont.) Platinum Agent Configuration Details

Platinum Field	Description
Allow Agent sudo Privileges	Grants the operating system user <code>sudo</code> privileges. Note: Selecting this option enables Oracle service engineers to collect logs and other information for a service request.
Agent OS Group Name	The operating system group for the user. The default name is <code>orarom</code> . The Platinum agent can read the Linux messages file as a member of this group.
Agent OS Group ID	The group ID for the user. The default ID is 1010.
Agent OS User Home	The home directory for the Platinum agent user.
Agent Software Home	The home directory for the agent software.
Agent Port	The port number for the Platinum agent. The default port is 1830.
SNMP Community String	The Simple Network Management Protocol (SNMP) trap community string. The default value is <code>public</code> .

5.19 Oracle Configuration Manager Page

Oracle Configuration Manager collects and uploads configuration information to an Oracle repository, for access when needed by Oracle Support.

Recovery Appliance does not require Oracle Configuration Manager, but you must configure one of the following alerting options:

- Auto Service Request
- Cell alerting
- Oracle Configuration Manager
- Oracle Platinum Services

The following table describes the Oracle Configuration Manager options.

Table 5-24 Oracle Configuration Manager Page

Oracle Configuration Manager Field	Description
Enable Oracle Configuration Manager	Select this option to enable Oracle Configuration Manager to collect configuration information from your Recovery Appliance. Then you can enter the configuration details.
Receive updates via MOS	Select this option to receive My Oracle Support updates automatically.
MOS Email Address	The address to receive updates from My Oracle Support.

Table 5-24 (Cont.) Oracle Configuration Manager Page

Oracle Configuration Manager Field	Description
Access Oracle Configuration Manager via Support Hub	Select this option to enable Support Hub access to Oracle Configuration Manager. Oracle Support Hub enables Oracle Configuration Manager instances to connect to a single internal port (the Support Hub) to upload configuration data. When Support Hub is used, the individual Oracle Configuration Manager instances in the compute servers do not require Internet access.
Support Hub Host Name	The host name for the Support Hub server. See Also: <i>Oracle Configuration Manager Companion Distribution Guide</i>
Hub User Name	The operating system user name for the Support Hub server.
HTTP Proxy used in upload to Oracle Configuration Manager	Select this option if you use an HTTP proxy to upload configuration information to the Oracle repository. Then you can enter the proxy information.
HTTP Proxy Host	The proxy server to connect to Oracle. The proxy can be between the following: <ul style="list-style-type: none"> • Compute servers and Oracle¹ • Compute servers and Support Hub² • Support Hub and Oracle Example: <code>[proxy_user@]proxy_host[:proxy_port]</code> The <code>proxy_host</code> and <code>proxy_port</code> entries are optional.
Proxy Port	The port number for the HTTP proxy server. The default is 80.
HTTP Proxy requires authentication	Select this option if the HTTP proxy requires authentication.
HTTP Proxy User	The user name for the HTTP proxy server. If the connection requires a password, then you must provide it during installation.

¹ Applicable when you do not have Oracle Support Hub.

² Applicable when you only have Oracle Support Hub.



See Also:

Oracle support policies at

<http://www.oracle.com/us/support/policies/index.html>

5.20 Auto Service Request Page

Auto Service Request (ASR) automatically opens service requests when specific hardware faults occur.

Recovery Appliance does not require Auto Service Request, but you must configure one of the following alerting options:

- Auto Service Request
- Cell alerting
- Oracle Configuration Manager
- Oracle Platinum Services

The following table describes the ASR configuration options.

Table 5-25 Auto Service Request Page

Auto Service Request Field	Description
Enable Auto Service Request	Select this option to use ASR with Recovery Appliance. Then you can enter the configuration details.
ASR Manager Host Name	The host name of the server where ASR Manager is installed. Note: Oracle recommends a standalone server that has connectivity to Recovery Appliance.
ASR Technical Contact Name	The name of the technical contact person for Recovery Appliance.
Technical Contact Email	The email address of the technical contact person.
My Oracle Support Account Name	The email address of the My Oracle Support account.
HTTP Proxy used in upload to ASR	Select this option if you use an HTTP proxy to upload to the Oracle ASR site. Then you can enter the proxy details.
HTTP Proxy Host	The host name of the proxy server.
Proxy Port	The port number for the HTTP proxy server. The default is 80.
HTTP Proxy requires authentication	Select this option if the HTTP proxy server requires authentication.
HTTP Proxy User	The user name for the proxy server. If the connection requires a password, then you must provide it during installation.



See Also:

[Setting Up Auto Service Request \(page 4-1\)](#)

5.21 Grid Control Agent Page

Oracle Enterprise Manager Cloud Control is the primary interface for administrators to use for creating, scheduling, and monitoring backups. When Oracle Enterprise Manager Grid Control agents are installed on Recovery Appliance, they can send information to the Oracle Management Server. Administrators can view this data from a Cloud Control console.

The following table describes the Oracle Enterprise Manager Grid Control agent options.

Table 5-26 Grid Control Agent Page

Grid Control Field	Description
Enable Oracle Enterprise Manager Grid Control Agent	Select this option to enable Oracle Enterprise Manager Grid Control agent to monitor Recovery Appliance. Then you can enter the configuration details.
EM Home Base Location	The Enterprise Manager base directory for the agents installed on the Recovery Appliance compute servers.
OMS Host Name	The host name where the Oracle Management Server repository is installed.
OMS HTTPS Upload Port	The port number that the Oracle Enterprise Manager Grid Control agent uses to upload information. The default is 4900.

5.22 Tape Library Connectivity Page

You can configure the tape library for Recovery Appliance only if you use Oracle Secure Backup as the media manager for tape backups. You can manually configure other media managers as clients only.

See "[About Tape Backup Infrastructure](#) (page 1-8)".

The following table describes the tape library configuration options.

Table 5-27 Tape Library Connectivity Page

Tape Library Connectivity Field	Description
Use Tape	Select this option if you plan to use Oracle Secure Backup as the media manager for tape backups. Then you can enter the configuration details.
Number of Tape Libraries	The number of tape libraries available to Recovery Appliance. All tape libraries must be hosted in the same domain. A separate subpage is created for each library, so that you can enter the configuration details.
Enable Admin VIP	Select this option to create a virtual IP (VIP) on the ingest network and enable Oracle Secure Backup as a Cluster Ready Services (CRS) service. Then you can enter the VIP details.
VIP Name	The alphanumeric name of the VIP.
VIP IP	The IP address of the VIP.

Table 5-27 (Cont.) Tape Library Connectivity Page

Tape Library Connectivity Field	Description
Library <i>n</i>	Select the appropriate software for each tape library: <ul style="list-style-type: none"><li data-bbox="776 405 1445 531">• ACSLs ACS API: The Automated Cartridge System Library Software (ACSLs) API has these fields:<ul style="list-style-type: none"><li data-bbox="816 468 1445 499">ACSLs Server: The name of the server running ACSLS<li data-bbox="816 506 1445 531">ACS ID: The ACS identification number<li data-bbox="776 537 1445 663">• ACSLs FC-SCSI: The Automated Cartridge System Library Software (ACSLs) FC-SCSI has these fields:<ul style="list-style-type: none"><li data-bbox="816 600 1445 632">ACSLs Server: The name of the server running ACSLS<li data-bbox="816 638 1445 663">Logical Library ID: The logical library identification number<li data-bbox="776 669 1445 716">• SCSI: The Simple SAN tape library has no configuration fields. Default.

5.23 Comments Page

Enter any additional information about the deployment in the text field, particularly anything of help to the person performing the installation. The information appears at the bottom of `InstallationTemplate.html`, which is one of the generated files.

Click **Reset List** to clear all comments.

5.24 Generate Page

Click **Next** to save the configuration files in your choice of locations. The default directory is named `ExadataConfigurations` in the current directory.

5.25 Finish Page

A summary page identifies the directory location and lists the names of the generated files. It also provides a link to the Installation Template.

Click **Finish** to close the utility.

Part II

Installation and Configuration

Part II provides information about installation and configuration of Recovery Appliance after it has arrived on site. This part contains the following chapters:

- [Installing the Recovery Appliance at the Site](#) (page 6-1)
- [Expanding a Recovery Appliance Rack with Additional Storage Servers](#) (page 7-1)
- [Extending Recovery Appliance with Multiple Racks](#) (page 8-1)
- [Configuring Ingest Over InfiniBand](#) (page 9-1)
- [Configuring a Recovery Appliance Rack](#) (page 10-1)
- [Installing the Recovery Appliance Software](#) (page 11-1)

6

Installing the Recovery Appliance at the Site

This chapter describes how to move and install the hardware, from unpacking the Recovery Appliance to powering it on. This chapter contains these sections:

- [Reviewing Safety Guidelines](#) (page 6-1)
- [Unpacking ZDLRA Rack](#) (page 6-2)
- [Placing ZDLRA Rack in Its Allocated Space](#) (page 6-5)
- [Installing the Fiber Channel Adapters](#) (page 6-11)
- [Acclimating the Recovery Appliance](#) (page 6-12)
- [Powering on the System for the First Time](#) (page 6-12)

6.1 Reviewing Safety Guidelines

Before ZDLRA 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 ZDLRA Rack in a location that is exposed to direct sunlight or near a device that may become hot.
- Do not install ZDLRA Rack in a location that is exposed to excessive dust, corrosive gases, or air with high salt concentrations.
- Do not install ZDLRA Rack in a location that is exposed to frequent vibrations. Install ZDLRA 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 ZDLRA Rack is used exclusively for ZDLRA Rack. Also be sure to observe the precautions, warnings, and notes about handling that appear on labels on the equipment.
- Do not place cables under the equipment or stretch the cables too tightly.
- Do not disconnect power cords from the equipment while its power is on.
- If you cannot reach the connector lock when disconnecting LAN cables, then press the connector lock with a flathead screwdriver to disconnect the cable. You could damage the system board if you force your fingers into the gap rather than using a flathead screwdriver.

- Do not place anything on top of ZDLRA 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 ZDLRA Rack. Allow for a sufficient warm-up period prior to server operation.
- Do not install ZDLRA 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 ZDLRA Rack can cause malfunctions. Static electricity is often generated on carpets.
- Confirm the supply voltage and frequency match the electrical ratings indicated for ZDLRA Rack.
- Do not insert anything into any ZDLRA Rack opening, unless doing so is part of a documented procedure. ZDLRA Rack contains high-voltage parts. If a metal object or other electrically-conductive object enters an opening in ZDLRA Rack, then it could cause a short circuit. This could result in personal injury, fire, electric shock, and equipment damage.
- When using single phase power distribution units (PDUs), note the following:
 - PDU A input 0 and PDU B input 2 must be on the same phase.
 - PDU A input 1 and PDU B input 1 must be on the same phase.
 - PDU A input 2 and PDU B input 0 must be on the same phase.

The inputs are labeled where they come out of the PDU. Connecting cables as described ensures the phases are balanced on both sides, A and B, in case of a failover.

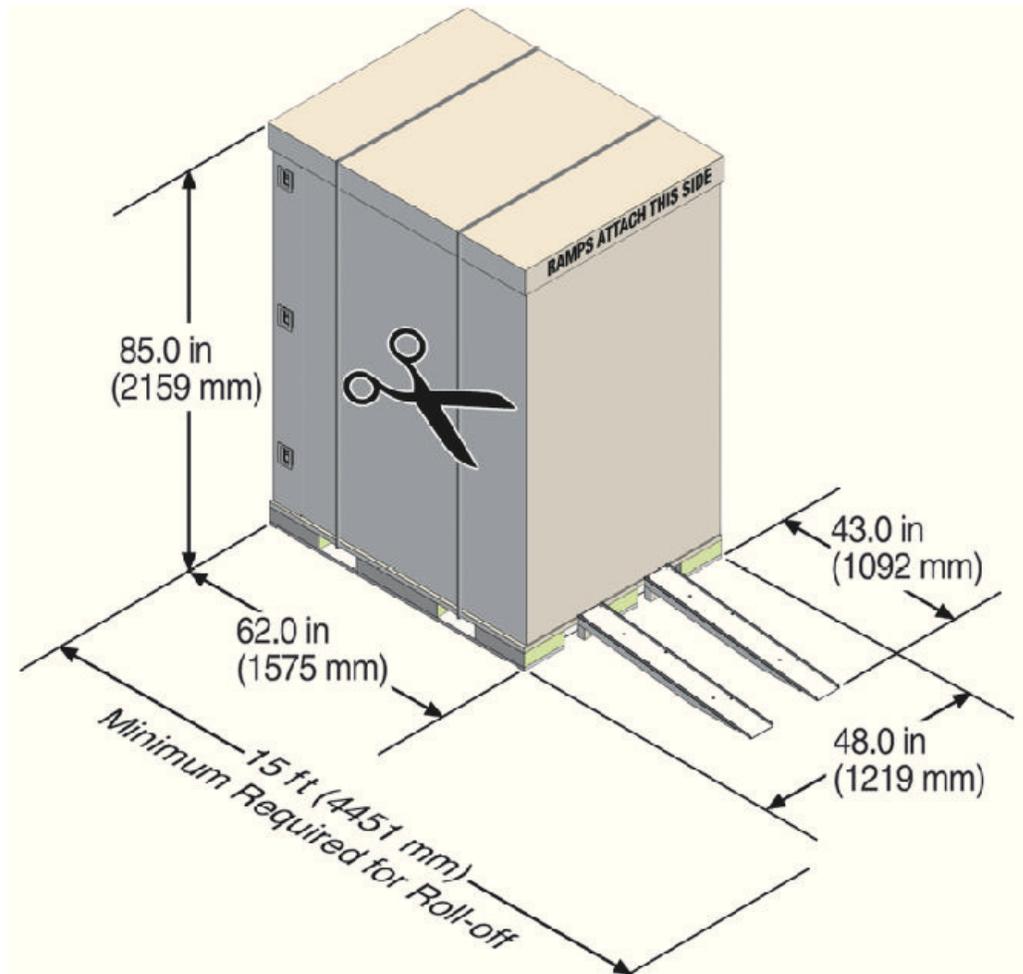
 **See Also:**

- *Important Safety Information for Sun Hardware Systems (816-7190)* included with the rack
- *Oracle Engineered System Safety and Compliance Guide, Compliance Model No.: ESY27*

6.2 Unpacking ZDLRA Rack

The unpacking location should be determined during the site planning process. The following image shows the ZDLRA Rack shipping crate.

Figure 6-1 ZDLRA Rack in Shipping Crate



6.2.1 Contents of the Shipping Kit

The ZDLRA Rack shipping kit contains the following tools and equipment to install and service the Recovery Appliance rack:

- 16 mm long No. 2 Phillips screw
- T30 Torx cranked wrench key
- T25 Torx cranked wrench key
- 6 mm hexagon Allen wrench key
- SW 12 mm single-headed wrench
- 2 square jumper brackets with 4 M5 Torx screws
- 2 cable management hooks with 4 spring nuts
- Side panel removal tool
- Keys to the front door, rear door, and side panel locks
- 32 M6 cage nuts

- 32 M6 screws
- Cage nut mounting tool
- SW 17 mm single-headed wrench is included on the shipping pallet

 **Note:**

The following items are not included in the shipping kit:

- No. 2 Phillips screwdriver
- Antistatic wrist strap

6.2.2 Removing ZDLRA Rack from the Shipping Crate

 **Note:**

Oracle strongly recommends that you use professional movers when unpacking and installing ZDLRA Rack.

The following procedure describes how to unpack ZDLRA Rack:

1. Unpack ZDLRA Rack carefully from the packaging and shipping pallet as follows:

 **Note:**

Before unpacking the rack from the shipping carton, refer to the labels on the carton and to the instructions that they provide.

- a. Remove the shipping carton bands.
- b. Remove the yellow fasteners and carton top.
- c. Remove the carton sides and inner top.

 **WARNING:**

Rocking or tilting the rack can cause the rack to fall over and cause serious injury or death.

2. Remove the shipping kit.
3. Attach the ramps to the shipping pallet as follows:
 - a. Remove the ramps from the pallet sides.
 - b. Obtain the parts bag from inside the cabinet.

- c. Adjust the leveling bolts on the ramps and connect the ramps to the pallet wheel track.
 - d. Verify that leveling feet are up before moving rack.
 4. Carefully roll ZDLRA Rack off the shipping pallet as follows:
 - a. Unfasten the exterior shipping brackets from the pallet.
 - b. Unfasten the interior shipping brackets from the pallet.

 **Note:**

Use care when removing the shipping brackets from underneath ZDLRA Rack. Access to the inside shipping brackets might be limited.

- c. Roll the cabinet down the ramps to the level floor. Oracle recommends having three people available to move the rack down the ramp: two people on both sides to help guide the rack and one person in back.
 5. Save the shipping brackets used to secure the rack to the shipping pallet. Do not dispose of these brackets, because you cannot order replacement brackets.
 6. Recycle the packaging properly. Follow local laws and guidelines to dispose of the material.

 **See Also:**

- *Sun Rack II User's Guide* at <https://docs.oracle.com/cd/E19657-01/html/E29153/index.html>
- *Oracle Rack Cabinet 1242 User's Guide* at http://docs.oracle.com/cd/E85660_01/html/E87280/index.html for X7 systems

6.3 Placing ZDLRA Rack in Its Allocated Space

This section describes how to position, stabilize, and ground ZDLRA Rack.

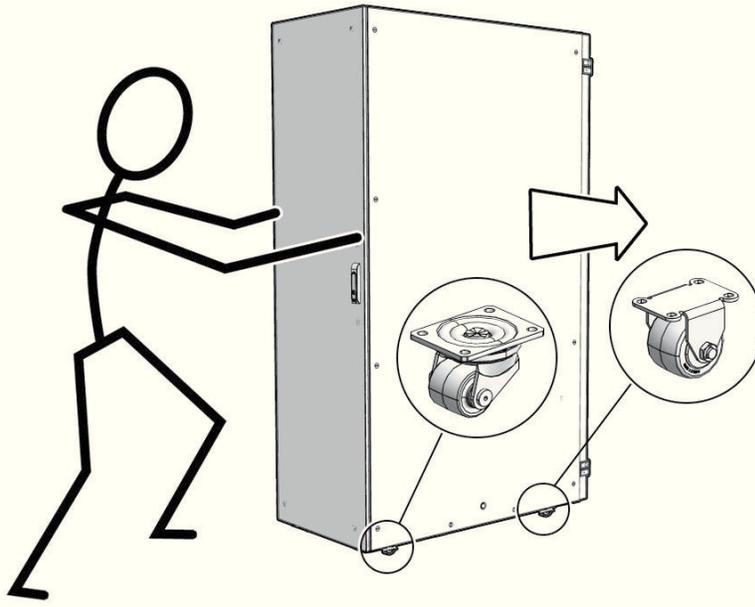
6.3.1 Moving ZDLRA Rack

The following procedure describes how to move ZDLRA Rack:

1. Ensure the doors are closed and secured.
2. Ensure the leveling and stabilizing feet on the rack are raised and out of the way.
3. Push ZDLRA Rack from behind to the installation site.

When moving ZDLRA Rack to the installation site, the front casters do not turn; you must steer the unit by moving the rear casters. You can safely maneuver ZDLRA Rack by carefully pushing it. [Figure 6-2](#) (page 6-6) shows the correct way to push the rack.

Figure 6-2 Carefully Push ZDLRA Rack from Behind



Oracle recommends having two people to move the rack: one person in front and one person in back to help guide the rack. When transporting configured racks from one location to another, take care to move them slowly, 0.65 meters per second (2.13 feet per second) or slower.

Carefully examine the transportation path. Avoid obstacles such as doorways or elevator thresholds that can cause abrupt stops or shocks. Go around obstacles by using ramps or lifts to enable smooth transport.

⚠ **WARNING:**

- Never attempt to move ZDLRA Rack by pushing on the side panels. Pushing on the side panels can tip the rack over. This can cause serious personal injury or death as well as damage to the equipment.
- Never tip or rock ZDLRA Rack because the rack can fall over.

6.3.2 Stabilizing ZDLRA Rack

After moving ZDLRA Rack to the installation site, stabilize the rack to ensure that it does not move or tip over.

Oracle does not recommend using the shipping brackets to permanently mount the rack to the floor.

▲ Caution:

Shipping brackets are not for use for bracing or anchoring the rack during seismic events.

You can stabilize the rack by extending the rack leveling feet.

- [Stabilize ZDLRA Rack with Leveling Feet](#) (page 6-7)
- [Stabilize ZDLRA Rack X7 with Leveling Feet](#) (page 6-8)

6.3.2.1 Stabilize ZDLRA Rack with Leveling Feet

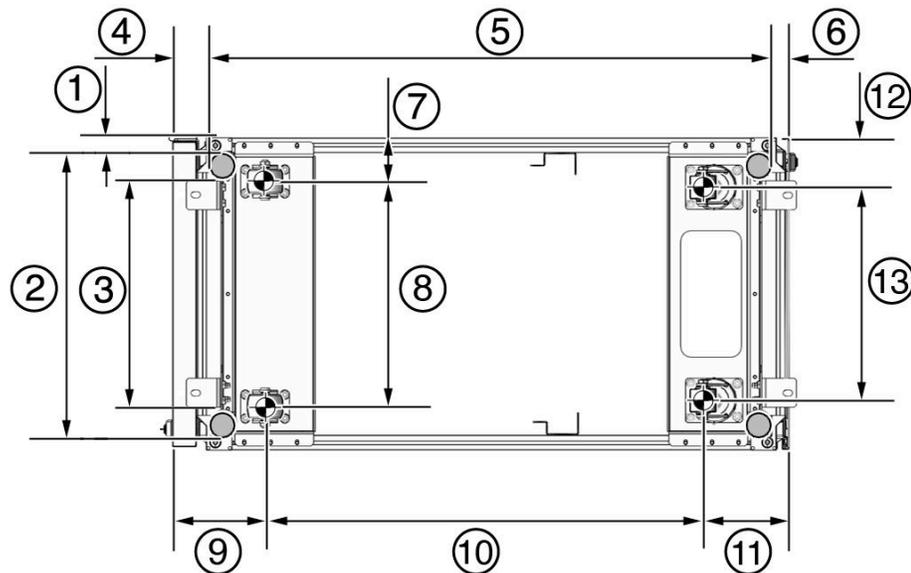
The ZDLRA Rack contains four leveling feet that can be lowered to stabilize the rack. The leveling feet can be used even when the rack is permanently secured to the floor. The four leveling feet should share the rack load with at least one other support. The following is an acceptable method to share the rack load:

- Casters and leveling feet

To adjust the leveling feet, do the following:

1. Locate the four leveling feet located at the bottom corners of the ZDLRA Rack. [Figure 6-3](#) (page 6-7) shows the location of the leveling feet on the bottom of the rack.

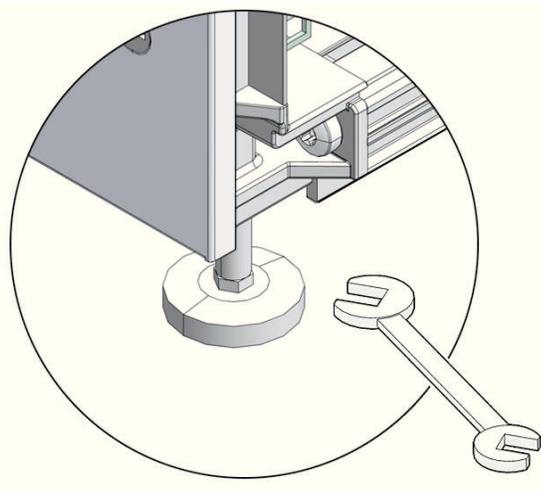
Figure 6-3 Location of Leveling Feet on ZDLRA Rack



- 1: Distance from the edge of the mounting feet to the side of the rack is 33.75 mm (1.33 inches)
- 2: Width from the outside edges of the leveling feet is 532.5 mm (20.96 inches)

- 3: Width from the inside edges of the leveling feet is 429 mm (16.89 inches)
 - 4: Distance from the edge of the feet to the front rack surface is 73.75 mm (2.90 inches)
 - 5: Depth of the outside edges of the leveling feet is 1058.5 mm (41.67 inches)
 - 6: Distance from the edge of the leveling feet to the rear rack surface is 32.5 mm (1.28 inches)
 - 7: Distance from the center of front casters to the side of the rack is 86.7 mm (3.41 inches)
 - 8: Width between the center of the front casters is 426.6 mm (16.80 inches)
 - 9: Distance from the center of the rear casters to the rear of the rack is 173.7 mm (6.83 inches)
 - 10: Depth between the front and rear casters is 828.6 mm (32.62 inches)
 - 11: Distance between the rear casters and the rear of the rack is 162.4 mm (6.39 inches)
 - 12: Distance from the center of rear casters to the side of the rack is 96.4 mm (3.80 inches)
 - 13: Width between the center of the rear casters is 407.2 mm (16.03 inches)
2. Lower the leveling feet to the floor to share the load with the casters, as shown in [Figure 6-4](#) (page 6-8) using the SW 12 mm wrench.

Figure 6-4 Securing ZDLRA Rack Using the Leveling Feet



6.3.2.2 Stabilize ZDLRA Rack X7 with Leveling Feet

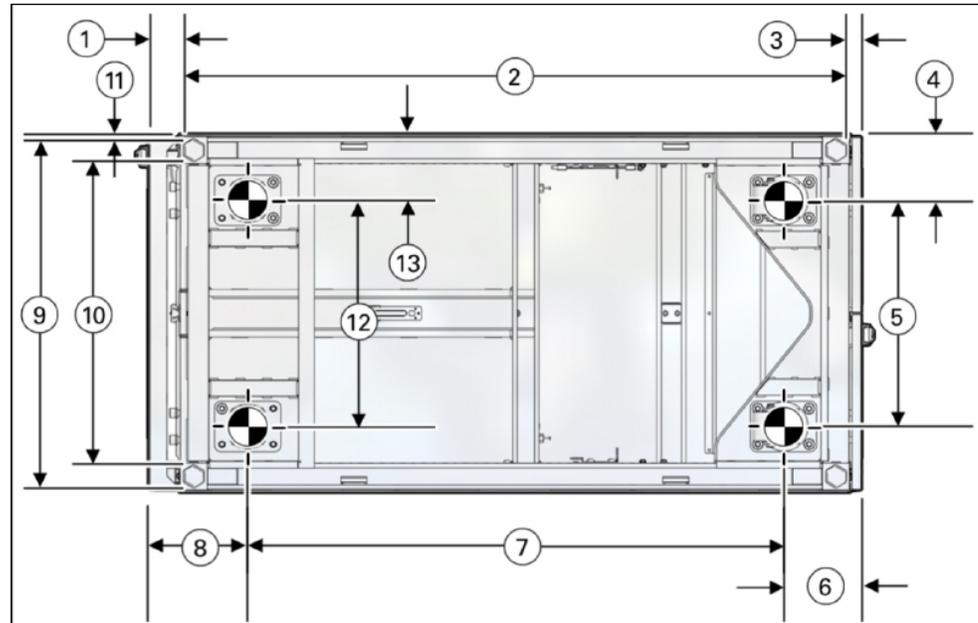
The rack contains four leveling feet that can be lowered to stabilize the rack.

The leveling feet can be used even when the rack is permanently secured to the floor. The four leveling feet share the load with the casters. This increases the footprint of the rack which improves stability and helps prevent rack movement.

To adjust the leveling feet, do the following:

1. Locate the four leveling feet located at the bottom corners of ZDLRA Rack. The following figure shows the location of the leveling feet on the bottom of ZDLRA Rack. The image dimensions include the doors attached to the rack.

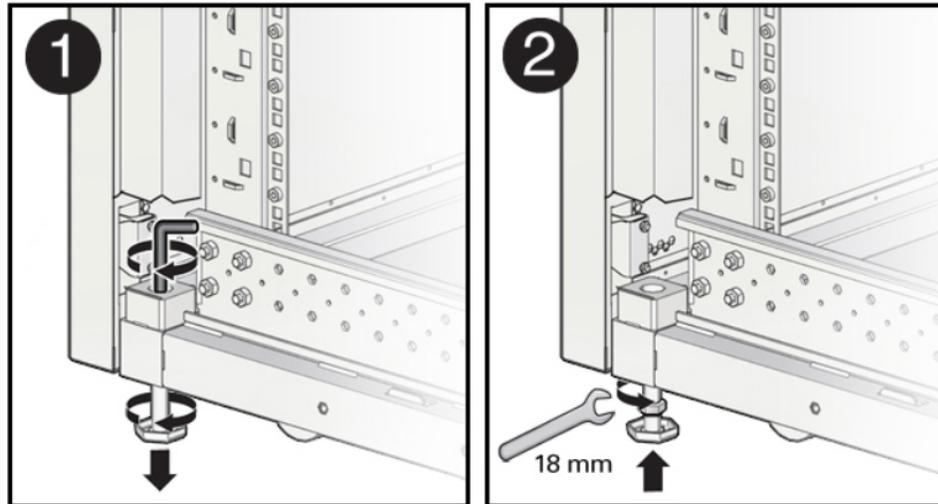
Figure 6-5 Location of Leveling Feet on ZDLRA Rack X7



- 1: Distance from the edge of the feet to the Front door surface is 58 mm (2.3 inches)
- 2: Depth of the outside edges of the leveling feet is 1113 mm (43.8 inches)
- 3: Distance from the edge of the leveling feet to the rear door surface is 26 mm (1.0 inches)
- 4: Distance from the center of rear casters to the side of the rack is 96.5 mm (3.8 inches)
- 5: Width between the center of the rear casters is 407 mm (16.0 inches)
- 6: Distance from the center of the rear casters to the rear door surface is 142 mm (5.6 inches)
- 7: Depth between the front and rear casters is 889 mm (35.0 inches)
- 8: Distance between the center of the front casters and the Front door surface is 166 mm (6.5 inches)
- 9: Width from the outside edges of the leveling feet is 583 mm (23.0 inches)
- 10: Width from the inside edges of the leveling feet is 513 mm (20.2 inches)
- 11: Distance from the edge of the mounting feet to the side of the rack is 8.5 mm (0.3 inches)
- 12: Width between the center of the front casters is 410 mm (16.1 inches)
- 13: Distance from the center of front casters to the side of the rack is 95 mm (3.7 inches)

2. Stabilize the rack by lowering the four leveling feet to the floor, using a 6-mm hex wrench or your fingers, as shown in the following image.

Figure 6-6 Securing ZDLRA Rack Using the Leveling Feet



6.3.3 Attaching a Ground Cable (Optional)

The ZDLRA Rack power distribution units (PDUs) achieve earth ground through their power cords. Final chassis ground is achieved by way of the ground prong when you connect the power cord to a socket. For additional grounding, attach a chassis earth ground cable to ZDLRA Rack. The additional ground point enables electrical current leakage to dissipate more efficiently.

⚠ WARNING:

The PDU power input lead cords and the ground cable must reference a common earth ground. If they do not, then a difference in ground potential can be introduced. If you are unsure of your facility's PDU receptacle grounding, then do not install a ground cable until you confirm that there is a proper PDU receptacle grounding. If a difference in ground potential is apparent, then you must take corrective action.

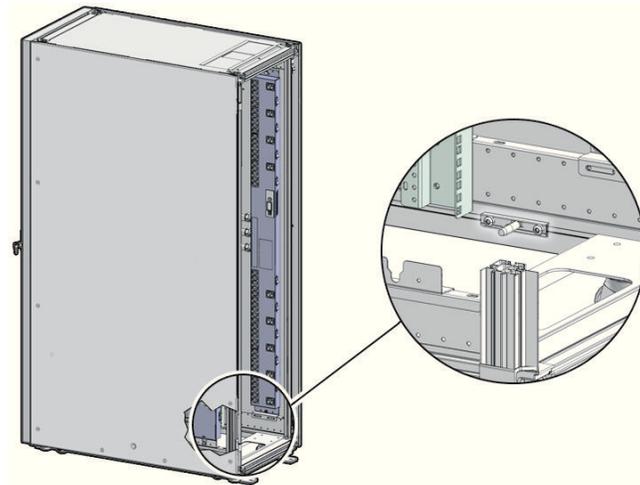
✎ Note:

A grounding cable is not shipped with the system.

1. Ensure the installation site has properly grounded the power source in the data center. The facility PDU must have earth ground.

2. Ensure all grounding points, such as raised floors and power receptacles, reference the facility ground.
3. Ensure that direct, metal-to-metal contact is made for this installation. During manufacturing, the ground cable attachment area might have been painted or coated.
4. Attach the ground cable to one of the attachment points located at the bottom rear of the system frame as shown in [Figure 6-7](#) (page 6-11). The attachment point is an adjustable bolt that is inside the rear of the ZDLRA Rack cabinet on the right side.

Figure 6-7 Earth Ground Attachment Bolt Location



 **See Also:**

Oracle Exadata Database Machine Extending and Multi-Rack Cabling Guide if ZDLRA Rack will be connected to another ZDLRA Rack.

6.4 Installing the Fiber Channel Adapters

The optional QLogic fiber channel adapters are shipped separately from the Recovery Appliance rack. Install them only when Recovery Appliance uses an Oracle tape system for backup.

You install a fiber channel adapter in the first PCIe slot of each compute server.

 **See Also:**

Oracle Server X7-2 Service Manual for instructions for installing a PCI card in slot 1:

https://docs.oracle.com/cd/E72435_01/html/E72445/grdej.html#scrolltoc

Oracle Server X6-2 Service Manual for instructions for installing a PCI card in slot 1:

https://docs.oracle.com/cd/E62159_01/html/E62171/z40007611438967.html

Oracle Server X5-2 Service Manual for instructions for installing a PCI card in slot 1:

http://docs.oracle.com/cd/E41059_01/html/E48312/napsm.z40007611438967.html#scrolltoc

6.5 Acclimating the Recovery Appliance

Oracle recommends an acclimation period of 24 hours for the Recovery Appliance.

This applies to all hardware, including racks and cables. Acclimation to the surrounding environment prevents damage as a result of condensation.

6.6 Powering on the System for the First Time

Before powering on the system for the first time, it is necessary to inspect the rack and connect the power cords. This section contains the following topics:

- [Inspecting Recovery Appliance After It Is in Place](#) (page 6-12)
- [Connecting Power Cords](#) (page 6-13)
- [Powering Up Recovery Appliance](#) (page 6-14)

 **Note:**

After powering on Recovery Appliance, configure it as described in [Configuring a Recovery Appliance Rack](#) (page 10-1).

6.6.1 Inspecting Recovery Appliance After It Is in Place

You should examine the Recovery Appliance physical system after it is in place, but before power is supplied.

To inspect the Recovery Appliance rack:

1. Check the rack for damage.
2. Check the rack for loose or missing screws.

3. Check the 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:
 - a. Check the power cables. Ensure that the connectors are the correct type 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 into the front of the rack.



See Also:

["Ventilation and Cooling Requirements \(page 2-14\)"](#)

6.6.2 Connecting Power Cords



WARNING:

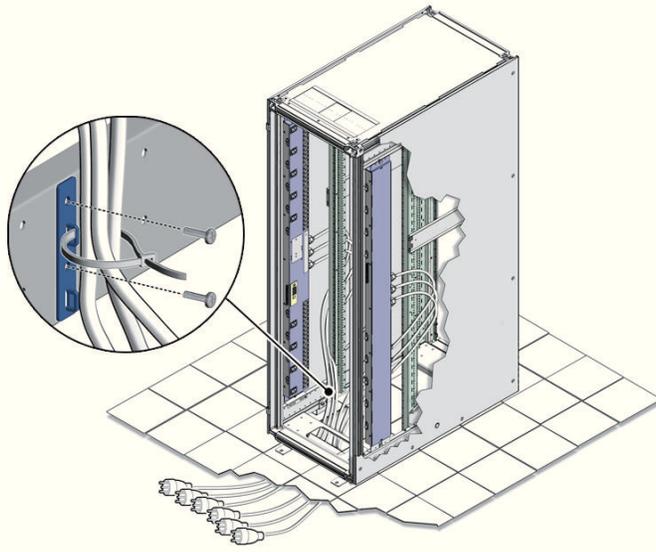
Confirm that all breaker switches are off before connecting the power cords. When off (0), the breakers extend beyond the side of the PDU.

There is one switch for each group and six switches for each PDU.

To connect power cords to the Recovery Appliance 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 through the bottom, as shown in [Figure 6-8](#) (page 6-14).
5. Secure the power cords in bundles as shown in [Figure 6-8](#) (page 6-14).
6. Ensure that the breaker switches are OFF before connecting the power cables.
7. Plug the PDU power cord connectors into the facility receptacles.

Figure 6-8 Power Cord Routing Example



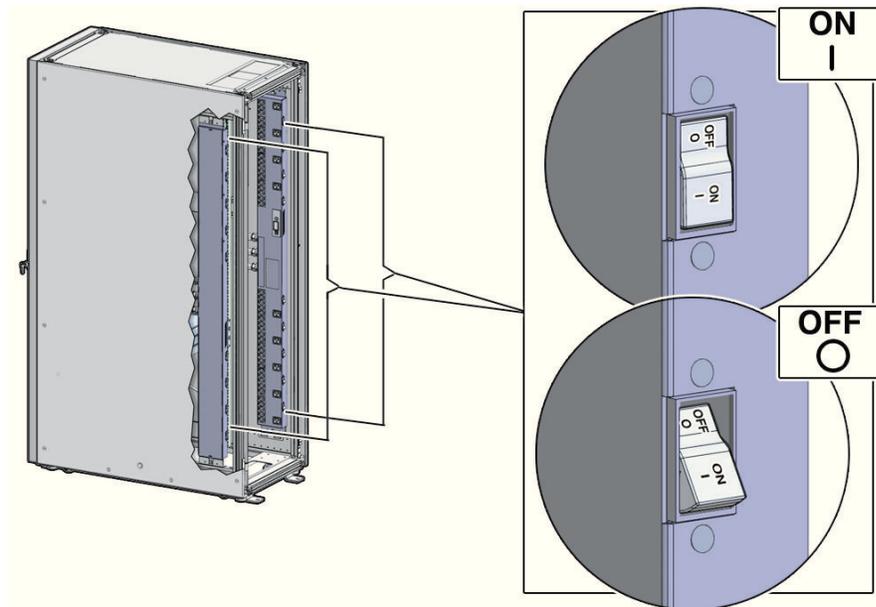
6.6.3 Powering Up Recovery Appliance

The compute servers may take up to 15 minutes to start through the normal BIOS POST tests.

To power up Recovery Appliance:

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 cabinet as shown in [Figure 6-9](#) (page 6-14). Press the ON (I) side of the toggle switches, so that the breakers are flush with the side of the PDU.

Figure 6-9 PDU Switch Locations



2. Verify that the expected power LEDs are on. The LEDs are located as follows:
 - **Compute servers:** Right LED
 - **Storage servers:** Top LED
 - **Cisco switch (viewed from front):** Left LED is green, and the other LEDs are red
 - **InfiniBand switches (viewed from front):** Left LED labeled PS0, or right LED labeled PS1
3. Connect power to PDU A. Ensure the breaker switches are in the OFF position before connecting the power cables.



See Also:

"[Reviewing Safety Guidelines](#) (page 6-1)" about single phase PDUs and cabling

4. Switch on the PDU A circuit breakers, one at a time.
5. Verify that the expected LEDs are on.
6. Perform a visual check of all cable connections in the rack. Do not press every connector to verify a connection.
7. Verify that the OK LED is blinking standby for all systems. The OK LED blinks every 3 seconds when in standby mode. It is dark for 2 to 3 minutes while Oracle ILOM is booting before going into standby mode.

7

Expanding a Recovery Appliance Rack with Additional Storage Servers

The Recovery Appliance racks have a flexible configuration that allows a minimum of three storage servers. You can extend a rack up to a total of 18 storage servers for Recovery Appliance X7, X6, or X5, and 14 storage servers for Recovery Appliance X4. All Recovery Appliance racks have two compute servers, regardless of the number of storage servers.

This chapter describes how to expand a Recovery Appliance rack by adding storage servers. It contains the following sections:

- [Preparing to Extend Recovery Appliance](#) (page 7-1)
- [Precautions and Guidelines](#) (page 7-4)
- [Adding Storage Servers to a Rack](#) (page 7-6)
- [Making New Storage Available](#) (page 7-18)

See Also:

- [In-Rack Cabling Tables for Recovery Appliance X7, X6, and X5](#) (page E-1)
- [In-Rack Cabling Tables for Recovery Appliance X4](#) (page F-1)

7.1 Preparing to Extend Recovery Appliance

Before expanding a rack, review the cabling requirements and collect information about the current rack. This section contains the following topics:

- [Estimating InfiniBand Cable Path Lengths](#) (page 7-1)
- [Obtaining the Current Configuration Information](#) (page 7-2)
- [Generating the Network Configuration Files](#) (page 7-3)
- [Moving the Audit and Diagnostic Files](#) (page 7-4)
- [Synchronizing Release and Patch Levels](#) (page 7-4)
- [Checking the Recovery Appliance Rack](#) (page 7-6)

7.1.1 Estimating InfiniBand Cable Path Lengths

Ensure that you have the required number and lengths of InfiniBand cables for your installation. Cable paths should be as short as possible. After you calculate the length of a cable path, consider the following ways that the length might increase. Then select the shortest cable to satisfy the length requirement.

Consider the following:

- Bends in the cable path increase the required length of the cable. A cable rarely travels in a straight line from connector to connector. Bends in the cable path are necessary, and each bend increases the total length.
- Bundling increases the required length of the cables. Bundling causes one or more cables to follow a common path. However, the bend radius is different in different parts of the bundle. If the bundle is large and unorganized, and there are many bends, one cable might bend only on the inner radius, while another cable might bend only on the outer radius. In this situation, the difference in the required cable lengths is substantial.
- If you are routing the InfiniBand cable under the floor, consider the height of the raised floor when calculating the cable path length.

7.1.2 Obtaining the Current Configuration Information

Use the current configuration of the Recovery Appliance rack to plan patching requirements, configure new IP addresses, and so on.

Log into the first compute server in the rack and collect the following information:

- The `exachk` report for the rack. See My Oracle Support Doc ID 1070954.1 for information about the `exachk` utility.
- Image history information. Use the following command:


```
# dcli -g ~/all_group -l root "imagehistory" > imagehistory.txt
```
- Current IP addresses of all storage servers and compute servers. Use the following command:


```
# dcli -g ~/all_group -l root "ifconfig" > ifconfig_all.txt
```
- Configuration of the storage servers, disks, flash logs, and IORM plans. Use the following commands:


```
# dcli -g ~/cell_group -l root "cellcli -e list cell detail" > cell_detail.txt

# dcli -g ~/cell_group -l root "cellcli -e list physicaldisk detail" > \
physicaldisk_detail.txt

# dcli -g ~/cell_group -l root "cellcli -e list griddisk attributes
name,offset,size,status,asmmodestatus,asmdeactivationoutcome" > griddisk.txt

# dcli -g ~/cell_group -l root "cellcli -e list flashcache detail" >
fc_detail.txt

# dcli -g ~/cell_group -l root "cellcli -e list flashlog detail" > fl_detail.txt

# dcli -g ~/cell_group -l root "cellcli -e list iormplan detail" >
iorm_detail.txt
```
- HugePages memory configuration on the compute servers. Use the following command:


```
dcli -g ~/dbs_group -l root "cat /proc/meminfo | grep 'HugePages'" >
hugepages.txt
```
- InfiniBand switch information. Use the following command:


```
# ibswitches > ibswitches.txt
```

- Firmware version of the Sun Datacenter InfiniBand Switch 36 switches. Use the `nm2version` command on each switch.
- The following network files on the first compute server in the rack:
 - `/etc/resolv.conf`
 - `/etc/ntp.conf`
 - `/etc/network`
 - `/etc/sysconfig/network-scripts/ifcfg-*`
- Any users, user identifiers, groups, and group identifiers created for cluster-managed services that must be created on the new servers:
 - `/etc/passwd`
 - `/etc/group`
- The current cluster status. Use the following command:

```
# crsctl stat res -t > crs_stat.txt
```
- Patch information from the Grid Infrastructure and Oracle homes. Use the following commands as the Grid Infrastructure home owner and the Oracle home owner:

```
/u01/app/oracle/product/11.2.0/dbhome_1/OPatch/opatch lsinventory -oh \
GRID_HOME -detail -all_nodes > opatch_grid.txt

/u01/app/oracle/product/11.2.0/dbhome_1/OPatch/opatch lsinventory -oh \
ORACLE_HOME -detail -all_nodes >> opatch_oracle.txt
```

In the preceding commands, `GRID_HOME` is the path of the Grid Infrastructure home directory, and `ORACLE_HOME` is the path of the Oracle home directory.

7.1.3 Generating the Network Configuration Files

Use the Oracle Exadata Deployment Assistant to list the IP addresses for the new servers. Specify only the number of new servers. Define all other items, such as the Oracle home location and owner, the same as the existing rack configuration. In particular, the bonding configuration of the new servers must match the existing servers in the rack. The Oracle Exadata Deployment Assistant InfiniBand Network Configuration Page enables you to select the type of bonding.

You must install the Recovery Appliance software in the new storage servers, after installing them in the rack.

See Also:

- [Using the Deployment Assistant](#) (page 5-1)
- [Installing the Recovery Appliance Software](#) (page 11-1)

7.1.4 Moving the Audit and Diagnostic Files

Move to delete files in the `$GRID_HOME/rdbms/audit` directory and the `$GRID_HOME/log/diagnostics` directory before extending a cluster. Oracle recommends doing this task a day or two before the planned extension to allow sufficient time.

7.1.5 Synchronizing Release and Patch Levels

The servers probably have a later release or patch level than the rack where they will be added. You can either update the rack to the later release or reimage the new servers to match the current rack. In either case, ensure that all servers, and all Sun Datacenter InfiniBand Switch 36 switches, are at the same patch level.

Also consider how to handle the Grid Infrastructure and database home releases and bundle patch updates. If you plan to apply the new patches, then Oracle recommends upgrading the existing servers, so that the new servers inherit the releases as part of the extension procedure. Then fewer servers are patched. Apply the patches to the existing servers in advance, thereby reducing the total amount of work required during the extension.

7.2 Precautions and Guidelines

This section lists safety precautions and guidelines for handling the equipment. Familiarize yourself with this information before extending the rack:

- [General Safety Precautions](#) (page 7-4)
- [InfiniBand Cable Precautions](#) (page 7-4)
- [InfiniBand Cable Bundling Guidelines](#) (page 7-5)
- [Cable Management Arm Guidelines](#) (page 7-5)

7.2.1 General Safety Precautions

Read the following safety guidelines before working on a Recovery Appliance rack:

- "[Reviewing Safety Guidelines](#) (page 6-1)"
- *Important Safety Information for Sun Hardware Systems* shipped with the rack
- *Zero Data Loss Recovery Appliance Safety and Compliance Guide* for safety notices

7.2.2 InfiniBand Cable Precautions

Review the following InfiniBand cable precautions before working with InfiniBand cables:

- Fiber optic InfiniBand cables with laser transceivers must be type Class 1.
- Do not allow any copper core InfiniBand cable to bend to a radius tighter than 127 mm (5 inches). Tight bends can damage the cables internally.
- Do not allow an optical InfiniBand cable to bend to a radius tighter than 85 mm (3.4 inches). Tight bends can damage the cables internally.

- Do not use zip ties to bundle or support InfiniBand cables. The sharp edges of the ties can damage the cables internally. Use hook-and-loop straps.
- Do not subject an InfiniBand cable to extreme tension. Do not pull or drag it. Pulling on an InfiniBand cable can damage it internally.
- Unroll the full length of an InfiniBand cable.
- Do not twist an InfiniBand cable more than one revolution over its entire length. Twisting an InfiniBand cable can damage it internally.
- Do not route InfiniBand cables where they can be stepped on or rolled over. Crushing can damage the cable internally.

7.2.3 InfiniBand Cable Bundling Guidelines

When bundling InfiniBand cables in groups, use hook-and-loop straps to keep the cables organized. If possible, use color-coordinated straps to help identify the cables and their routing. The InfiniBand splitter and 4X copper conductor cables are thick and heavy for their length. Consider the retention strength of the hook-and-loop straps when supporting cables. Bundle as few cables as reasonably possible. Oracle recommends that no more than eight cables be bundled together.

If the InfiniBand cables break free of their straps and fall, then they might break internally from striking the floor or from sudden changes in tension. Place the hook-and-loop straps as close together as reasonably possible, for example, one strap every foot (0.3 m). If a cable breaks free from a strap, then the cable cannot fall far before it is supported by another strap.

7.2.4 Cable Management Arm Guidelines

Review the following cable management arm (CMA) guidelines before routing the cables:

- Remove all required cables from the packaging, and allow cables to acclimate or reach operating temperature, if possible. This improves the ability to manipulate the cables.
- Label both ends of each cable using a label stock that meets the ANSI/TIA/EIA 606-A standard, if possible.
- Begin the installation procedure in ascending order.
- Only slide out one server at a time. Sliding out more than one server can cause cables to drop or cause problems when sliding the servers back.
- Separate the installation by dressing cables with the least stringent bend radius requirements first. The following bend radius requirements are based on EIA/TIA 568-x standards, and may vary from the manufacturer's requirements:
 - CAT5e UTP: 4 x diameter of the cable or 1 inch/25.4 mm minimum bend radius
 - AC power cables: 4 x diameter of the cable or 1 inch/ 25.4 mm minimum bend radius
 - TwinAx: 5 x diameter of the cable or 1.175 inch/33 mm.
 - Quad Small Form-factor Pluggable (QSFP) InfiniBand cable: 6 x diameter of the cable or 2 inch/55 mm.

- Fiber core cable: 10 x diameter of the cable or 1.22 inch/31.75 mm for a 0.125 cable.
- Install the cables with the best longevity rate first.

7.3 Adding Storage Servers to a Rack

The following sections describe how add storage servers to a Recovery Appliance rack:

- [Removing the Doors](#) (page 7-8)
- [Adding New Servers](#) (page 7-10)
- [Cabling the Storage Servers](#) (page 7-15)
- [Closing the Rack](#) (page 7-17)

Note:

- Apply any patches to the switches and servers in the rack before extending the hardware.
- You can extend the hardware while Recovery Appliance is online, and with no downtime. However, take extreme care.

7.3.1 Checking the Recovery Appliance Rack

Examine the Recovery Appliance physical systems for any problems:

1. Check the rack for damage.
2. Check the rack for loose or missing screws.
3. Check Recovery Appliance for the ordered hardware configuration.
4. Check that all cable connections are secure and well seated.
5. Check the power cables.
6. Ensure that the correct connectors were supplied for the data center facility power source.
7. Check the network data cables.
8. Check the site location tile arrangement for cable access and airflow.
9. Check the data center airflow into the front of Recovery Appliance.

7.3.2 Unpacking and Labeling the New Storage Servers

The Recovery Appliance racks are supplied with the InfiniBand and Cat5e harness connections for the maximum number of storage servers.

To prepare the new servers for installation:

1. Unpack all components from the packing carton:
 - 1 storage server

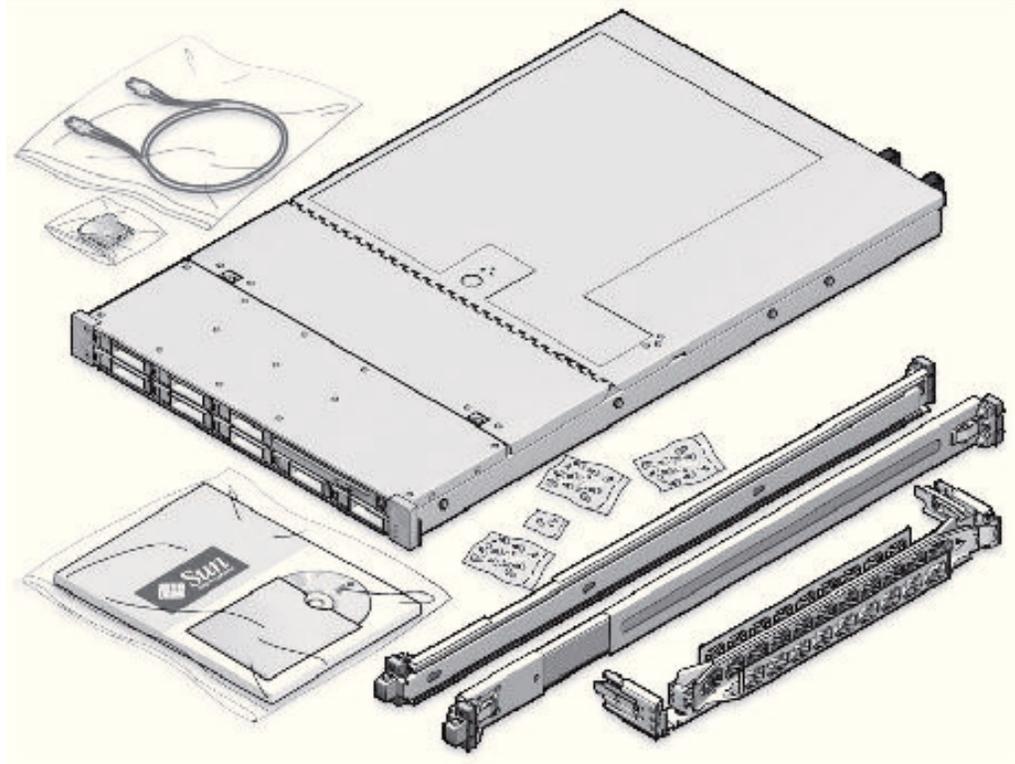
- 2 power cords and labeling kit
 - 1 cable management arm (CMA)
 - 2 rack rails
 - Component documentation
2. Lay out the cables for the servers.
 3. Unroll the cables and stretch them to remove the bends.
 4. Apply the cable labels. Oracle recommends labeling all cables before installation.

 **See Also:**

"[Parts for InfiniBand Switches](#) (page D-33)" for information about cable labels

Figure 7-1 (page 7-7) shows the components of the server expansion kit.

Figure 7-1 Server Expansion Kit Components



 **Note:**

All new equipment receives a Customer Support Identifier (CSI), including the new equipment for Recovery Appliance. Contact Oracle Support Services to reconcile the new CSI with the existing CSI for the Recovery Appliance rack. You must have the original instance or serial numbers and the new numbers when contacting Oracle Support Services.

7.3.3 Recovery Appliance X7, X6, and X5 Population Rules

When adding servers to a Recovery Appliance X7, X6, or X5 rack, follow the population rules listed in [Table 7-1](#) (page 7-8).

Table 7-1 Recovery Appliance X7, X6, and X5 Rack Population Rules

Install Order	Rack Units	Rule
1	U08 to U15	Populate toward the top of the rack (4 units)
2	U27 to U40	Populate toward the top of the rack (7 units)
3	U26 to U23	Populate toward the bottom of the rack (2 units)
4	U19 to U18	(1 unit)
5	U41 to U42	(1 unit)

 **See Also:**

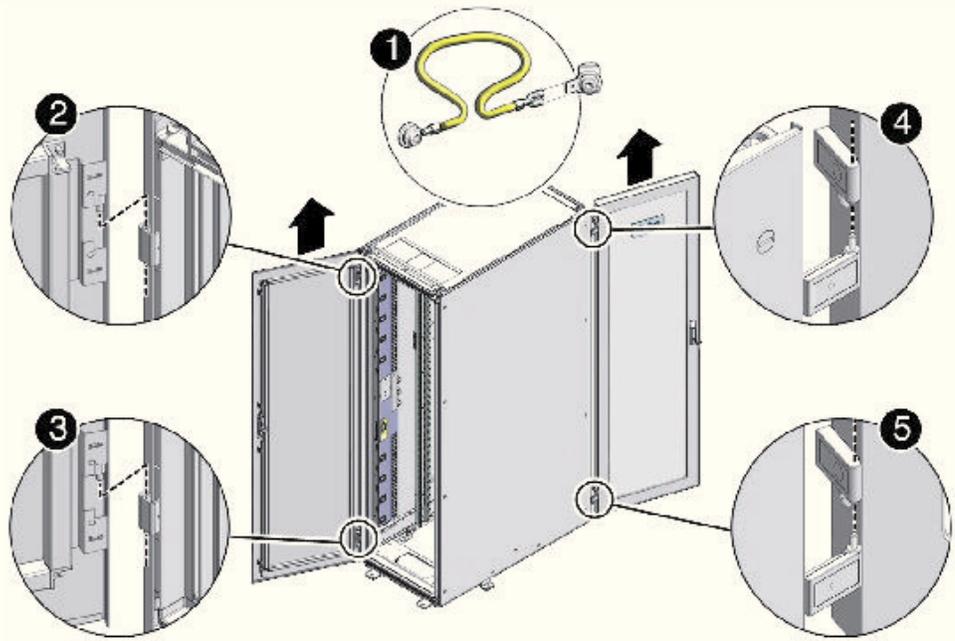
[In-Rack Cabling Tables for Recovery Appliance X7, X6, and X5](#) (page E-1)

7.3.4 Removing the Doors

To remove the doors on the Recovery Appliance rack:

1. Remove the front and rear doors, as follows:
 - a. Unlock the front and rear doors. The keys are in the shipping kit.
 - b. Open the doors.
 - c. Detach the grounding straps connected to the doors by pressing down on the tabs of its quick-release connectors, and pull the straps from the frame.
 - d. Lift the doors up and off their hinges.

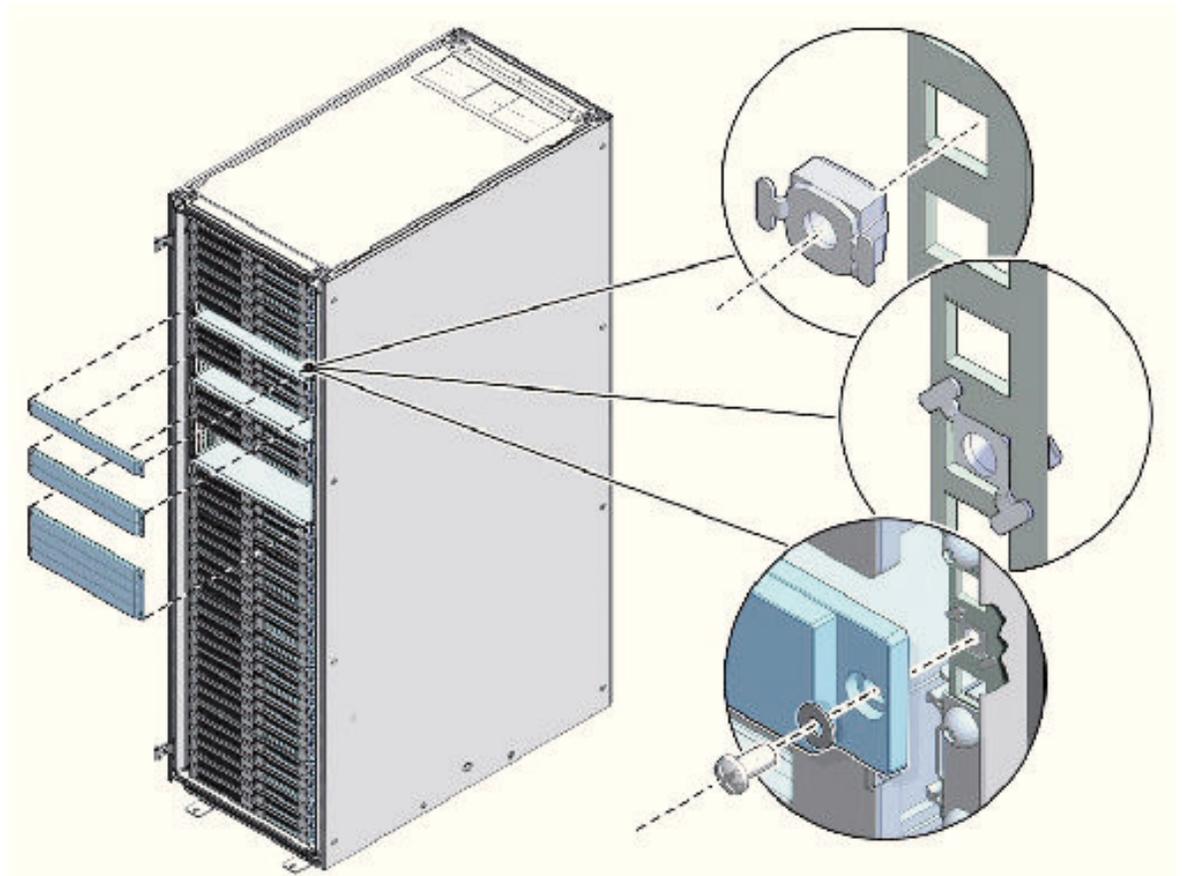
Figure 7-2 Removing the Rack Doors



Description of the callouts in [Figure 7-2](#) (page 7-9):

- 1: Detaching the grounding cable
 - 2: Top rear hinge
 - 3: Bottom rear hinge
 - 4: Top front hinge
 - 5: Bottom front hinge
2. To remove the filler panels where the servers will be installed, use a No. 2 screwdriver to remove the M6 screws. The number of screws depends on the type of filler panel. Save the screws for future use.

Figure 7-3 Removing the Filler Panels



 **Note:**

If you are replacing the filler panels, then do not remove the Duosert cage nuts from the Radio Electronics Television Manufacturers Association (RETMA) rail holes.

7.3.5 Adding New Servers

You can extend Recovery Appliance X7, X6, or X5 to a maximum of 18 storage servers and Recovery Appliance X4 to a maximum of 14 storage servers. You can add any number of storage servers at a time, from one to 15 servers when starting with a minimum configuration.

To prepare for the installation:

1. Identify the rack unit where you will install the server. Fill the first available unit, starting from the bottom of the rack.

See "[Recovery Appliance X7, X6, and X5 Rack Layout](#) (page E-1)" or "[Recovery Appliance X4 Rack Layout](#) (page F-1)", as applicable.

2. Remove and discard the trough, which attaches the cable harness when no server is installed in the unit.

Figure 7-4 Location of the Trough

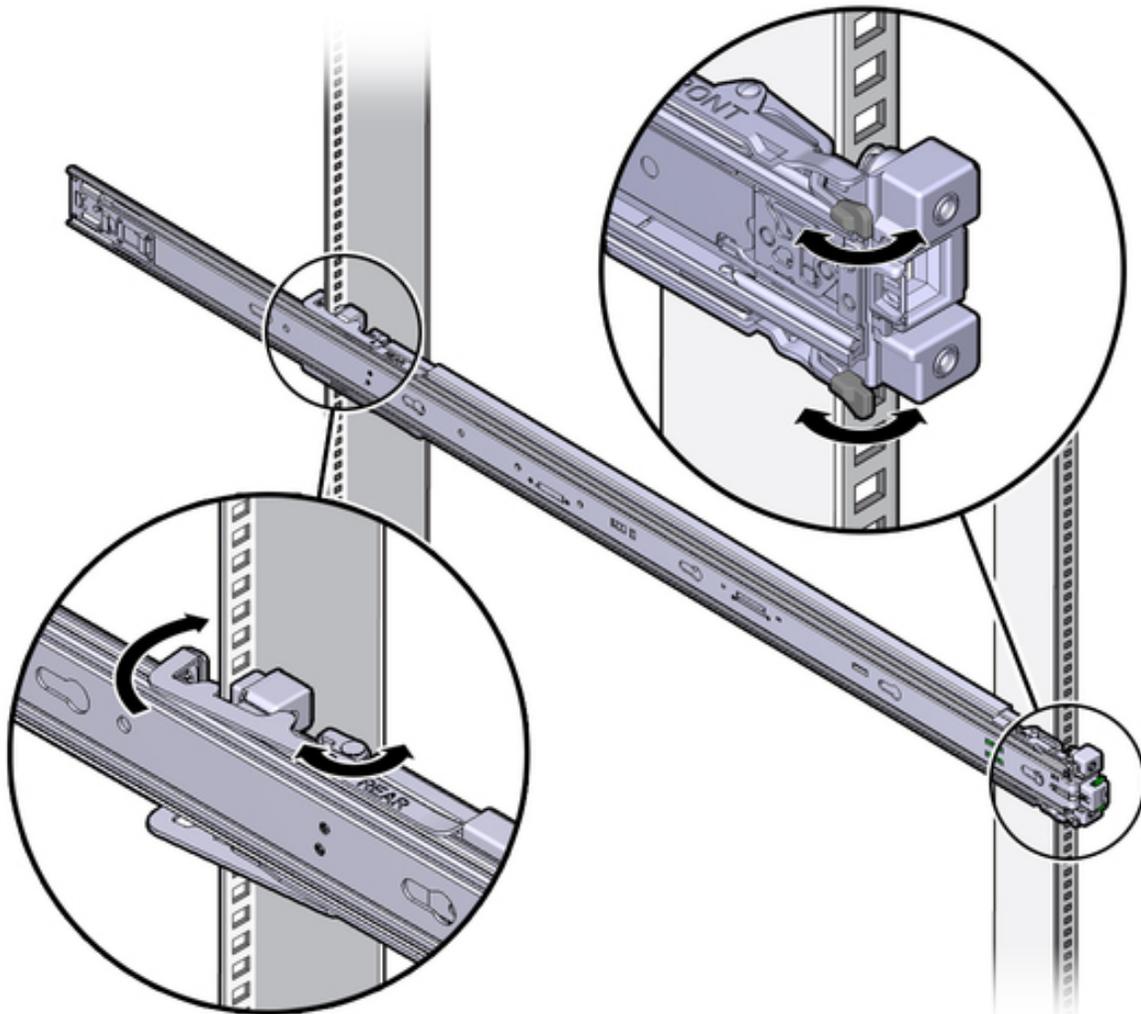


3. Remove and discard the solid filler.

To install the rail assembly:

1. Position a mounting bracket against the chassis so that the slide-rail lock is at the server front, and the five keyhole openings on the mounting bracket are aligned with the five locating pins on the side of the chassis.
2. Orient the slide-rail assembly so that the ball-bearing track is forward and locked in place.
3. Starting on either side of the rack, align the rear of the slide-rail assembly against the inside of the rear rack rail, and push until the assembly locks into place with an audible click.

Figure 7-5 Aligning the Side-Rail Assembly With the Rack



4. Align the front of the slide-rail assembly against the outside of the front rack rail, and push until the assembly locks into place and you hear the click.
5. Repeat steps 2 (page 7-11) to 4 (page 7-12) on the other side on the rack.

 **WARNING:**

- Installing a server requires a minimum of two people or a lift because of the weight of each server. Attempting this procedure alone can result in equipment damage, personal injury, or both.
- Extend the rack anti-tip bar to prevent the rack from tipping during equipment installation.
- For Recovery Appliance X4 racks, load the servers from the bottom up, so that the rack does not become top-heavy and tip over. See "[Recovery Appliance X4 Rack Layout](#) (page F-1)" for the designated location of each server.
- For Recovery Appliance X6 or X5 racks, load the servers following the population rules. See "[Recovery Appliance X7, X6, and X5 Population Rules](#) (page 7-8)".

To install the server:

1. Read the service label on the top cover of the server before installing a server into the rack.
2. Push the server into the slide rail assembly:
 - a. Push the slide rails into the slide rail assemblies as far as possible.
 - b. Position the server so the rear ends of the mounting brackets are aligned with the slide rail assemblies mounted in the equipment rack.

Figure 7-6 Inserting the Server with Mounting Brackets into the Slide Rails

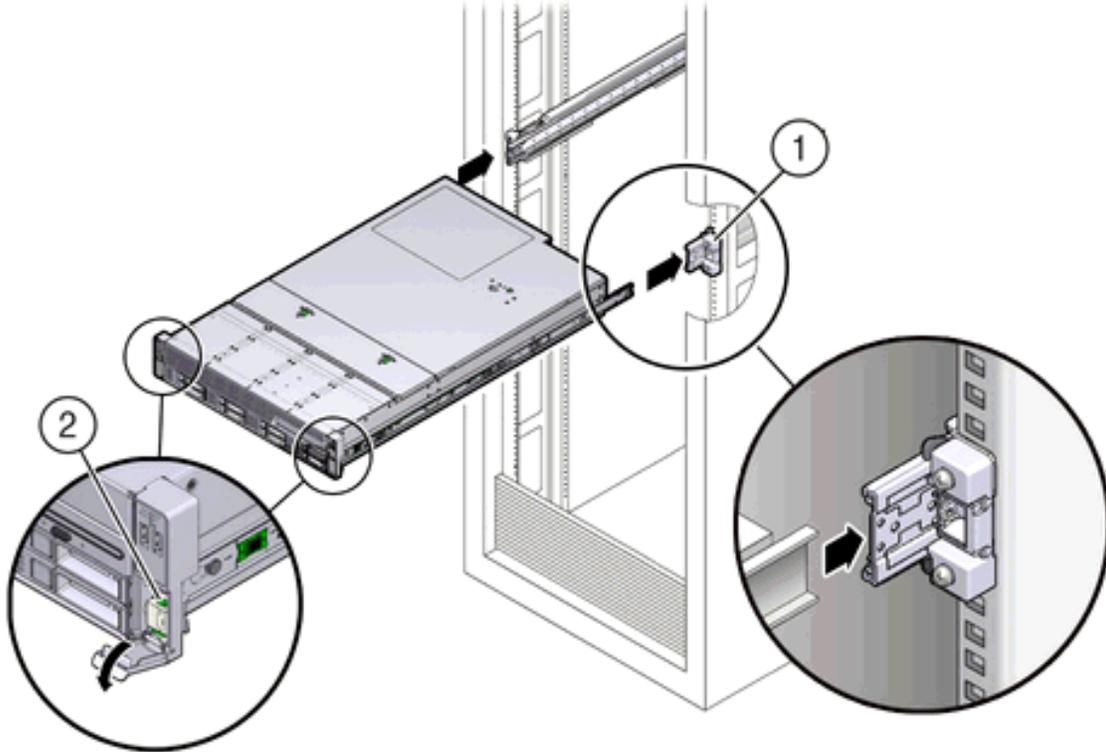


Figure 7-6 (page 7-14) callouts highlight the following:

- 1: Mounting bracket inserted into slide rail
- 2: Slide-rail release lever
- c. Insert the mounting brackets into the slide rails, and push the server into the rack until the mounting brackets encounter the slide rail stops, approximately 30 cm (12 inches).
- d. Simultaneously push down and hold the slide rail release levers on each mounting bracket while pushing the server into the rack.
- e. Continue pushing until the slide rail locks on the front of the mounting brackets engage the slide rail assemblies, and you hear the click.
3. Cable the new server as described in "[Cabling the Storage Servers](#) (page 7-15)".

 **Note:**

Oracle recommends that two people push the servers into the rack: one person to move the server in and out of the rack, and another person to watch the movement of the cables and the CMA.

 **See Also:**

"Installing the Server Into a Rack" in *Sun Server X4-2L Installation Guide* at
http://docs.oracle.com/cd/E36974_01/html/E38145/z4000087165586.html#scrolltoc

7.3.6 Cabling the Storage Servers

After installing the new storage servers, cable them to the existing equipment.

 **Note:**

- The existing cable connections in the rack do not change.
- The blue cables connect to the compute servers, and the black cables connect to the storage servers. These network cables are for the NET0 Ethernet interface port.
- Attach and route the management cables on the CMA and rear panel for one server at a time. Do not slide out multiple servers.
- Start from the bottom of the rack, and work upward.
- Longer hook and loop straps are needed when cabling three CAT5e cables or two twinaxial cables.

To cable the new servers in the rack:

1. Install the cable management arm (CMA):
 - a. Into the left slide rail, slide CMA connector A until it locks into place with an audible click.
 - b. Into the right slide rail, slide CMA connector B until it locks into place with an audible click. Then do the same with CMA connector C.
2. Insert the cables into their ports through the hook and loop straps, then route the cables into the CMA in this order:
 - a. Power
 - b. Ethernet
 - c. InfiniBand



3. Route the cables through the CMA and secure them with hook and loop straps on both sides of each bend in the CMA.



4. Close the crossbar covers to secure the cables in the straightaway.
5. Verify operation of the slide rails and the CMA for each server:

 **Note:**

Oracle recommends that two people do this step: one person to move the server in and out of the rack, and another person to watch the movement of the cables and the CMA.

- a. Slowly pull the server out of the rack until the slide rails reach their stops.
- b. Inspect the attached cables for any binding or kinks.
- c. Verify that the CMA extends fully from the slide rails.
6. Push the server back into the rack:
 - a. Release the two sets of slide rail stops.
 - b. Locate the levers on the inside of each slide rail, just behind the back panel of the server. They are labeled PUSH.

- c. Simultaneously push in both levers and slide the server into the rack, until it stops in approximately 46 cm (18 inches).
 - d. Verify that the cables and CMA retract without binding.
 - e. Locate the slide rail release buttons near the front of each mounting bracket.
 - f. Simultaneously push in both slide rail release buttons and slide the server completely into the rack, until both slide rails engage.
7. Dress the cables, and then tie off the cables with the straps. Oracle recommends that you dress the InfiniBand cables in bundles of eight or fewer.
8. Slide each server out and back fully to ensure that the cables are not binding or catching.
9. Repeat the procedure for all servers.
10. Connect the power cables to the power distribution units (PDUs). Ensure the breaker switches are in the OFF position before connecting the power cables. Do not plug the power cables into the facility receptacles now.

 **See Also:**

- [In-Rack Cabling Tables for Recovery Appliance X4](#) (page F-1)
- "Install the Second-Generation Cable Management Arm" in *Sun Server X4-2L Installation Guide* at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000151427029.html#scrolltoc

7.3.7 Closing the Rack

To close the rack after installing new equipment:

1. Replace the rack front and rear doors:
 - a. Retrieve the doors and place them carefully on the door hinges.
 - b. Connect the front and rear door grounding strap to the frame.
 - c. Close the doors.
 - d. (Optional) Lock the doors. The keys are in the shipping kit.
2. (Optional) Replace the side panels, if they were removed for the upgrade:
 - a. Lift each side panel up and onto the side of the rack. The top of the rack can support the weight of the side panel. Ensure that the panel fasteners line up with the grooves in the rack frame.
 - b. Turn each side panel fastener 1/4 turn clockwise using the side panel removal tool. Turn the fasteners next to the panel lock clockwise. There are 10 fasteners per side panel.
 - c. (Optional) Lock each side panel. The key is in the shipping kit. The locks are located on the bottom center of the side panels.
 - d. Connect the grounding straps to the side panels.

After closing the rack:

- If you are installing a new system, proceed to "[Installing the Software on Recovery Appliance](#) (page 11-3)"
- If you added the server to increase storage on an existing system, proceed to "[Making New Storage Available](#) (page 7-18)"

7.4 Making New Storage Available

After you add storage servers to a rack to increase the storage capacity of an existing system, you must run `racli` commands to make the additional storage available for backups.

Caution:

After you add a storage cell to the Recovery Appliance, you cannot remove it.

Before you begin the procedure, ensure that:

- The new storage cells have the same Exadata image version as the existing storage cells.
- The new storage cells do not have any associated databases. (You can run `cellcli -e list database` to verify this is the case.)

To make the new storage available:

1. Log in to Recovery Appliance as `root`.
2. Change to the `bin` directory:

```
# cd /opt/oracle.RecoveryAppliance/bin
```

3. Run the command to add the new storage cells.

The following command adds storage cells named `cellhost1` and `cellhost2` to the Recovery Appliance:

```
# ./racli add storage_cells --cell=cellhost1,cellhost2
```

4. If you are prompted to do so, run the command to add storage disks.

The following command adds storage disks with the newly added cells `cellhost1` and `cellhost2`:

```
# ./racli add storage_disks --cell=cellhost1,cellhost2
```

Note:

You can execute this command while the system is running but the operation can take a while to complete and can negatively impact performance. Avoid running other I/O intensive tasks such as patching or maintenance during this time.

5. If you are prompted to update the storage location:
 - a. Run the `status` command to ensure that the rebalance operation is complete.

```
# ./racli status diskgroup
```

- b.** Run the command to update the storage location.

The following command increases the size of the storage location using the newly added cells `cellhost1` and `cellhost2`:

```
# ./racli update storage_location --cells=cellhost1,cellhost2
```

- 6.** (Optional) When the procedure is complete, check the size of the storage location.

 **See Also:**

- [RACLI Overview](#) (page A-1) for details about the commands used in the procedure.
- *Zero Data Loss Recovery Appliance Administrator's Guide* for information on how to view storage location information.

8

Extending Recovery Appliance with Multiple Racks

This chapter describes how to cable together multiple Recovery Appliance racks. It contains the following sections:

- [About Multirack Configurations](#) (page 8-1)
- [Extending a Rack by Adding Another Rack](#) (page 8-2)
- [Adding the Spine Switches](#) (page 8-2)
- [Floor and Underfloor Delivery of InfiniBand Cables](#) (page 8-4)
- [Cabling Two Racks Together](#) (page 8-4)
- [Cabling Several Racks Together](#) (page 8-5)
- [Configuring Multiple Racks](#) (page 8-7)



See Also:

[Multirack Cabling Tables](#) (page G-1)

8.1 About Multirack Configurations

You can add racks together to increase system capacity and performance. Recovery Appliance has a flexible configuration. A rack has two compute servers and a minimum of three storage servers. A Recovery Appliance X6 or X5 full rack has 18 storage servers and a Recovery Appliance X4 full rack has 14 storage servers. The maximum configuration is 18 full racks.

You can cable together multiple Recovery Appliance racks that have any number of storage servers; they do not have to be full racks. Spine switches connect the InfiniBand network among multiple racks. You must order the spine switches and cables.

Before cabling the racks together:

1. Determine the number of racks to be cabled together.
2. Obtain the total number of InfiniBand cables required for the installation. See [Multirack Cabling Tables](#) (page G-1).
3. Ensure that each rack has a Sun Datacenter InfiniBand Switch 36 spine switch installed in the bottom rack unit (U1). See "[Recovery Appliance X7, X6, and X5 Rack Layout](#) (page E-1)" or "[Recovery Appliance X4 Rack Layout](#) (page F-1)", as applicable.
4. Gather the following tools:

- Pen
- Note pad
- Flashlight
- Labels for cables and rack
- Cable ties (Oracle recommends hook and loop straps)



See Also:

"[Parts for InfiniBand Switches](#) (page D-33)" for information about cable labels

5. Review the safety guidelines. See *Zero Data Loss Recovery Appliance Safety and Compliance Guide*.

8.2 Extending a Rack by Adding Another Rack

You can extend Recovery Appliance by adding it to another Recovery Appliance rack. The procedure consists of cabling and configuring the racks together. You can cable the racks with no downtime.

Note the following:

- There is some performance degradation while cabling the racks together. This degradation results from reduced network bandwidth, and retransmitting data because of packet loss when a cable is unplugged.
- The environment is not a high-availability environment because one leaf switch must be off. All traffic goes through the remaining leaf switch.
- Only the existing rack is operational, and the new rack is powered down.
- The software running on the systems must not have problems related to InfiniBand restarts.
- The new racks must be configured with the appropriate IP addresses to be migrated into the expanded system before cabling. Duplicate IP addresses are not allowed.
- The existing spine switch is set to priority 10 during the cabling procedure. This setting gives the spine switch a higher priority than any other switch in the fabric. The spine switch is the first to take the Subnet Manager Master role whenever a new Subnet Manager Master is set during the cabling procedure.

8.3 Adding the Spine Switches

A spine switch is required in each rack when configuring multiple racks on the same InfiniBand fabric. A spine switch can be factory installed before shipment, or it can be installed on site. Recovery Appliance uses a third Sun Datacenter InfiniBand Switch 36 as the spine switch.

 **Note:**

This procedure is specific to Recovery Appliance. It differs from the steps in the Sun Datacenter InfiniBand Switch 36 manual.

1. Review the guidelines for adding a rack component. See "[Precautions and Guidelines](#) (page 7-4)".
2. Unpack the Sun Datacenter InfiniBand Switch 36 switch components from the packing cartons:
 - Sun Datacenter InfiniBand Switch 36 switch
 - Cable bracket and rackmount kit
 - Cable management bracket and cover
 - 2 rack rail assemblies
 - Assortment of screws and captive nuts
 - Sun Datacenter InfiniBand Switch 36 documentation

The service label procedure on top of the switch includes descriptions of the preceding items.
3. Install cage nuts in each rack rail in the appropriate holes.
4. Attach the brackets with cutouts to the power supply side of the switch.
5. Attach the C-brackets to the switch on the side of the InfiniBand ports.
6. Slide the switch halfway into the rack from the front. You must keep it to the left side of the rack as far as possible while pulling the two power cords through the C-bracket on the right side.
7. Use a long-shaft, 16-inch, No. 3, Phillips screwdriver to install the slide rails from the rear of the rack into the C-brackets on the switch.
8. Attach the cable arm brackets to the slide rail to the rear of the rack.
9. Install the lower screw loosely, with the cable arm bracket rotated 90 degrees downward to allow access to the screw.
10. Rotate the cable arm bracket to the correct position.
11. Install the upper screw.
12. Tighten both screws.
13. Repeat steps [7](#) (page 8-3) through [12](#) (page 8-3) on the other side.
14. Push the switch completely into the rack from the front, routing the power cords through the cutout on the rail bracket.
15. Tighten the screws, using the Phillips screwdriver.
16. Slide the server in rack location U2 out to the locked service position, to improve access to the switch.
17. Install the lower part of the cable management arm across the back of the switch.
18. Connect the cables to the appropriate ports.

 **See Also:**

["InfiniBand Switch-to-Server Cabling Connections \(page F-7\)"](#)

19. Install the upper part of the cable management arm.
20. Slide the server in rack location U2 into the rack.
21. Install power cords into the InfiniBand switch power supply slots on the front.
22. Install the vented filler panel in front of the switch.

8.4 Floor and Underfloor Delivery of InfiniBand Cables

Sun Datacenter InfiniBand Switch 36 switch accepts InfiniBand cables from floor or underfloor delivery. Floor and underfloor delivery limits the tension in the InfiniBand cable to the weight of the cable for the rack height of the switch.

 **Note:**

Overhead cabling details are not included in this guide. For details on overhead cabling, contact your Oracle service engineer.

8.5 Cabling Two Racks Together

This section describes how to cable together two adjacent racks. The existing rack is R1, and the new rack is R2.

To cable two racks together:

1. Set the priority of the current, active Subnet Manager Master to 10 on the spine switch:
 - a. Log in to any InfiniBand switch on the active system.
 - b. Use the `getmaster` command to verify that the Subnet Manager Master is running on the spine switch. If it is not, then follow the procedure in ["Setting a Spine Switch as the Subnet Manager Master \(page 10-8\)"](#).
 - c. Log in to the spine switch.
 - d. Use the `disableesm` command to stop Subnet Manager.
 - e. Use the `setsmpriority 10` command to set the priority to 10.
 - f. Use the `enableesm` command to restart the Subnet Manager.
 - g. Repeat step 1.b (page 8-4) to ensure the Subnet Manager Master is running on the spine switch.
2. Ensure the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
3. Completely shut down the new rack (R2).
4. Cable the leaf switch in the new rack according to [Table G-4 \(page G-8\)](#).

5. Power off leaf switch R1 IB2. This action causes all compute servers and storage servers to fail over their InfiniBand traffic to R1 IB3.
6. Disconnect all interswitch links between R1IB2 and R1 IB3.
7. Cable leaf switch R1 IB2 according to [Table G-3](#) (page G-7).
8. Power on leaf switch R1 IB2.
9. Wait three minutes for R1 IB2 to become completely operational.
To check the switch, log in to it and run the `ibswitches` command. The output should show three switches, R1 IB1, R1 IB2, and R1 IB3.
10. Power off leaf switch R1 IB3. This action causes all compute servers and storage servers to fail over their InfiniBand traffic to R1 IB2.
11. Cable leaf switch R1 IB3 according to [Table G-3](#) (page G-7).
12. Power on leaf switch R1 IB3.
13. Wait three minutes for R1 IB3 to become completely operational.
14. To check the switch, log in to it and run the `ibswitches` command. The output should show three switches, R1 IB1, R1 IB2, and R1 IB3.
15. Power on all InfiniBand switches in R2.
16. Wait for three minutes for the switches to become completely operational.
17. To check the switch, log in to it and run the `ibswitches` command. The output should show six switches, R1 IB1, R1 IB2, R1 IB3, R2 IB1, R2 IB2, and R2 IB3.
18. Run the `getmaster` command from any switch to ensure that the Subnet Manager Master is running on R1 IB1.
19. Power on all servers in R2.
20. Log in to spine switch R1 IB1, and lower its priority to 8:
 - a. Use the `disableism` command to stop the Subnet Manager.
 - b. Use the `setsmpriority 8` command to set the priority to 8.
 - c. Use the `enableism` command to restart the Subnet Manager.
21. Ensure that the Subnet Manager Master is running on one of the spine switches.

After cabling the racks together, configure them as described in [Configuring a Recovery Appliance Rack](#) (page 10-1).

8.6 Cabling Several Racks Together

This section describes how to cable several adjacent racks together. The existing racks are R1, R2,... Rn, the new rack is Rn+1, and the Subnet Manager Master is running on R1 IB1.

To cable several racks together:

1. Set the priority of the current, active Subnet Manager Master to 10 on the spine switch:
 - a. Log in to any InfiniBand switch on the active system.

- b. Use the `getmaster` command to verify that the Subnet Manager Master is running on the spine switch. If it is not, then follow the procedure in "[Setting a Spine Switch as the Subnet Manager Master](#) (page 10-8)".
 - c. Log in to the spine switch.
 - d. Use the `disableesm` command to stop the Subnet Manager.
 - e. Use the `setsmpriority 10` command to set the priority to 10.
 - f. Use the `enableesm` command to restart the Subnet Manager.
 - g. Repeat step 1.b (page 8-4) to ensure that the Subnet Manager Master is running on the spine switch.
2. Ensure that the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
3. Completely shut down the new rack (R_{n+1}).
4. Cable the leaf switch in the new rack according to the appropriate table in [Multirack Cabling Tables](#) (page G-1). For example, if rack R_{n+1} is R4, then use [Table G-11](#) (page G-11).
5. Complete the following procedure for each of the original racks. In these steps, Rx represents a rack number from R1 to R_n .
 - a. Power off leaf switch Rx IB2. This action causes all compute servers and storage servers to fail over their InfiniBand traffic to Rx IB3.
 - b. Cable leaf switch Rx IB2 according to [Multirack Cabling Tables](#) (page G-1).
 - c. Power on leaf switch Rx IB2.
 - d. Wait three minutes for Rx IB2 to become completely operational.
 - e. To check the switch, log in to it and run the `ibswitches` command. The output should show $n*3$ switches for IB1, IB2, and IB3 in racks R1, R2,... R_n .
 - f. Power off leaf switch Rx IB3. This action causes all compute servers and storage servers to fail over their InfiniBand traffic to Rx IB2.
 - g. Cable leaf switch Rx IB3 according to [Multirack Cabling Tables](#) (page G-1).
 - h. Power on leaf switch Rx IB3.
 - i. Wait three minutes for Rx IB3 to become completely operational.
 - j. To check the switch, log in to the switch and run the `ibswitches` command. The output should show $n*3$ switches for IB1, IB2, and IB3 in racks R1, R2,... R_n .
 - k. Repeat these steps for each rack until all of them are recabled according to [Multirack Cabling Tables](#) (page G-1).
6. Power on all InfiniBand switches in R_{n+1} .
7. Wait three minutes for the switches to become completely operational.
8. To check the switch, log in to it and run the `ibswitches` command. The output should show $(n+1)*3$ switches for IB1, IB2, and IB3 in racks R1, R2,... R_{n+1} .
9. Run the `getmaster` command from any switch to ensure that the Subnet Manager Master is running on R1 IB1.
10. Power on all servers in R_{n+1} .
11. Log in to spine switch R1 IB1, and lower its priority to 8:

- a. Use the `disableism` command to stop the Subnet Manager.
 - b. Use the `setsmpriority 8` command to set the priority to 8.
 - c. Use the `enableism` command to restart the Subnet Manager.
12. To ensure that the Subnet Manager Master is running on one of the spine switches, use the `getmaster` command from any switch.
 13. To ensure that the Subnet Manager is running on every spine switch, use the following command from any switch:

```
ibdiagnet -r
```

The Summary Fabric SM-state-priority section of the output should show that the spine switches are running.

14. If a spine switch is not running, then log in to the switch and use the `enableism` command to enable the Subnet Manager.
15. If there are now four or more racks, then log in to the leaf switches in each rack and use the `disableism` command to disable the Subnet Manager

8.7 Configuring Multiple Racks

In addition to the hardware installation of the new racks, you must do the following procedures:

- Generate the installation files. See [Using the Deployment Assistant](#) (page 5-1).
- Configure the rack components. See [Configuring a Recovery Appliance Rack](#) (page 10-1).
- Install the Recovery Appliance software. See [Installing the Recovery Appliance Software](#) (page 11-1).

9

Configuring Ingest Over InfiniBand

The ingest network connects the protected Oracle Database servers to Recovery Appliance for backup within the same data center. You can optionally configure the ingest network over InfiniBand instead of or in addition to a private, high-speed Ethernet network.

Ingest Over InfiniBand is supported only with compute servers running release March 2015 or later with database patch 12.1.0.2 DBBP4 applied.

This chapter contains the following sections:

- [About Ingest Over InfiniBand](#) (page 9-1)
- [Cabling Recovery Appliance to Oracle Exadata Database Machine](#) (page 9-2)
- [Adding VIPs to the Recovery Appliance InfiniBand Network](#) (page 9-4)
- [Updating Recovery Appliance With the New VIPs for InfiniBand](#) (page 9-5)
- [Update the Database Initialization Parameters](#) (page 9-6)

9.1 About Ingest Over InfiniBand

Oracle recommends that the Recovery Appliance receives backups from all Oracle Exadata Database Machine racks over an Ethernet network. This configuration, shown in "[Example of Network Connections for Recovery Appliance](#) (page 3-5)", supports thousands of protected databases equally. Nonetheless, you can configure local Oracle Exadata Database Machine racks to back up to Recovery Appliance over the InfiniBand network. This configuration provides the local Exadata racks with priority access to Recovery Appliance. Redo and SQL*Net control information from the protected databases are still sent over the 10 Gb Ethernet network.

 **Note:**

Do not attempt to use the InfiniBand network for redo transport as this configuration is not supported.

To configure Ingest Over InfiniBand:

1. Ensure that all racks in the InfiniBand fabric are in the same subnet and have unique IP addresses. The IP addresses can be registered in the DNS, but it is not mandatory.

When deploying a new Recovery Appliance rack, identify the same subnet as Oracle Exadata Database Machine and specify a unique range of IP addresses in Oracle Exadata Deployment Assistant.

After deployment, correct any issues with the subnet or IP addresses. See *Oracle Exadata Database Machine Maintenance Guide* for instructions on changing InfiniBand IP addresses and host names.

2. Edit the `/etc/hosts` file on all servers in the InfiniBand fabric to list all IP addresses and host names.
3. Create a VIP IP address for each compute server in the Recovery Appliance rack. See "[Adding VIPs to the Recovery Appliance InfiniBand Network](#) (page 9-4)".
4. Cable the Recovery Appliance and Oracle Exadata Database Machine racks together.

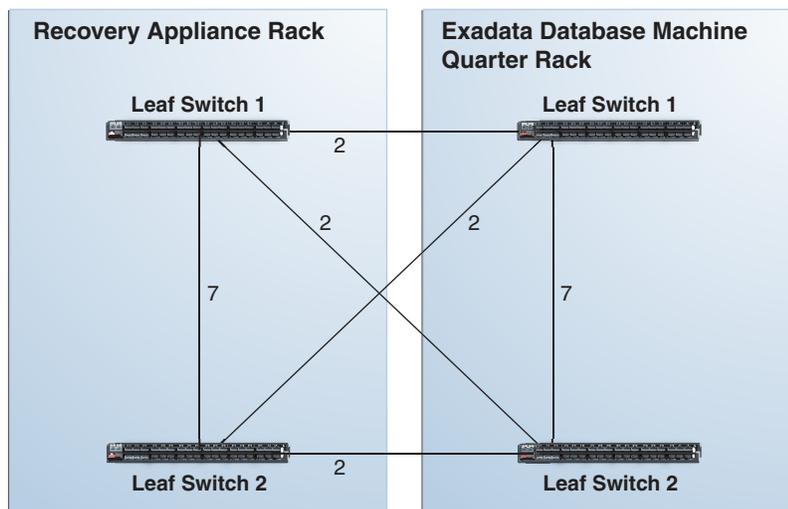
9.2 Cabling Recovery Appliance to Oracle Exadata Database Machine

When configuring Ingest Over InfiniBand, you can connect Recovery Appliance to one or more local Oracle Exadata Database Machine racks using one of these methods:

- If the Recovery Appliance rack has a spine switch, then use the multitrack cabling tables for the appropriate number of racks. See [Multitrack Cabling Tables](#) (page G-1).
- If the Recovery Appliance rack has no spine switch, then you can cable it in a way that is similar to cabling together multiple Oracle Exadata Database Machine quarter racks.

[Figure 9-1](#) (page 9-2) shows the cable connections between a Recovery Appliance rack and an Oracle Exadata Database Machine quarter rack. The leaf switches within each rack retain their existing seven connections. The leaf switches interconnect the racks with two links, each using the ports reserved for external connectivity.

Figure 9-1 Recovery Appliance Connected to an Exadata Quarter Rack



[Figure 9-2](#) (page 9-3) shows the cable connections from Recovery Appliance to an Oracle Exadata Database Machine half rack or full rack. The leaf switches within each rack retain their existing seven connections. The leaf switches interconnect the racks with two links, each using the ports reserved for external connectivity.

Figure 9-2 Recovery Appliance Connected to an Exadata Half or Full Rack

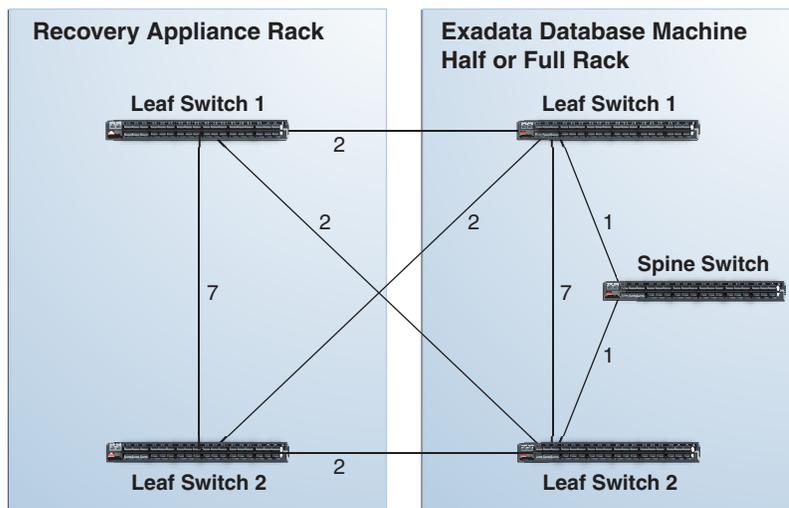
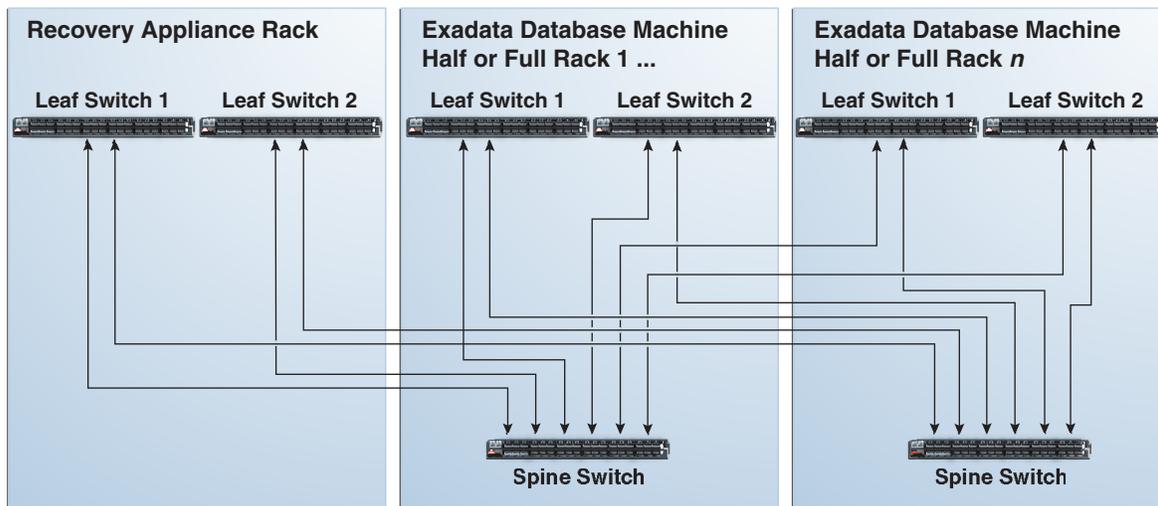


Figure 9-3 (page 9-3) shows the cable connections from Recovery Appliance to two or more Exadata racks. All Oracle Exadata Database Machine racks must be half or full racks, and must be interconnected using a fat-tree topology.

Each Recovery Appliance leaf switch has two links to the spine switches in the Exadata racks. When connecting five or more racks, use one link to the spine switches.

Figure 9-3 Recovery Appliance Connected to Multiple Exadata Racks



9.3 Adding VIPs to the Recovery Appliance InfiniBand Network

After cabling the Recovery Appliance rack with one or more Oracle Exadata Database Machine racks, you must add the virtual IP addresses (VIPs) to the Recovery Appliance InfiniBand network.

To add the VIPs:

1. Log in to a Recovery Appliance compute server.
2. Verify the subnet information for the Recovery Appliance InfiniBand network:

```
$ oifcfg iflist -p -n | grep PRIVATE
eth0 10.10.40.0 PRIVATE 255.255.248.0
ib0 192.168.40.0 PRIVATE 255.255.248.0
ib1 192.168.40.0 PRIVATE 255.255.248.0
bondeth0 10.10.62.0 PRIVATE 255.255.255.0
bondeth1 10.10.241.192 PRIVATE 255.255.255.192
```

3. Change to the `root` user.
4. Add a new network to Oracle Cluster Ready Services (CRS):

```
# srvctl add network -k 3 -S 192.168.40.0/255.255.248.0/ib0\|ib1
```

Note:

You must specify network #3 (use the "-k 3" option) when you configure Ingest Over InfiniBand.

5. Verify that the new network is configured correctly:

```
# srvctl config network -k 3
Network 3 exists
Subnet IPv4: 192.168.40.0/255.255.248.0/ib0:ib1, static
Subnet IPv6:
Ping Targets:
Network is enabled
Network is individually enabled on nodes:
Network is individually disabled on nodes:
```

6. Add the VIPs to the new CRS network:

```
# srvctl add vip -n raadm01 -A raadm01-priv-vip.example.com/255.255.248.0/ib0\|
ib1 -k 3
#
# srvctl add vip -n raadm02 -A raadm02-priv-vip.example.com/255.255.248.0/ib0\|
ib1 -k 3
```

7. Start the VIPs on the Recovery Appliance InfiniBand network:

```
# srvctl start vip -i raadm01-priv-vip.example.com
# srvctl start vip -i raadm02-priv-vip.example.com
```

8. Change to the `oracle` user.
9. Verify that the VIPs are running:

```
$ srvctl status vip -vip raadm01-priv-vip.example.com
VIP 192.168.40.79 is enabled
VIP 192.168.40.79 is running on node: raadm01

$ srvctl status vip -vip raadm02-priv-vip.example.com
VIP 192.168.40.80 is enabled
VIP 192.168.40.80 is running on node: raadm02
```

9.4 Updating Recovery Appliance With the New VIPs for InfiniBand

Recovery Appliance recognizes a preferred network and a default network for receiving backups into the appliance. The deployment procedure configures the 10 Gb Ethernet network as the default network. The InfiniBand network becomes the preferred network. The protected databases use the preferred network if they can contact it; otherwise, they use the default network.

The following procedure updates a table in Oracle Database with the VIPs created in ["Adding VIPs to the Recovery Appliance InfiniBand Network \(page 9-4\)"](#).

To update Recovery Appliance with the Virtual IPs:

1. While logged in to a compute server in the Recovery Appliance rack, open SQL*Plus and connect as `RASYS`.
2. View the current values in the `RAI_HOST` table:

```
SQL> SELECT * FROM rai_host
```

NODE_NAME	ADMIN_IP_ADDRESS	BACKUP_IP_ADDRESS
raadm01.example.com	10.10.40.171	10.10.62.139
raadm02.example.com	10.10.40.172	10.10.62.140

```

REPLICATION_IP_ADDRESS
-----
10.10.241.204
10.10.241.205

```

3. Update the `BACKUP_IP_ADDRESS` column with the VIP associated with the InfiniBand network for each host. Use the appropriate values for the IP addresses and compute server (node) names for your installation:

```
SQL> UPDATE rai_host
      SET backup_ip_address='192.168.40.79,'|| backup_ip_address
      WHERE node_name = 'raadm01.example.com';

SQL> UPDATE rai_host
      SET backup_ip_address='192.168.40.80,'||backup_ip_address
      WHERE node_name = raadm02.example.com'
```

4. Save your changes to the table:

```
SQL> commit;
```

5. Query `RAI_HOST` to verify that the settings are correct:

```
SQL> SELECT * FROM rai_host;
```

NODE_NAME	ADMIN_IP_ADDRESS	BACKUP_IP_ADDRESS
raadm01.example.com	10.10.40.171	192.168.40.79,10.10.62.139
raadm02.example.com	10.10.40.172	192.168.40.80,10.10.62.140

REPLICATION_IP_ADDRESS
10.10.241.204
10.10.241.205

9.5 Update the Database Initialization Parameters

Recovery Appliance processes database backup and restore requests using dispatchers, which prevent an excessive number of dedicated connections to the appliance database.

The `init.ora` file defines the dispatchers for the default ingest network as shown in this example:

```
*.dispatchers=''
ral.dispatchers='(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=raingest01-
vip.example.com))(SDU=65536))(SERVICE=ZDLRAXDB)(DISPATCHERS=4)'
ra2.dispatchers='(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=raingest02-
vip.example.com))(SDU=65536))(SERVICE=ZDLRAXDB)(DISPATCHERS=4)'
```

You can use the SQL `ALTER SYSTEM` command to add the dispatchers for the InfiniBand network to the `init.ora` parameter settings.

To add the dispatchers for InfiniBand:

- While still connected to Oracle Database as `RASYS`, configure the dispatchers for the InfiniBand VIPs as shown in these examples:

```
SQL> ALTER SYSTEM SET dispatchers=
'(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
(HOST=raadm01-priv-vip.example.com))(SDU=65536))
(SERVICE=ZDLRAXDB)(DISPATCHERS=4)' SCOPE=BOTH SID='ral';
```

```
SQL> ALTER SYSTEM SET dispatchers=
'(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)
(HOST=raadm02-priv-vip.example.com))(SDU=65536))
(SERVICE=ZDLRAXDB)(DISPATCHERS=4)' SCOPE=BOTH SID='ra2';
```

10

Configuring a Recovery Appliance Rack

This chapter describes how to configure the hardware components of a Recovery Appliance rack. It contains the following sections:

- [Supporting Auto Service Request](#) (page 10-1)
- [Installing the Tape Hardware](#) (page 10-2)
- [Connecting Recovery Appliance to Your Networks](#) (page 10-3)
- [Configuring the InfiniBand Switches](#) (page 10-3)
- [Configuring the Cisco Ethernet Switch](#) (page 10-10)
- [Configuring the Power Distribution Units](#) (page 10-16)
- [Checking the Health of the Compute Servers](#) (page 10-19)
- [Checking the Health of the Storage Servers](#) (page 10-22)
- [Verifying the InfiniBand Network](#) (page 10-24)
- [Configuring the IP Addresses of the Servers](#) (page 10-24)

Note:

The procedures in this chapter use the files generated by Oracle Exadata Deployment Assistant. You must run this utility before doing the procedures in this chapter.

10.1 Supporting Auto Service Request

Auto Service Request is an optional component of Recovery Appliance. To configure Recovery Appliance for Auto Service Request, ASR Manager must be installed first.

10.1.1 Prerequisites for Using Auto Service Request

Verify that Auto Service Request was selected for use in Oracle Exadata Deployment Assistant. Recovery Appliance cannot also be used with Oracle Advanced Support Gateway or Oracle Platinum Gateway.

You must know the IP address and the root password of the ASR Manager host.

10.1.2 Checking an Existing ASR Manager Installation

If ASR Manager is already operating at the site, then verify that it is version 4.5 or higher. Otherwise, you must upgrade it.

To obtain the version number of ASR Manager:

- On a Linux system:

```
# rpm -qa | grep SUNWswasr  
SUNWswasr-2.7-1
```

- On a Solaris system:

```
# pkginfo -l SUNWswasr  
PKGINST: SUNWswasr  
NAME: SASM ASR Plugin  
CATEGORY: application  
ARCH: all  
VERSION: 2.6  
BASEDIR: /  
VENDOR: Sun Microsystems, Inc.  
.  
.  
.
```

The output from the previous examples indicate that ASR Manager must be updated to 4.5 or higher.

10.1.3 Installing ASR Manager

If ASR Manager is not already installed, then follow the instructions in [Setting Up Auto Service Request](#) (page 4-1). After you register ASR Manager with the Oracle ASR back end, return to these instructions for configuring Recovery Appliance.

10.2 Installing the Tape Hardware

Oracle Secure Backup tape backup is an option to Recovery Appliance. You must install the QLogic ZLE8362 fiber cards and transceivers on site; they are not factory installed.

The QLogic fiber cards are shipped from Oracle as ride-alongs with the rack. The transceivers are shipped directly from the supplier.

To install the tape networking hardware:

1. Insert a fiber card into PCIe Slot 1 of each compute server.
2. Verify that the cards are seated properly and align with the adjacent cards.
3. Install the transceivers in the SAN switch or tape library, and cable them to the fiber cards.

If the SAN switch and tape library are not installed yet, then contact your supplier.



See Also:

My Oracle Support Doc ID 1592317.1 for full instructions about replacing a PCIe card

10.3 Connecting Recovery Appliance to Your Networks

Before configuring the individual devices in the Recovery Appliance rack, ensure that there are no IP address conflicts between the factory settings for the rack and the existing network.

To prepare the Recovery Appliance rack for configuration:

1. Connect the Recovery Appliance rack to the data center networks.
2. Log in to the system where you ran the Oracle Exadata Deployment Assistant, or to another system where you copied the generated files. The system must be on the same subnet as Recovery Appliance.

See [Using the Deployment Assistant](#) (page 5-1) if you have not run Oracle Exadata Deployment Assistant.

3. Run the checkip script, using a command like the following:

```
# ./customer_name-rack_name-checkip.sh -cf configuration_file.xml
```

In the preceding command, *configuration_file* is the name of the XML configuration file generated by the Oracle Exadata Deployment Assistant. On some platforms, the script has a *.cmd* file name extension.

4. Resolve all reported conflicts and rerun the `checkip` script to ensure that no conflicts remain.

If you are unable to resolve all conflicts, then you can perform the procedures in this chapter without connecting to the network. However, this method takes more time to complete, because of time outs from attempts to connect to services such as Domain Name System (DNS) and network time protocol (NTP).

5. To configure the rack units as described in this chapter, plug a laptop in to a unit to configure.
6. To complete the configuration procedures in this chapter, connect a laptop to the device to configure.

See Also:

"[Factory IP Address Settings](#) (page 3-10)"

10.4 Configuring the InfiniBand Switches

The two Sun Datacenter InfiniBand Switch 36 leaf switches are identified in Recovery Appliance as *iba* and *ibb*. Complete these configuration procedures for both switches:

- [Configuring an InfiniBand Switch](#) (page 10-4)
- [Setting the Time on an InfiniBand Switch](#) (page 10-5)
- [Checking the Health of an InfiniBand Switch](#) (page 10-7)

10.4.1 Configuring an InfiniBand Switch

The default identifier for leaf switch 1 in U20 is `iba`, and for leaf switch 2 in U22 is `ibb`.

To configure a Sun Datacenter InfiniBand Switch 36 switch:

1. For a one-rack installation, unplug the InfiniBand cable from Port 8B on the InfiniBand leaf switches. Use hook-and-loop tape to hang it out of the way.

This cable is preinstalled for a connection to the spine switch in a multirack installation. However, in a one-rack installation, the monitoring software might show it as a down link.

2. Connect a serial cable between your laptop and the InfiniBand switch USB serial adapter. Use these terminal settings, as needed:

```
TERM=vt100; export TERM
```

3. Log in to Oracle ILOM:

```
localhost: ilom-admin  
password: welcome1
```

4. Ensure that the firmware version is 2.1.5-1 or later:

```
-> version
```

If the switch has a lower version than 2.1.5-1, then contact Oracle Support Services.

5. Set the switch host name, without the domain name. The following example assigns the name `ralsw` to the first gateway switch (`iba`):

```
-> set /SP hostname=ralsw-iba  
-> show /SP hostname  
/SP  
Properties:  
hostname = ralsw-iba
```

See the Installation Template for the name of the switch.

6. Set the DNS server and domain names. In the following syntax, *IP_addresses* can have up to three IP addresses, separated by commas, in the preferred search order.

```
-> set /SP/clients/dns auto_dns=enabled  
-> set /SP/clients/dns nameserver=IP_addresses  
-> set /SP/clients/dns searchpath=domain_name
```

7. Verify the settings:

```
-> show /SP/clients/dns  
/SP/clients/dns  
Targets:  
Properties:  
auto_dns = enabled  
nameserver = 10.196.23.245, 138.2.202.15  
retries = 1  
searchpath = example.com  
timeout = 5  
.  
.  
.
```

- Configure the switch management network settings. In the following commands, *pending_ip*, *pending_gw*, and *pending_nm* are IP addresses defined by the network administrator:

```
-> cd /SP/network
-> set pendingipaddress=pending_ip
-> set pendingipgateway=pending_gw
-> set pendingipnetmask=pending_nm
-> set pendingipdiscovery=static
-> set commitpending=true
```

- Verify the settings:

```
-> show
/SP/network
Targets:
test
Properties:
commitpending = (Cannot show property)
dhcp_server_ip = none
ipaddress = 10.196.16.152
ipdiscovery = static
ipgateway = 10.196.23.254
ipnetmask = 255.255.248.0
macaddress = 00:E0:4B:38:77:7E
pendingipaddress = 10.196.16.152
pendingipdiscovery = static
pendingipgateway = 10.196.23.254
pendingipnetmask = 255.255.248.0
state = enabled
.
.
.
```

- If any of the values are wrong, repeat the `set pendingipparameter` command, and then the `commitpending=true` command.

10.4.2 Setting the Time on an InfiniBand Switch

To set the time on an InfiniBand switch:

- Set the time zone, using the value shown in the Installation Template. The following commands display the current setting, change the time zone, and verify the new setting:

```
-> show /SP/clock
-> set /SP/clock timezone=zone identifier
-> show /SP/clock
```

The Oracle Exadata Deployment Assistant generates the Installation Template. See "[About the Configuration Files](#) (page 5-4)".

- Set the SP clock to the current time. Use the time format *MMddHHmmCCyy*, indicating the month, day, hour, minute, century, and year. The following commands display the current setting, change the time, and verify the new setting:

```
-> show /SP/clock
-> set /SP/clock datetime=MMddHHmmCCyy
-> show /SP/clock
```

- Configure NTP. The following commands configure both the primary (1) and the secondary (2) NTP servers:

```
-> set /SP/clients/ntp/server/1 address=IP_address
-> set /SP/clients/ntp/server/2 address=IP_address
-> set /SP/clock usntpserver=enabled
```

 **Note:**

If the network does not use NTP, then configure the first compute server (U16) as an NTP server before you install the software in [Installing the Recovery Appliance Software](#) (page 11-1).

4. Verify the IP address of the primary NTP server:

```
-> show /SP/clients/ntp/server/1
/SP/clients/ntp/server/1
Targets:

Properties:
  address = 10.204.74.2

Commands:
  cd
  set
  show
```

5. Verify the IP address of the secondary NTP server:

```
-> show /SP/clients/ntp/server/2
/SP/clients/ntp/server/2
Targets:

Properties:
  address = 10.196.16.1
  .
  .
  .
```

6. Verify the time:

```
-> show /SP/clock
/SP/clock
Targets:

Properties:
  datetime = Mon Nov 04 11:53:19 2013
  timezone = EST (US/Eastern)
  usntpserver = enabled
  .
  .
  .
```

10.4.3 Setting the Serial Number on a Spine Switch

In a multitrack configuration, set the rack master serial number in the ILOM of the spine switch. Skip this procedure when configuring the leaf switches.

To set the serial number on the spine switch:

1. Set the system identifier to 40 characters or fewer:

```
-> set /SP system_identifier="Oracle ZDLRA X5 serial_number"
```

An invalid property value error indicates too many characters.

2. Verify that the value is set:

```
-> show /SP system_identifier
    /SP
    Properties:
        system_identifier = Oracle ZDLRA X5 AK012345678
```

10.4.4 Checking the Health of an InfiniBand Switch

To check the health of an InfiniBand switch:

1. Open the fabric management shell:

```
-> show /SYS/Fabric_Mgmt
NOTE: show on Fabric_Mgmt will launch a restricted Linux shell.
User can execute switch diagnosis, SM Configuration and IB
monitoring commands in the shell. To view the list of commands,
use "help" at rsh prompt.
Use exit command at rsh prompt to revert back to
ILOM shell.
FabMan@hostname->
```

The prompt changes from -> to FabMan@hostname->

2. Check the general health of the switch:

```
FabMan@ralsw-iba-> showunhealthy
OK - No unhealthy sensors
```

3. Check the general environment.

```
FabMan@ralsw-iba-> 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
.
.
.
```

The report should show that fans 1, 2, and 3 are present, and fans 0 and 4 are not present. All OK and Passed results indicate that the environment is normal.

4. Determine the current InfiniBand subnet manager priority of the switch. Leaf switches must have an smpriority of 5, and spine switches must have a smpriority of 8. The sample output shown here indicates the correct priority for a leaf switch.

```
FabMan@ralsw-iba-> setsmpriority list
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xfe80000000000000
```

5. If the priority setting is incorrect, then reset it:

- a. Disable the subnet manager:

```
FabMan@ralsw-iba-> disablesm
Stopping partitiond daemon.          [ OK ]
Stopping IB Subnet Manager..        [ OK ]
```

- b. Reset the priority. This example sets the priority on a leaf switch:

```
FabMan@ralsw-iba->setsmpriority 5
Current SM settings:
smpriority
5 controlled_handover TRUE
subnet_prefix 0xfe80000000000000
```

- c. Restart the subnet manager:

```
FabMan@ralsw-iba->enablesm
Starting IB Subnet Manager.          [ OK ]
Starting partitiond daemon.         [ OK ]
```

6. Log out of the Fabric Management shell and the Oracle ILOM shell:

```
FabMan@ralsw-iba-> exit
-> exit
```

7. Log in to Linux as root and restart the switch:

```
localhost: root
password: welcome1
[root@localhost ~]# reboot
```

8. Disconnect your laptop from the InfiniBand switch.
9. Repeat these procedures for the second InfiniBand leaf switch.

10.4.5 Setting a Spine Switch as the Subnet Manager Master

The InfiniBand switch located in rack unit 1 (U1) is the spine switch. Recovery Appliance has a spine switch only when it is connected to another Recovery Appliance. It is not included as a basic component of the rack.

Perform these steps after the racks are cabled together

The spine switch is the Subnet Manager Master for the InfiniBand subnet. The Subnet Manager Master has priority 8.

To verify the priority setting of the spine switch:

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 other two InfiniBand switches are the leaf switches. The leaf switches are located in rack units 20 and 22 (U20 and U22). They are the Standby Subnet Managers with a priority of 5. You can verify the status using the preceding procedure, substituting a value of 5 in the command shown in step 2 (page 10-8) 2.b (page 10-8).

To determine the Subnet Manager Master:

1. Log in as the `root` user on any InfiniBand switch.
2. Display the location of the Subnet Manager Master.

```
# getmaster
20100701 11:46:38 OpenSM Master on Switch : 0x0021283a8516a0a0 ports 36 Sun DCS
36
QDR switch ra01sw-ib1.example.com enhanced port 0 lid 1 lmc 0
```

The preceding output shows the proper configuration. The Subnet Manager Master is running on spine switch ra01sw-ib1.example.com.

If the spine switch is not the Subnet Manager Master, then reset 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.

 **See Also:**

"Disable the Subnet Manager" in *Sun Datacenter InfiniBand Switch 36 User's Guide* at

<http://docs.oracle.com/cd/E19197-01/835-0784-05/z4001de61813698.html#z40003f12047367>

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 (page 10-9) and 3 (page 10-9) until the spine switch is the Subnet Manager Master.
5. Enable Subnet Manager on the leaf switches that were disabled during this procedure.

 **See Also:**

"Enable the Subnet Manager" in *Sun Datacenter InfiniBand Switch 36 User's Guide* at

<http://docs.oracle.com/cd/E19197-01/835-0784-05/z4001de61707660.html#z40003f12047359>

 **Note:**

If the InfiniBand network consists of four or more racks cabled together, then only the spine switches run Subnet Manager. Disable the Subnet Manager on the leaf switches.

10.5 Configuring the Cisco Ethernet Switch

The Cisco Catalyst 4948 Ethernet switch supplied with Recovery Appliance has IPBASEK9-MZ firmware. The switch is minimally configured during installation. These procedures configure the Cisco Ethernet switch into one large virtual LAN.

The Cisco Ethernet switch configuration consists of these topics and procedures:

- [Scope of the Configuration](#) (page 10-10)
- [Prerequisites for Configuring the Ethernet Switch](#) (page 10-10)
- [Configuring the Ethernet Switch on the Data Center Network](#) (page 10-11)
- [Setting the Clock and Time Zone on the Ethernet Switch](#) (page 10-14)
- [Verifying the Ethernet Configuration](#) (page 10-16)

10.5.1 Scope of the Configuration

The minimal configuration disables IP routing, and sets the following:

- Host name
- IP address
- Subnet mask
- Default gateway
- Domain name
- Name server
- NTP server
- Time
- Time zone

10.5.2 Prerequisites for Configuring the Ethernet Switch

To avoid disrupting the customer network, observe these prerequisites:

- Do not connect the Cisco Ethernet switch until the network administrator has verified the running configuration and has made any necessary changes.
- Do not connect the Cisco Ethernet switch until the IP addresses on all components of Recovery Appliance are configured. This sequence prevents any duplicate IP address conflicts, which might occur with the factory settings.
- Configure the Cisco Ethernet switch with the network administrator.

See Also:

- ["Network Connection Requirements for Recovery Appliance](#) (page 2-15)"
- [Zero Data Loss Recovery Appliance Site Checklists](#)

10.5.3 Configuring the Ethernet Switch on the Data Center Network

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 RJ45 to DB9 serial cable is included with the Cisco documentation package. The default serial port speed is 9600 baud, 8 bits, no parity, 1 stop bit, and no handshake.

The switch displays status lights and the following prompt:

```
Switch con0 is now available
Press RETURN to get started.
```

2. Change to enable mode:

```
Switch> enable
Password: welcome1
```

3. Verify that the switch is running the current firmware version used by Oracle:

```
Switch# show version
Cisco IOS Software, Catalyst 4500 L3 Switch Software (cat4500e-
IPBASEK9-M), Version 15.2(1)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
ralsw-ip uptime is 2 weeks, 6 hours, 9 minutes
System returned to ROM by power-on
System restarted at 13:06:29 EDT Tue Oct 29 2013
System image file is "bootflash:cat4500e-ipbasek9-mz.152-1.E2.bin"
Hobgoblin Revision 22, Fortooin Revision 1.40
.
.
.
```

The previous output shows the current version, which is IPBASE9-M. The full release version name is cat4500e-ipbasek9-mz.152-1.E2.bin.

4. Configure the network for a single VLAN. The following example uses 10.7.7.34 for the IP address of the switch and 255.255.255.0 for the netmask:

```
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 ]
```

5. If the data center network does *not* require IP routing, then disable the default IP routing setting 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)# 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 ]
```

6. If IP routing is required, 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 ]
```

7. Set the host name of the switch. This example sets the name to ra1sw-ip:

```
Switch# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
Switch(config)# hostname ra1sw-ip
ra1sw-ip(config)# end
ra1sw-ip# write memory
Building configuration...
Compressed configuration from 3789 bytes to 1469 bytes [OK ]
```

The prompt changes to the system host name.

8. Configure up to three DNS servers. Replace the domain name and IP addresses used in this example with the values for the installation:

```
ra1sw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
ra1sw-ip(config)# ip domain-name example.com
ra1sw-ip(config)# ip name-server 10.7.7.3
ra1sw-ip(config)# ip name-server 129.148.5.5
ra1sw-ip(config)# ip name-server 10.8.160.1
ra1sw-ip(config)# end
*Sep 15 14:26:37.045:%SYS-5-CONFIG_I:Configured from console by console
ra1sw-ip# write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes [OK ]
```

If no DNS service is available, then set just the domain name, so that you can configure the SSH keys.

9. Set a temporary password, if it was not factory set:

```
ra1sw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
ra1sw-ip(config)# enable password password
ra1sw-ip(config)# enable secret password
ra1sw-ip(config)# end
ra1sw-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 ]
```

10.5.4 Disabling Telnet Connections

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.

To disable Telnet connections:

1. Verify that telnet is disabled:

```
ralsw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ralsw-ip(config)# line vty 0 15
ralsw-ip(config-line)# 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
% end
```

The previous example shows that no password is set, and therefore telnet is disabled.

2. Disable login if a telnet password is set:

```
ralsw-ip(config-line)# no password
ralsw-ip(config-line)# end
ralsw-ip# write memory
Building configuration...
Compressed configuration from 3786 bytes to 1468 bytes[OK]
```

10.5.5 Configuring SSH on the Ethernet Switch

To configure a secure shell (SSH) on the Ethernet switch:

- Enter the commands shown in this example:

```
ralsw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ralsw-ip(config)# crypto key generate rsa
% You already have RSA keys defined named ralsw-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]
ralsw-ip(config)# username admin password 0 welcome1
ralsw-ip(config)# line vty 0 15
ralsw-ip(config-line)# transport input ssh
ralsw-ip(config-line)# exit
ralsw-ip(config)# aaa new-model

ralsw-ip(config)# ip ssh time-out 60
ralsw-ip(config)# ip ssh authentication-retries 3
ralsw-ip(config)# ip ssh version 2
ralsw-ip(config)# end
*Sep 15 14:26:37.045: %SYS-5-CONFIG_I: Configured from console by console
```

```
ralsw-ip# write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes[OK]
```

10.5.6 Setting the Clock and Time Zone on the Ethernet Switch

To set the time on the Cisco Ethernet switch:

1. Set the clock and time zone. The switch keeps internal time in Coordinated Universal Time (UTC) format.

- To set the clock time to UTC:

```
no clock timezone global configuration
```

- To set the clock to use a time zone

```
clock timezone zone hours-offset [minutes-offset]
```

In the preceding command, *zone* is the time zone to display when standard time is in effect, *hours-offset* is the hours offset from UTC, and *minutes-offset* is the minutes offset from UTC.

- To set summer time hours:

```
clock summer-time zone recurring [week day month hh:mm week day month \
hh:mm [offset]]
```

In the preceding command, *zone* is the time zone to use when summer time is in effect, *week* is the week of the month (1 to 5, or *last*), *day* is the day of the week, *month* is the month, *hh:mm* is the time 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:

```
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 of the month, *month* is the month, and *year* is the year. The time is relative to the configured time zone.

When you set the local time and the time zone, ordering is important. The following example sets the local time to the US Eastern time zone:

```
Switch# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
rasw-ip(config)# clock timezone EST -5
rasw-ip(config)# clock summer-time EDT recurring
rasw-ip(config)# end
rasw-ip# clock set 21:00:00 July 17 2014
rasw-ip# write memory
Building configuration...
Compressed configuration from 3784 bytes to 1465 bytes [OK ]
rasw-ip# show clock
21:00:06.643 EST Thu Jul 17 2014
```

2. Configure up to two NTP servers. Replace the IP addresses used in this example with the values for the installation:

```
rasw-ip# configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
rasw-ip(config)# ntp server 10.7.7.32 prefer
```

```
rasw-ip(config)# ntp server 129.148.9.19
rasw-ip(config)# end
*Sep 15 14:51:08.665:%SYS-5-CONFIG_I:Configured from console by console
rasw-ip# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes [OK ]
rasw-ip# show ntp status
.
.
.
rasw-ip# show clock
20:59:06.643 EST July 17 2014
```

The NTP server is synchronized to local time when you connect the Cisco switch to the network and it has access to NTP. Symbols that precede the `show clock` display indicate that the time is the following:

- * Not authoritative
- . Authoritative, but NTP is not synchronized.
- Authoritative (blank space).

10.5.7 Disabling Spanning Tree on the Ethernet Switch

Spanning tree is enabled by default on Cisco switches. If you add a switch with spanning tree enabled to the network, then you might cause network problems. As a precaution, you can disable spanning tree from the uplink port VLAN before connecting the switch to the network. Alternatively, you can turn on spanning tree protocol with specific protocol settings either before or after connecting to the network.

To disable spanning tree on the uplink port VLAN:

1. Disable spanning tree on the uplink port VLAN:

```
rasw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
rasw-ip(config)# no spanning-tree vlan 1
rasw-ip(config)# end
rasw-ip# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes[OK]
```

2. Verify that spanning tree is disabled:

```
rasw-ip# show spanning-tree vlan 1
Spanning tree instance(s) for vlan 1 does not exist.
```

To re-enable spanning tree protocol with the default protocol settings:

- Use the commands shown in this example:

```
ralsw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ralsw-ip(config)# spanning-tree vlan 1
ralsw-ip(config)# end
ralsw-ip# write memory
```

 **See Also:**

Cisco Switch Configuration Guide to enable spanning tree protocol with the specific protocol settings required by the data center Ethernet network

10.5.8 Verifying the Ethernet Configuration

To verify the Ethernet configuration:

1. Verify the configuration:

```
rasw-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
.
.
.
```

2. If any setting is incorrect, then repeat the appropriate command. To erase a setting, enter `no` before the same command. For example, the following commands erase the default gateway setting:

```
rasw-ip# no ip default-gateway 10.7.7.1
rasw-ip# end
rasw-ip# write memory
```

3. Save the current configuration:

```
rasw-ip# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
Compressed configuration from 2654 bytes to 1189 bytes[OK]
```

4. Exit from the session:

```
rasw-ip# exit
rasw-ip con0 is now available
Press RETURN to get started.
```

5. Disconnect the cable from the Cisco console.
6. (Optional) Attach your laptop to port 48 and ping the IP address of the switch on the internal management network. Do not connect the switch to the management network.

10.6 Configuring the Power Distribution Units

The power distribution units (PDUs) are configured with static IP addresses to connect to the network for monitoring.

10.6.1 Assigning Network Addresses to the PDUs

To configure the PDU network addresses:

1. Use an RS-232 cable to connect your laptop to the Cisco Ethernet switch SER MGT port.
2. Configure your laptop's terminal emulator to use these settings:
 - 9600 baud
 - 8 bit
 - 1 stop bit
 - No parity bit
 - No flow control
3. Log in to the PDU metering unit as the `admin` user with password `welcome1`.
Change this password after configuring the network.
4. Enter the network settings for the IP address, subnet mask, and default gateway:

```
pducli -> set net_ipv4_dhcp=Off
set OK
pducli -> set net_ipv4_ipaddr=ip_address
set OK
pducli -> set net_ipv4_subnet=subnet_mask
set OK
pducli -> set net_ipv4_gateway=default_gateway
set OK
```

5. (Optional) Configure the PDU with the DNS server IP addresses:

```
pducli -> set net_ipv4_dns1=domain_name_1
set OK
pducli -> set net_ipv4_dns2=domain_name_2
set OK
```

6. Reset the PDU metering unit:

```
pducli -> reset=yes
set OK
```

7. Remove the RS-232 cable from the SER MGT port.
8. Repeat these steps for the second PDU metering unit.

10.6.2 Configuring the PDU System Time Settings

To configure the PDUs:

1. Connect your laptop to the Ethernet switch.
2. Open a browser and connect to the PDU, using its IP address:

```
https://pdu_ip_address
```

Accept the security note. The Metering Overview page is displayed.

3. Click **Net Configuration** in the upper left, and log in as user `admin` with the password `welcome1`.

4. Select the **System Time** tab.
5. Configure Manual Settings with the current date and time, and then click **Submit**.
6. Configure NTP Server Settings, and then click **Submit**:
 - Select the Enable option.
 - Enter an NTP server IP address, which is listed on the Installation Template.
 - Select **Time Zone** from the menu.
7. Select the **PDU Information** tab.
8. Enter these values, and then click **Submit**:
 - **Name**: PDU host name, provided in the Installation Template; for example, ra5sw-pdua0
 - **Product Identifier** (case sensitive): ZDLRA X5
 - **Rack Serial Number**: Serial number similar to AK12345678
 - **Location** (optional): Site identifier
9. On the Metering Overview page, select **Module Info**.
10. Confirm that the firmware version is 2.01 or higher. If it is not, then upgrade the firmware after you finish this procedure.
11. Click **Logout** to log out of the PDU.
12. Repeat these steps for the second PDU metering unit.
13. Disconnect the PDU metering units from the Cisco Ethernet switch, and connect them to the data center management network.

10.6.3 Upgrading the PDU Firmware

If the PDU firmware is out of date, with a version earlier than 2.01, then download and install the current version.

To upgrade the PDU firmware:

1. Download the current firmware for Enhanced PDUs from My Oracle Support to your laptop.
2. Unzip the downloaded file on your laptop.
3. Open a browser and connect to the PDU, using its IP address:
`https://pdu_ip_address`
Accept the security note. The Metering Overview page is displayed.
4. Click **Net Configuration** in the upper left, and log in as user `admin` with the password `welcome1`.
5. Select the **Firmware Update** tab.
6. Click **Browse**, and select `MKAPP_V2.x.DL` from the unzipped, downloaded firmware files on your laptop.
7. Click **Submit** to update the firmware.
The PDU reboots automatically when the update is complete.
8. Reconnect your browser to the PDU.

9. Click Module Info on the Metering Overview page, and confirm that the firmware was updated successfully.
10. Click **Logout** to log out of the PDU.
11. Repeat these steps for the second PDU metering unit.

10.7 Checking the Health of the Compute Servers

To check the two compute servers in U16 and U17:

1. Power on both compute servers if they are not up already, and wait while they initialize the BIOS and load the Linux operating system.
2. Use a serial cable to connect your laptop to the first compute server's serial MGT port.
3. Configure your laptop's terminal emulator to use these settings:
 - 9600 baud
 - 8 bit
 - 1 stop bit
 - No parity bit
 - No handshake
 - No flow control
4. Log in as the `root` user with the `welcome1` password.
 - On the first compute server (which is connected to your laptop), open the Oracle ILOM console, and then log in:

```
-> start /SP/console
```
 - On the second compute server, use SSH to log in. The default factory IP address is 192.168.1.109.
5. Verify that the rack master and host serial numbers are set correctly. The first number must match the rack serial number, and the second number must match the SysSN label on the front panel of the server.

```
# ipmitool sunoem cli "show /System" | grep serial
serial_number = AK12345678
component_serial_number = 1234NM567H
```

6. Verify that the model and rack serial numbers are set correctly:

```
# ipmitool sunoem cli "show /System" | grep model
model = ZDLRA X5
# ipmitool sunoem cli "show /System" | grep ident
system_identifiar = Oracle Zero Data Loss Recovery Appliance X5 AK12345678
```

7. Verify that the management network is working:

```
# ethtool eth0 | grep det
Link detected: yes
```

8. Verify that the ILOM management network is working:

```
# ipmitool sunoem cli 'show /SP/network' | grep ipadd
ipaddress = 192.168.1.108
pendingipaddress = 192.168.1.108
```

9. Verify that Oracle ILOM can detect the optional QLogic PCIe cards, if they are installed:

```
# ipmitool sunoem cli "show /System/PCI_Devices/Add-on/Device_1"
Connected. Use ^D to exit.
-> show /System/PCI_Devices/Add-on/Device_1
/System/PCI_Devices/Add-on/Device_1
Targets:

Properties:
  part_number = 7101674
  description = Sun Storage 16 Gb Fibre Channel PCIe Universal FC HBA,
                Qlogic
  location = PCIE1 (PCIe Slot 1)
  pci_vendor_id = 0x1077
  pci_device_id = 0x2031
  pci_subvendor_id = 0x1077
  pci_subdevice_id = 0x024d

Commands:
  cd
  show

-> Session closed
Disconnected
```

See "[Installing the Tape Hardware](#) (page 10-2)" for information about the QLogic PCIe cards.

10. Verify that all memory is present (256 GB):

```
# grep MemTotal /proc/meminfo
MemTotal: 264232892 kB
[
```

The value might vary slightly, depending on the BIOS version. However, if the value is smaller, then use the Oracle ILOM event logs to identify the faulty memory.

11. Verify that the four disks are visible, online, and numbered from slot 0 to slot 3:

```
# cd /opt/MegaRAID/MegaCli/
# ./MegaCli64 -Pdlist -a0 | grep "Slot\Firmware state"
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
```

12. Verify that the hardware logical volume is set up correctly. Look for Virtual Disk 0 as RAID5 with four drives and no hot spares:

```
[root@db01 ~]# cd /opt/MegaRAID/MegaCli
[root@db01 MegaCli]# ./MegaCli64 -LdInfo -lAll -a0
Adapter 0 -- Virtual Drive Information:
Virtual Drive: 0 (Target Id: 0)
Name :DBSYS
RAID Level : Primary-5, Secondary-0, RAID Level Qualifier-3
Size : 1.633 TB
Physical Sector Size: 512
```

```
Logical Sector Size : 512
VD has Emulated PD : No
Parity Size : 557.861 GB
State : Optimal
Strip Size : 1.0 MB
Number Of Drives : 4
Span Depth : 1
.
.
.
```

13. Verify that the hardware profile is operating correctly:

```
# /opt/oracle.SupportTools/CheckHWnFWProfile
[SUCCESS] The hardware and firmware matches supported profile for
server=ORACLE_SERVER_X5-2
```

The previous output shows correct operations. However, the following response indicates a problem that you must correct before continuing:

```
[WARNING] The hardware and firmware are not supported. See details below
[InfinibandHCAPCIeSlotWidth]
Requires:
x8
Found:
x4
[WARNING] The hardware and firmware are not supported. See details above
```

Use the `--help` argument to review the available options, such as obtaining more detailed output.

14. When connected to the first compute server only:

a. Verify the IP address of the first compute server:

```
# ifconfig eth0
eth0 Link encap:Ethernet HWaddr 00:10:E0:3C:EA:B0
      inet addr:172.16.2.44 Bcast:172.16.2.255 Mask:255.255.255.0
      inet6 addr: fe80::210:e0ff:fe3c:eab0/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:7470193 errors:0 dropped:0 overruns:0 frame:0
      TX packets:4318201 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:872195171 (831.7 MiB) TX bytes:2444529519 (2.2 GiB)
```

b. Verify the IP address of the second compute server:

```
# ibhosts
Ca : 0x0010e0000159c61c ports 2 "node4 elasticNode 172.16.2.40,172.16.2.40
ETH0"
Ca : 0x0010e000015a46f0 ports 2 "node10 elasticNode 172.16.2.46,172.16.2.46
ETH0"
Ca : 0x0010e0000159d96c ports 2 "node1 elasticNode 172.16.2.37,172.16.2.37
ETH0"
Ca : 0x0010e0000159c51c ports 2 "node2 elasticNode 172.16.2.38,172.16.2.38
ETH0"
Ca : 0x0010e000015a5710 ports 2 "node8 elasticNode 172.16.2.44,172.16.2.44
ETH0"
```

15. Disconnect from the server:

- First compute server: `exit`
- Second compute server: `logout`

16. Repeat these steps for the second compute server.

10.8 Checking the Health of the Storage Servers

A Recovery Appliance X5 rack has three to 18 storage servers, and a Recovery Appliance X4 rack has three to 14 storage servers. Begin at the bottom of the rack and check each server.

To check a storage server:

1. Power on all storage servers if they are not already on, and wait while the servers initialize the BIOS and load the Linux operating system.
2. Use SSH to connect your laptop to the first storage server. Use its factory IP address.

See "[Factory IP Address Settings \(page 3-10\)](#)".

3. Log in as the `root` user with the `welcome1` password.

The terminal emulation settings are the same as for the compute servers. See "[Checking the Health of the Compute Servers \(page 10-19\)](#)".

4. Verify that the rack master and host serial numbers are set correctly. The first number must match the rack serial number, and the second number must match the SysSN label on the front panel of the server.

```
# ipmitool sunoem cli "show /System" | grep serial
serial_number = AK01234567
component_serial_number = 1234NM5678
```

5. Verify that the model and rack serial numbers are set correctly:

```
# ipmitool sunoem cli "show /System" | grep model
model = ZDLRA X5
# ipmitool sunoem cli "show /System" | grep ident
system_identifier = Oracle Zero Data Loss Recovery Appliance X5 AK01234567
```

6. Verify that the management network is working:

```
# ethtool eth0 | grep det
Link detected: yes
```

7. Verify that the ILOM management network is working:

```
# ipmitool sunoem cli 'show /SP/network' | grep ipadd
ipaddress = 192.168.1.101
pendingipaddress = 192.168.1.101
```

8. Verify that all memory is present (96 GB):

```
# grep MemTotal /proc/meminfo
MemTotal: 98757064 kB
[
```

If the value is smaller, then use the Oracle ILOM event logs to identify the faulty memory.

9. Verify that the hardware profile is operating correctly:

```
# /opt/oracle.SupportTools/CheckHwFWProfile
[SUCCESS] The hardware and firmware matches supported profile for
server=ORACLE_SERVER_X5-2L_EXADATA_HIGHCAPACITY
```

The previous output shows correct operations. However, the following response indicates a problem that you must correct before continuing:

```
[WARNING] The hardware and firmware are not supported. See details below
[InfinibandHCAPCIeSlotWidth]
Requires:
x8
Found:
x4
[WARNING] The hardware and firmware are not supported. See details above
```

Use the `--help` argument to review the available options, such as obtaining more detailed output.

10. Verify that 12 disks are visible, online, and numbered from slot 0 to slot 11:

```
# cd /opt/MegaRAID/MegaCli
# ./MegaCli64 -Pdlist -a0 | grep "slot\Firmware state"
Slot Number: 0
Firmware state: Online, Spun Up
Slot Number: 1
Firmware state: Online, Spun Up
.
.
.
```

11. Verify that there are four NVME logical devices:

```
# ls -l /dev | grep nvme | grep brw
brw-rw---- 1 root disk 259, 0 Nov 12 19:10 nvme0n1
brw-rw---- 1 root disk 259, 1 Nov 12 19:10 nvme1n1
brw-rw---- 1 root disk 259, 2 Nov 12 19:10 nvme2n1
brw-rw---- 1 root disk 259, 3 Nov 12 19:10 nvme3n1
```

12. Confirm the healthy status of the AIC card:

```
# nvmecli --identify --all | grep -i indicator
Health Indicator      : Healthy
Health Indicator      : Healthy
Health Indicator      : Healthy
Health Indicator      : Healthy
```

13. Verify that the boot order is USB (Oracle Unigen), RAID, and PXE:

```
[# ubiosconfig export all > /tmp/bios.xml
[# grep -m1 -A20 boot_order /tmp/bios.xml
<boot_order>
  <boot_device>
    <description>USB:USBIN0:ORACLE SSM UNIGEN-UFD PMAP</description>
    <instance>1</instance>
  </boot_device>
  <boot_device>
    <description>RAID:PCIE6:(Bus 50 Dev 00)PCI RAID Adapter</description>
    <instance>1</instance>
  </boot_device>
  <boot_device>
    <description>PXE:NET0:IBA XE Slot 3A00 v2320</description>
    <instance>1</instance>
  </boot_device>
  <boot_device>
    <description>PXE:NET1:IBA XE Slot 4001 v2196</description>
    <instance>1</instance>
  </boot_device>
```

14. If the boot order is wrong, then restart the server and fix the order in the BIOS setup:

```
# ipmitool chassis bootdev bios
# shutdown -r now
```
15. Exit or log out of SSH.
16. Repeat these steps for the next storage server until you have checked all of them.

10.9 Verifying the InfiniBand Network

To verify the InfiniBand network:

1. Visually check all the InfiniBand 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 to any system in the rack.
3. Verify the InfiniBand topology:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify-topology]

[DB Machine Infiniband Cabling Topology Verification Tool ]
Every node is connected to two leaf switches in a single rack.....[SUCCESS]
Every inter-leaf switch link is connected correctly in a single rack.[SUCCESS]
Every leaf switch in an interconnected quarter rack is correctly connected to
other rack in a multi-rack group..... [NOT APPLICABLE]
Every leaf switch is connected to every spine switch in a multi-rack
group..... [SUCCESS]
Every rack has balanced inter-leaf-and-spine switch links in a multi-rack
group..... [NOT APPLICABLE]
No spine switch is connected to another spine switch in a multi-rack
group..... [NOT APPLICABLE]
Every spine switch is connected to two external spine switches in a multi-rack
group..... [NOT APPLICABLE]
No external spine switch is connected to a leaf switch in a multi-rack
group..... [NOT APPLICABLE]
No external spine switch is connected to another external spine switch in a
multi-rack group..... [NOT APPLICABLE]
```

The previous sample output indicates that the network components are cabled correctly, with a status of `SUCCESS` or `NOT APPLICABLE` for each check. A problem appears as a `FAILED` check with additional diagnostic information.

10.10 Configuring the IP Addresses of the Servers

Before installing Recovery Appliance software, you must run a script to configure the compute and storage servers with proper IP addresses. Otherwise, the `install.sh` script will fail when it tries to configure the networks.

Follow the instructions in My Oracle Support Doc ID 1953915.1 to configure the IP addresses.

11

Installing the Recovery Appliance Software

This chapter describes the procedures for installing the Recovery Appliance software. It contains the following sections:

- [About the Installation Software](#) (page 11-1)
- [Software Installation Prerequisites](#) (page 11-2)
- [Reclaiming Disk Space on the Compute Servers](#) (page 11-3)
- [Installing the Software on Recovery Appliance](#) (page 11-3)
- [Install Utility Syntax](#) (page 11-6)
- [Setting Up Enterprise Manager to Administer Recovery Appliance](#) (page 11-10)

Note:

Recovery Appliance software is supported only when installed on Recovery Appliance hardware. See the *Zero Data Loss Recovery Appliance Licensing Information User Manual* for information about software licensing.

11.1 About the Installation Software

You use the following utilities to install the Recovery Appliance software:

- `ra_preinstall.pl`: Performs initial checks and prepares the Recovery Appliance for new installations and upgrades. This script was introduced in Zero Data Loss Recovery Appliance software update 12.1.1.1.7. See *Zero Data Loss Recovery Appliance Upgrade and Patching* (Doc ID 2028931.1) for more information.
- `install.sh`: Uses the files generated by Oracle Exadata Deployment Assistant to configure the networks, Oracle Database, and other software components.

The OneCommand ZIP file that contains Oracle Exadata Deployment Assistant for use in [Using the Deployment Assistant](#) (page 5-1) also contains this script. You can run this script from any system on the same network as Recovery Appliance. Some platforms use `install.cmd`.

- `racli`: Installs the Recovery Appliance software. You must run `racli` while logged in to a Recovery Appliance compute server, and not from another server on the same network.



Note:

Whenever the deployment configuration file is modified such as through a re-import and re-generation of XML, run `ra_preinstall.pl` again before running `install.sh`.

11.2 Software Installation Prerequisites

Before you begin the software installation, ensure that these procedures are complete:

- You used Oracle Exadata Deployment Assistant to generate the configuration files. See [Using the Deployment Assistant](#) (page 5-1).
- You assembled the installation files into the `WorkDir` directory under the OneCommand installation directory. (The OneCommand installation directory is the directory to which you extracted the OneCommand ZIP file, such as `linux-x64`.) See "[Generating the Configuration Files](#) (page 5-2)" for information about the OneCommand installation directory. For the list of installation files, see My Oracle Support Doc ID 1927416.1.
- The `checkip` script ran successfully with no remaining errors. See "[Connecting Recovery Appliance to Your Networks](#) (page 10-3)".
- All racks of a multirack installation are cabled together. See [Extending Recovery Appliance with Multiple Racks](#) (page 8-1).
- All components of all new racks are configured. See [Configuring a Recovery Appliance Rack](#) (page 10-1).
- Oracle Management Server is operating on the same network as Recovery Appliance.
- Reclaim the disk space on all compute servers. See "[Reclaiming Disk Space on the Compute Servers](#) (page 11-3)".
- The Oracle tape library is installed and operational. A tape library is an optional component of Recovery Appliance.

11.2.1 About Adding a Rack to an Existing Rack

When you add a new Recovery Appliance rack to an older rack, the factory installed images are typically different. You can either update the older rack to the new image, or retain the older version and reimage the new rack to match it. In either case, ensure that all servers and all Sun Datacenter InfiniBand Switch 36 switches are at the same patch level.

11.2.2 About Adding New Storage Servers to an Existing Rack

Additional patching considerations include the Grid Infrastructure and database home releases and bundle patch updates. If the software installation applies new patches, then Oracle recommends changing the existing servers so the new servers inherit the releases as part of the extension procedure. This way, the number of servers that require patches is lower. Apply any patches to the existing servers in advance, so that they are at the desired level when the extension work is scheduled, thereby reducing the amount of work required by the extension.

11.3 Reclaiming Disk Space on the Compute Servers

Before you install the Recovery Appliance software, you must run the `reclaimdisks` script on each compute server. The script checks the volume configuration and reallocates all space correctly for the Recovery Appliance. Reclaiming the disk space converts the Linux system to four disks in a RAID-5 configuration at the disk controller level.

Perform the following steps to configure the disks on a compute server:

1. Log in to the compute server as the `root` user.
2. Change to the `/opt/oracle.SupportTools` directory.
3. Start the disk reclamation process:

```
# ./reclaimdisks.sh -free -reclaim
```

The command frees any preconfigured disks, schedules reclamation of all free disks to Linux on the next boot, and restarts Recovery Appliance. The process may take two hours to complete.

Caution:

Do not restart Recovery Appliance or interrupt the reclamation process.

4. Validate the Linux configuration:

```
./reclaimdisks.sh -check
```

The last lines of the output should be the following:

```
[INFO] Valid single boot configuration found for Linux: RAID5 from 4 disks  
with no global and dedicated hot spare disks  
Valid. Booted: Linux. Layout: Linux.
```

5. Repeat these steps on the other compute servers.

11.4 Installing the Software on Recovery Appliance

After completing the prerequisites and assembling the necessary files, you are ready to install the software.

To install the software on a new Recovery Appliance rack:

1. Log in as `root` to the first compute server.
2. Change to the `u01` directory:

```
# cd /u01
```

3. Create a directory named `ra_install` directly under `u01`:

```
# mkdir ra_install
```

4. Extract the latest Linux version of the OneCommand ZIP file into `/u01/ra_install/`, using the default directory name, `linux-x64`. (The ZIP file can be downloaded from My Oracle Support Doc ID 1927416.1.)

5. Copy the XML configuration files that were generated by using the Oracle Exadata Deployment Assistant into the `ExadataConfigurations` subdirectory under `/u01/ra_install/linux-x64`. See "[Generating the Configuration Files](#) (page 5-2)" for information about the configuration files.
6. Run the Recovery Appliance preinstallation script:
 - a. Obtain the latest Recovery Appliance patch file from My Oracle Support Doc ID 1927416.1, and extract it into `/u01/ra_install/linux-x64/WorkDir`.

- b. Change to the `WorkDir` directory:

```
# cd /u01/ra_install/linux-x64/WorkDir
```

Notes:

- The preinstallation script is available with Zero Data Loss Recovery Appliance software update 12.1.1.1.7 and later. See Zero Data Loss Recovery Appliance Upgrade and Patching (Doc ID 2028931.1) for information about the latest patch file
- The files you extracted to `WorkDir` should include a readme file (README.txt) that contains important information to review before you proceed

- c. Run the script with the path of the XML configuration file and the Oracle Exadata Deployment Assistant directory:

```
# /usr/bin/perl ra_preinstall.pl --config_xml=/u01/ra_install/linux-x64/ExadataConfigurations/Example_Inc-ra01.xml
--oeda_dir=/u01/ra_install/linux-x64
```

Note:

Before the `install.sh` has completed all steps, if the fix to a validation issue or any other issue requires re-generating the OEDA XML configuration file, run `ra_preinstall.pl` on the generated XML as given above before re-starting `install.sh` from the beginning.

7. Change to the `linux-x64` directory:

```
# cd /u01/ra_install/linux-x64
```

8. List the steps you need to run.

This example runs the `install` script from a Linux system, using an XML file named `ExadataConfigurations/Example_Inc-ra01.xml`:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -l
```

9. Run each step in numeric order *except* the last step (Resecure Machine), and verify that it completed successfully before continuing to the next step. You will resecure the machine in a later step.

This example runs step 1 from a Linux system, using an XML file named `ExadataConfigurations/Example_Inc-ra01.xml`:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -s 1
```

See "[About the Install Utility Steps](#) (page 11-7)".

10. Change to the Recovery Appliance bin directory:

```
# cd /opt/oracle.RecoveryAppliance/bin
```

11. Run the `racli` utility to install the Recovery Appliance software.

This command runs all of the steps necessary to install the software based on the system and configuration settings:

```
# ./racli install appliance
```

Alternatively, you can run each step performed by this command individually by using the `step` option. However, ensure that you run all of the required steps in order and that each step completes successfully before you run the next step. The following example runs only the first installation step:

```
# ./racli install appliance --step=1
```

See "[racli install appliance](#) (page A-16)" for a description of each step.

 **Note:**

If you encounter a problem while running the `install appliance` command, contact Oracle Support Services at <http://support.oracle.com> for assistance.

12. Change back to the `linux-x64` directory:

```
# cd /u01/ra_install/linux-x64
```

13. Use the `install.sh` utility to run the last step (Resecure Machine) to secure Recovery Appliance.

14. To complete the Recovery Appliance installation:

- a. Change the default passwords for the Recovery Appliance database user (`rasys`) and the OSB tape backup users, if applicable.
- b. Use Oracle Enterprise Manager Cloud Control to deploy the agents and discover the Recovery Appliance targets. Then you can use Cloud Control to monitor and administer the Recovery Appliance environment.

See "[Setting Up Enterprise Manager to Administer Recovery Appliance](#) (page 11-10)".

- c. If you will use VLAN tagging for the Recovery Appliance ingest network, configure and test the VLAN tagging.

See [Enabling 8021.Q VLAN Tagging in Zero Data Loss Recovery Appliance over ingest networks \(Doc ID 2047411.1\)](#) for instructions.

- d. Enroll the Oracle databases, so that they can use Recovery Appliance.

See *Zero Data Loss Recovery Appliance Protected Database Configuration Guide*.

 **See Also:**

- [About Tape Backup Infrastructure](#) (page 1-8)
- [Changing OSB Tape Backup Passwords](#) (page 13-5)
- [Default User Accounts](#) (page B-1)

11.5 Install Utility Syntax

The `install` utility has the following syntax:

```
./install.sh -cf config_filename.xml option1 option2...
```

11.5.1 Install Utility Parameters

The syntax of the `install` utility enables you to run one step or multiple steps, and to undo one or more steps that ran with errors. You can then fix whatever problems caused the errors and run the step again.

[Table 11-1](#) (page 11-6) describes the options.

Table 11-1 Options of the Install Utility

Option	Description
<code>-cf config_filename.xml</code>	Identifies the XML configuration file. Required.
<code>-h</code>	Describes the script options.
<code>-l</code>	List the steps needed to install and configure the software for the configuration described in the XML file.
<code>-r {n-N n}</code>	Runs steps <i>n</i> through <i>N</i> , or step <i>n</i> .
<code>-s n</code>	Runs step <i>n</i> .
<code>-u {n-N n}</code>	Undoes steps <i>n</i> through <i>N</i> , or step <i>n</i> .

11.5.2 Install Utility Syntax Examples

The following examples provide examples of the syntax and sample output. The XML configuration file is named `Example_Inc-ra01.xml` (`linux-x64/ExadataConfigurations/Example_Inc-ra01.xml`).

This command displays Help for the utility:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -h
install.sh -cf <config.xml> -l [options]
install.sh -cf <config.xml> -s <step #> | -r <num-num>
install.sh
ARGUMENTS:
  -l                List all the steps that exist
  -cf              config file name [Full path please]
  -s <step #>     Run the step # at a time
  -r <num-num>    Run the steps one after the other as long as no errors
                  are encountered
```

```
-u <num-num> | <step#> Undo a range of steps or a particular step
-h Usage
Version : 14.254.04:00
```

This command lists the steps:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -l

1. Validate Configuration File
2. Setup Required Files
3. Create Users
4. Setup Cell Connectivity
5. Verify Infiniband
6. Calibrate Cells
7. Create Cell Disks
8. Create Grid Disks
9. Configure 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. Install Exachk
18. Setup ASR Alerting
19. Create Installation Summary
20. Resecure Machine
```

The next example runs step 7:

```
./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -s 7

Executing Create Cell Disks
Check physical disks for errors before creating celldisks.....
Restarting cell services.....
Initializing cells.....
Cleaning cells.....
Restarting cell services.....
Creating cell disks.....
Creating flashlog on cells.....
Creating flashcache on cells.....
Successfully completed execution of step Create Cell Disks
]
```

This command reverses step 15:

```
$ ./install.sh -cf ExadataConfigurations/Example_Inc-ra01.xml -u 15

Undoing Create Databases
Deleting databases...
Deleting database zdlra
.
.
.
```

11.5.3 About the Install Utility Steps

The following descriptions show how each step contributes to the complete software installation. The install script lists the substeps as it performs them.

Step 1 Validate Configuration File

Runs a series of checks to ensure that the configuration file contains settings that are appropriate for the specified cluster on the current network:

- Validates the host names and IP addresses
- Verifies the operating system
- Validates the cluster networks and network connectivity
- Validates the network time protocol (NTP) setup
- Validates the physical disks in the storage servers
- Validates the compute servers for database readiness

Step 2 Setup Required Files

Ensures that the required files are accessible, copies them to the second compute server, unzips them as needed, and creates symbolic links.

See "[Software Installation Prerequisites](#) (page 11-2)" if this step fails.

Step 3 Create Users

Creates the required users and groups for the cluster, updates `/etc/hosts`, and sets up a secure shell (SSH).

Step 4 Setup Cell Connectivity

Generates several files that are used by other steps, including `cellip.ora`, `cellinit.ora`, and `cellaffinity`, if required.

Step 5 Verify InfiniBand

Performs a series of checks on the InfiniBand fabric to ensure that the compute servers, storage servers, and InfiniBand switches are cabled together correctly.

If errors occur, a diagnostic ZIP file is generated. The output from this step includes the location of this file.

Step 6 Calibrate Cells

Uses the I/O calibration feature of Oracle Database to assess the I/O capability of the storage subsystem.

Step 7 Create Cell Disks

Performs a series of steps on the storage server disks that include creating the cell disks, and then creating a flash log and a flash cache on each cell.

Step 8 Create Grid Disks

Creates the grid disks for the cluster.

In this release, you must follow this step with additional manual procedures.

Step 9 Configure Alerting

Configures email on the storage servers, so that they can send alerts when problems occur.

Step 10 Install Cluster Software

Uses an Oracle Database utility to obtain the grid disk identifiers, and then runs the clusterware installer. It patches the software as needed.

Step 11 Initialize Cluster Software

If a replication network is being configured, this step creates the single client access name (SCAN) and virtual IP (VIP) for it. Then it initializes the cluster.

Step 12 Install Database Software

Installs the Oracle Database software and creates the Oracle Database home directory on both compute servers. It also patches the software as required.

In this release, you must follow this step with additional manual procedures.

Step 13 Relink Database with RDS

Relinks Oracle Database with Reliable Datagram Sockets (RDS).

Step 14 Create ASM Disk Groups

Obtains the grid disk identifiers and creates the Automatic Storage Management (ASM) disk groups.

Step 15 Create Databases

Creates the Recovery Appliance catalog on both compute servers.

Step 16 Apply Security Fixes

Restarts the cluster software to set the security parameters, and enables turbo mode if required.

Step 17 Install Exachk

Installs the `exachk` utility on the Recovery Appliance. See My Oracle Support Doc ID 1070954.1 for information about the utility.

Step 18 Setup ASR Alerting

Installs and configures Auto Service Request (ASR). Optional.

This step does the following:

- Installs the required software packages
- Configures the trap destinations
- Starts the monitoring daemon

Step 19 Create Installation Summary

Obtains the system details, calibrates the storage servers, and generates a summary XML report. The output from this step includes the location of this file.

Step 20 Resecure Machine

Replaces the default passwords with custom passwords.

11.6 Setting Up Enterprise Manager to Administer Recovery Appliance

To enable Enterprise Manager to administer Recovery Appliance, you must deploy the management agents to each compute server, then discover the targets for the Recovery Appliance. The targets include the cluster, database, listeners, Automatic Storage Management (ASM), and the appliance itself.

Work with your Oracle support engineer to perform the tasks in this section.



Note:

See Doc ID 1929507.1 for information about the plug-ins and patches required for Zero Data Loss Recovery Appliance management, monitoring, and protected database management.

11.6.1 Deploying the Enterprise Manager Agents

Perform the following steps to deploy a management agent to each Recovery Appliance compute server:

1. On each compute server, create a directory owned by oracle:oinstall in which to install the agent.
2. Log in to Oracle Enterprise Manager Cloud Control (Cloud Control) as a user with privileges to discover and manage a Recovery Appliance target.
3. From the **Setup** menu, choose **Add Target**, select **Add Targets Manually**, choose **Add Host Targets**, then click **Add Host...**
The Add Host Targets: Host and Platform page appears.
4. From the Platform list, select **Same for All Hosts**.
5. For each compute server, click **Add**, enter the fully qualified host name, and select **Linux x86-64** in the Platform list. Then click **Next**.
The Add Host Targets: Installation Details page appears.
6. Specify the Agent installation directory that you created in Step 1, **Named Credential** for the user that will own the Agent installation, and other requested properties, then click **Next**.
The Add Host Targets: Review page appears.

7. Click **Deploy Agent** to begin deploying Agents to all Recovery Appliance compute servers.
The Add Host page appears, displaying warnings about remote prerequisite checks in the Agent Deployment Summary section. The Remote Prerequisite Check Details section indicates that the root.sh script could not be run due to insufficient privileges. You will fix this in a subsequent step.
8. To resume the deployment, click **Continue**, and select **All Hosts**.
The Add Host page displays a message that the deployment was successful.
9. Follow to the on screen instructions to run root.sh manually on each compute server.
10. Return to the Add Host page, and click **Done**.

11.6.2 Discovering the Cluster

Perform the following steps to discover the Recovery Appliance cluster:

1. From the **Setup** menu in Cloud Control, choose **Add Target**, select **Add Targets Manually**.
The Add Targets Manually page appears.
2. Select **Add Targets Using Guided Process**.
3. From the Target Types drop-down list, select **Oracle Cluster and High Availability Service**, then click **Add Using Guided Process...**
The Add Target: Cluster and Oracle High Availability Service page appears.
4. Enter the host name of one of the compute servers on the Recovery Appliance, then click the Search icon.
The Select Targets dialog appears.
5. Select the target host, and click **Select**.
6. Click **Discover Target**.
7. Review the automatically detected cluster information on the Cluster and Oracle High Availability Service page. Verify that all compute servers on the Recovery Appliance cluster are included in the list of cluster hosts and that the SCAN Name and SCAN Ports display the correct values for the ingest network, then click **Save**.
When the cluster target is created successfully, the Confirmation window appears.
8. Close the Confirmation window.

11.6.3 Discovering the Cluster Database Targets

Before you discover the Recovery Appliance itself, perform the following steps to discover the Recovery Appliance cluster database, listener, and ASM targets:

1. From the **Setup** menu in Cloud Control, choose **Add Target**, select **Add Targets Manually**.
The Add Targets Manually page appears.
2. Select **Add Targets Using Guided Process**.

3. From the Target Types drop-down list, choose **Oracle Database, Listener, and Automatic Storage Management**, then click **Add Using Guided Process...**
The Database Discovery: Search Criteria page appears.
4. Use the Recovery Appliance cluster name or the host name of one of the compute servers as the criteria to perform the search, then click **Next**.
The Database Discovery: Results page appears, showing the Recovery Appliance cluster database is listed, along with the ASM instances and listeners on all compute servers.
5. Select the cluster database, and click **Configure**.
6. Modify the settings to use fully qualified names in the Listener Machine Name fields and 1521 in the Port fields, then click **Save**.
7. Supply the monitoring credentials for the cluster database, and click **Test Connection**.
8. Select the cluster ASM, configure the appropriate monitoring credentials, and click **Test Connection**.
9. Ensure that the management database under Single Instance Databases (-MGMTDB) is not selected.
10. Select all listeners except the management listener (MGMTLSNR), and click **Next**.
The Database Discovery: Review page appears.
11. Verify the information, then click **Save** to start monitoring the targets.
12. Close the Confirmation window.

11.6.4 Discovering the Recovery Appliance

After discovering the Recovery Appliance cluster and cluster database targets, the Recovery Appliance target itself can be discovered. As part of discovering the top-level Recovery Appliance target that will include all Enterprise Manager Recovery Appliance software and hardware management functionality, a separate Recovery Appliance hardware target will also be discovered.

Perform the following steps to manually discover the Recovery Appliance target:

1. From the **Setup** menu in Cloud Control, choose **Add Target**, then select **Add Targets Manually**.
The Add Targets Manually page appears.
2. Select **Add Targets Using Guided Process**.
3. From the **Target Types** drop-down list, choose **Recovery Appliance**, then click **Add Using Guided Process...**
The Recovery Appliance Hardware Discovery page appears.
4. Select **Discover new Recovery Appliance hardware components as targets**, and click **Discover Targets**.
The Recovery Appliance Hardware Discovery: Discovery Inputs page appears.
5. Specify the URL of the agent on the compute server on which the schematic file (databasemachine.xml) resides. You can use the host name of the compute server on the Recovery Appliance to search for the correct value to use.

Information about the schematic file appears.

6. Set the credential for the host, and click **Next**.

The Recovery Appliance Hardware Discovery: Infiniband Discovery page appears.

7. Follow the detailed instructions in all subsequent pages of the Recovery Appliance Hardware Discovery wizard, supplying all requested credentials for hardware components. On the Review page, review all hardware component details, then click **Submit**. The Database Machine target representing the Recovery Appliance hardware will be created, along with targets for all hardware components.

The Target Creation Summary page appears, showing a summary of all hardware targets created, including the Database Machine target representing the Recovery Appliance hardware and all member targets for all hardware components.

8. Click **Continue With Recovery Appliance Discovery**.

The Recovery Appliance Discovery: Properties page appears.

9. Select the **Target Name** for the Recovery Appliance. The Recovery Appliance Hardware target name is already filled-in with the target name discovered in a previous step.

(If the Recovery Appliance hardware discovery was completed without completing the full Recovery Appliance discovery process, that may leave one or more Recovery Appliance hardware targets in a state where they are not associated with a Recovery Appliance target. In this case, when Recovery Appliance discovery is re-initiated from the Add Targets Manually page, the discovery process will not automatically go in to the Recovery Appliance hardware discovery wizard. Instead, the process will go directly to the Properties page to allow selection of an unassociated Recovery Appliance hardware target. In this case, the hardware target name will not be pre-filled. Click the **Select Target** icon to launch the Search and Select Targets popup, which shows a list of Recovery Appliance hardware targets that have not yet been associated with a Recovery Appliance target. Select the correct **Recovery Appliance hardware target**.)

10. In the Recovery Appliance Administrator Credentials section, specify the database user credentials for the Recovery Appliance recovery catalog owner. This user is also the Recovery Appliance administrator. You can use **Named** or **New** credentials. Click **More Details** to see detailed information about the credentials.
11. In the Recovery Appliance Monitoring Credentials section, specify the database user credentials that will be used to monitor the Recovery Appliance. If necessary, that user will be granted the role required to access monitoring information in the Recovery Appliance database. Alternatively, you can select **Use base catalog user credentials as monitoring credentials** which disables the Username and Password fields.

12. Click **Next**.

The Recovery Appliance Discovery: Oracle Secure Backup Domain page appears.

13. If Oracle Secure Backup is installed on the Recovery Appliance, enter `/usr/local/oracle/backup` in the Installation Home field and specify the monitoring credentials for the Oracle Secure Backup domain. Otherwise, select **Skip Oracle Secure Backup Domain Discovery**.

14. Click **Next**.

The Recovery Appliance Discovery: Review page appears.

15. Review the target discovery selections, and click **Submit**.

The Recovery Appliance target is created.

Part III

Maintenance of Recovery Appliance

Part III provides information about maintenance procedures for Recovery Appliance. It contains the following chapters:

- [Maintaining the Recovery Appliance Hardware](#) (page 12-1)
- [Reconfiguring the Hardware](#) (page 13-1)

12

Maintaining the Recovery Appliance Hardware

This chapter describes how to maintain the Recovery Appliance rack components. It contains the following topics:

- [Cautions and Warnings](#) (page 12-1)
- [Determining the Server Model](#) (page 12-2)
- [Powering On and Off a Recovery Appliance Rack](#) (page 12-2)
- [Replacing the Disk Controller Batteries](#) (page 12-7)
- [Replacing a Power Distribution Unit](#) (page 12-9)
- [Resetting a Non-Responsive Oracle ILOM](#) (page 12-11)
- [Maintaining the Compute Servers](#) (page 12-13)
- [Reimaging a Compute Server](#) (page 12-14)
- [Maintaining the Storage Servers](#) (page 12-24)
- [Maintaining the Physical Disks of Storage Servers](#) (page 12-26)
- [Maintaining the Flash Disks of Storage Servers](#) (page 12-37)
- [Replacing a Disk Controller Battery Backup Unit](#) (page 12-43)
- [Using the Storage Server Rescue Procedure](#) (page 12-45)



See Also:

[Replacement Units](#) (page D-1)

12.1 Cautions and Warnings

When maintaining the Recovery Appliance hardware, observe the following precautions:



WARNING:

Do not touch the parts of this product that use high-voltage power. Touching them might result in serious injury.

 **Caution:**

- Do not power off Recovery Appliance unless there is an emergency. In that case, follow "[Emergency Power-Off Procedure](#) (page 12-2)".
- Keep the front and rear cabinet doors closed. Failure to do so might cause system failure or result in damage to the hardware components.
- Keep the top, front, and back of the cabinets clear to allow proper airflow and to prevent the components from overheating.

12.2 Determining the Server Model

Use the following command to determine the model of a compute server or a storage server:

```
dmidecode -s system-product-name
```

12.3 Powering On and Off a Recovery Appliance Rack

This section includes the following topics:

- [Emergency Power-Off Procedure](#) (page 12-2)
- [Shutting Down Recovery Appliance](#) (page 12-3)
- [Starting Up Recovery Appliance](#) (page 12-5)

12.3.1 Emergency Power-Off Procedure

In an emergency, halt power to Recovery Appliance immediately. The following emergencies might require powering off Recovery Appliance:

- Natural disasters, such as earthquake, flood, hurricane, tornado, or cyclone
- Abnormal noise, smell, or smoke coming from the system
- Threat to human safety

12.3.1.1 Powering Off in an Emergency

In an emergency, do one of the following:

- Turn off power at the circuit breaker.
- Pull the emergency power-off switch in the computer room.

After the emergency, contact Oracle Support Services about restoring power to the system.

12.3.1.2 About the Emergency Power-Off Switch

You can use the emergency power-off (EPO) switch to remove power from Recovery Appliance.

EPO switches are required when computer equipment contains batteries capable of supplying more than 750 volt-amperes for more than five minutes. Systems that have these batteries include internal EPO hardware for connecting to a site EPO switch or relay.

12.3.2 Shutting Down Recovery Appliance

Under normal, nonemergency conditions, you can power down the software services and hardware gracefully.

Stop all software services before shutting down the rack components.

12.3.2.1 Stopping Recovery Appliance Services

You must stop the Recovery Appliance services, Oracle Database File System, Oracle Database, and the cluster services.

To stop the Recovery Appliance services:

1. Log in as `oracle` to either Recovery Appliance compute server.
2. Open a SQL connection to Oracle Database as the `rasys` user:

```
$ sqlplus rasys
```

3. Check the status of the services:

```
SQL> SELECT state FROM ra_server;

STATE
-----
ON
```

4. Shut down Recovery Appliance services:

```
SQL> exec dbms_ra.shutdown;
```

5. Disconnect from Oracle Database:

```
SQL> exit
```

6. If Oracle Secure Backup is configured:

- a. Switch to the `root` user.

- b. Check the current status of Oracle Secure Backup:

```
# $GRID_HOME/bin/crsctl status res osbadmin
NAME=osbadmin
TYPE=cluster_resource
TARGET=ONLINE
STATE=ONLINE on example01adm04
```

- c. If Oracle Secure Backup is online, then stop it:

```
# $GRID_HOME/bin/crsctl stop res osbadmin
```

- d. Switch back to the `oracle` user.

7. Check the status of Oracle Database:

```
$ srvctl status database -d zdlra5
Instance zdlra51 is running on node radb07
Instance zdlra52 is running on node radb08
```

8. Stop Oracle Database:

```
$ srvctl stop database -d zdlra5
```

9. Verify that Oracle Database is stopped:

```
$ srvctl status database -d zdlra5
Instance zdlra51 is not running on node radb07
Instance zdlra52 is not running on node radb08
```

10. Switch to the root user.

11. Stop the Oracle Clusterware stack on all nodes in the cluster:

```
# $GRID_HOME/bin/crsctl stop cluster -all
CRS-2673: Attempting to stop 'ora.crsd' on 'zdlradb07'
CRS-2790: Starting shutdown of Cluster Ready Services-managed resources on
'zdlradb07'
CRS-2673: Attempting to stop 'ora.LISTENER_SCAN2.lsnr' on 'zdlradb07'
CRS-2673: Attempting to stop 'ora.LISTENER_SCAN1.lsnr' on 'zdlradb07'
.
.
.
#
```

If the command fails, reenter it with the `-f` option.

12. On each compute server, run the following command to stop the Oracle Cluster Ready Services (CRS):

```
# $GRID_HOME/bin/crsctl stop crs
CRS-2791: Starting shutdown of Oracle High Availability Services-managed
resources on 'radb08'
CRS-2673: Attempting to stop 'ora.crf' on 'radb08'
CRS-2673: Attempting to stop 'ora.mdnsd' on 'radb08'
.
.
.
CRS-2677: Stop of 'ora.crf' on 'radb08' succeeded
CRS-2677: Stop of 'ora.mdnsd' on 'radb08' succeeded
CRS-2793: Shutdown of Oracle High Availability Services-managed resources on
'radb08' has completed
CRS-4133: Oracle High Availability Services has been stopped.
```

13. Shut down or reboot the hardware as required, in the following order:

- a. Compute servers
- b. Storage servers
- c. Rack and switches

12.3.2.2 Powering Down the Servers

Before powering down a server, stop the services running on it, as described in "[Shutting Down Recovery Appliance](#) (page 12-3)".

To shut down a compute server or a storage server:

1. Log in to the server as `root`.

2. Stop the operating system:

```
# shutdown -h -y now
```

Or restart the operating system:

```
# shutdown -r -y now
```

Example 12-1 Powering Off Recovery Appliance Using the dcli Utility

1. Stop Oracle Clusterware on all compute servers:

```
# GRID_HOME/grid/bin/crsctl stop cluster -all
```

2. Shut down the other compute server in the rack:

```
# dcli -l root -g ra-adm02 shutdown -h -y now
```

In the preceding command, `ra01adm02` is the name of the second compute server.

3. Shut down all storage servers:

```
# dcli -l root -g cell_group shutdown -h -y now
```

In the preceding command, `cell_group` is a file that lists all storage servers.

4. Shut down the local compute server:

```
shutdown -h -y now
```

5. Power off the rack.

Use the `dcli` utility to run the `shutdown` command on multiple servers simultaneously. Do not run `dcli` from a server that will be powered off by the command.

The following example shuts down a group of storage servers listed in a file named `cell_group`:

```
# dcli -l root -g cell_group shutdown -h -y now
```

[Example 12-1](#) (page 12-5) shows the power off procedure for the rack when using the `dcli` utility to shut down multiple servers simultaneously. The commands run from a compute server.

12.3.2.3 Powering the Network Switches

The gateway and spine switches do not have power controls. They power off when power is removed, by turning off a PDU or a breaker in the data center.

12.3.3 Starting Up Recovery Appliance

Turn on the rack components first, then start the software services.

12.3.3.1 Starting Up Recovery Appliance Components

To power on the rack components, use one of the following methods:

- Press the power button on the front of the component.
- Log in to Oracle ILOM and apply power to the system. See "[Powering On Servers Remotely](#) (page 12-6)".

12.3.3.2 Startup Sequence

Power on the rack components in this sequence:

1. Rack and switches
Allow the switches a few minutes to initialize, before you start the storage servers.
2. Storage servers
Allow five to 10 minutes for the storage servers to start all services. Ensure that they finish initializing before you start the compute servers.
3. Compute servers
When a compute server is powered on, the operating system and Oracle Clusterware start automatically. Oracle Clusterware then starts all resources that are configured to start automatically.

12.3.3.3 Powering On Servers Remotely

You can use the Oracle ILOM interface to power on the Recovery Appliance servers remotely. To access Oracle ILOM, use the web console, the command-line interface (CLI), intelligent platform management interface (IPMI), or simple network management protocol (SNMP).

For example, to apply power to server `ra01ce101` using IPMI, you use its Oracle ILOM with a command like the following:

```
# ipmitool -H ra01ce101-ilom -U root chassis power on
```

IPMITool must be installed on the server where you use the command.



See Also:

Oracle Integrated Lights Out Manager (ILOM) 3.0 documentation for additional information about using Oracle ILOM to power on the servers:

<http://docs.oracle.com/cd/E19860-01/index.html>

12.3.3.4 Starting the Recovery Appliance Software

1. Log in as `root` to a Recovery Appliance compute server.
2. Confirm that Oracle Cluster Ready Services (CRS) is running:

```
# $GRID_HOME/bin/crsctl status server
NAME=radb07
STATE=ONLINE

NAME=radb08
STATE=ONLINE
```

3. If CRS is not running, then start it:

```
# $GRID_HOME/bin/crsctl start cluster -all
CRS-2672: Attempting to start 'ora.evmd' on 'radb07'
```

```
CRS-2672: Attempting to start 'ora.cssdmonitor' on 'radb07'  
CRS-2672: Attempting to start 'ora.cssdmonitor' on 'radb08'  
.  
.  
.  
#
```

4. Switch to the `oracle` user.

5. Verify that Oracle Database is running:

```
$ srvctl status database -d zdlra5  
Instance zdlra51 is not running on node radb07  
Instance zdlra52 is not running on node radb08
```

6. If Oracle Database is not running:

a. Start Oracle Database:

```
$ srvctl start database -d zdlra5
```

b. Confirm that Oracle Database is running:

```
$ srvctl status database -d zdlra5  
Instance zdlra51 is running on node radb07  
Instance zdlra52 is running on node radb08
```

7. If Oracle Secure Backup is enabled, then start it:

```
# $GRID_HOME/bin/crsctl start res osbadm
```

8. Connect to Oracle Database as the `RASYS` user:

```
$ sqlplus rasys
```

9. Check the status of Recovery Appliance services:

```
SQL> SELECT state FROM ra_server;
```

```
STATE  
-----  
OFF
```

10. If the services are off, then start them:

```
SQL> exec dbms_ra.startup;
```

11. Confirm that the services are started:

```
SQL> /  
  
STATE  
-----  
ON
```

12.4 Replacing the Disk Controller Batteries

The disk controllers in storage servers and compute servers have battery-backed write cache to accelerate write performance. If the battery charge capacity degrades, so that the battery cannot protect the cached data for a power loss of 48 hours or more, then the write cache is disabled and the disk controller switches to write-through mode. Write performance is reduced, but no data is lost.

The battery charge capacity degrades over time, and its life expectancy is inversely proportional to the operating temperature. [Table 12-1](#) (page 12-8) shows the worst case life expectancy of a battery in Recovery Appliance.

Table 12-1 Battery Life Expectancy

Inlet Ambient Temperature	Battery Lifetime
< 25 degrees Celsius (77 degrees Fahrenheit)	3 years
< 32 degrees Celsius (89.6 digresses Fahrenheit)	2 years

12.4.1 Monitoring Batteries in the Compute Servers

To monitor the battery charge capacity in the compute servers:

```
# /opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -a0 | grep "Full Charge" -A5 | sort \
| grep Full -A1
```

The following is an example of the output from the command:

```
Full Charge Capacity: 1357 mAh
Max Error: 2 %
```

You should proactively replace batteries that have a capacity less than 800 milliamper hour (mAh) and a maximum error less than 10%. Immediately replace any battery that has less than 674 mAh or a maximum error greater than 10%.

To monitor the battery temperature:

```
/opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -a0 | grep BatteryType; \
/opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -a0 | grep -i temper
```

The following is an example of the output from the command:

```
BatteryType: iBBU08
Temperature: 38 C
  Temperature           : OK
  Over Temperature      : No
```

If the battery temperature is greater than or equal to 55 degrees Celsius, then determine the cause, and correct the problem.



Note:

Storage servers generate an alert when the battery charge capacity is insufficient, the temperature is high, or the battery should be replaced.

12.4.2 Replacing Batteries in Disk Controllers

Oracle replaces the failed batteries at no extra charge under these conditions:

- The battery charge capacity in the disk controllers falls below the minimum threshold

- The system is covered either by the Oracle Premier Support for Systems or occurs during the warranty period.

For customers with Premier Support for Systems, Oracle attempts to proactively replace the batteries in Recovery Appliance before the end of the estimated lifetime, on a best effort basis.

12.5 Replacing a Power Distribution Unit

A power distribution unit (PDU) can be replaced while Recovery Appliance is online. The second PDU in the rack maintains the power to all components in the rack. PDU-A is on the left, and PDU-B is on the right, when viewing the rack from the rear.

12.5.1 PDU Replacement Guidelines

Before replacing a PDU, review the following guidelines to ensure that you can perform the procedure safely and without disrupting availability:

- Unlatching the InfiniBand cables while removing or inserting PDU-A might remove servers from the cluster and thus make the rack unavailable. Be careful when handling the InfiniBand cables, which are normally latched securely. Do not place excessive tension on the InfiniBand cables by pulling them.
- Unhooking the wrong power feeds shuts down the rack. Trace the power cables that will be replaced from the PDU to the power source, and only unplug those feeds.
- Allow time to unpack and repack the PDU replacement parts. Notice how the power cords are coiled in the packaging, so you can repack the failed unit the same way.
- Removing the side panel decreases the time needed to replace the PDU. However, removing the side panel is optional.
- Using a cordless drill or power screwdriver decreases the time needed to replace the PDU. Allow more time for the replacement if you use a hand wrench. A screwdriver requires Torx T30 and T25 bits.
- You might need to remove the server cable arms to move the power cables. In that case, twist the plug connection and flex the cable arm connector, to avoid having to unclip the cable arm. If you must unclip the cable arm, then support the cables with one hand, remove the power cord, and then clip the cable arm. Do not leave the cable arm hanging.
- When you remove the T30 screws from the L-bracket, do not remove the T25 screws or nuts that attach the PDU to the bracket, until the PDU is out of the rack.

12.5.2 Replacing a PDU

To replace a PDU:

1. Restart the PDU monitor to identify the network settings:
 - a. Press the reset button for 20 seconds, until it starts to count down from 5 to 0. While it is counting, release the button, and then press it once.
 - b. Record the network settings, firmware version, and so on, displayed on the LCD screen as the monitor restarts.

If the PDU monitor is not working, then retrieve the network settings by connecting to the PDU over the network, or from the network administrator.

2. Turn off all PDU breakers.
3. Unplug the PDU power plugs from the AC outlets.

If the rack is on a raised floor, then move the power cords out through the floor cutout. You might need to maneuver the rack over the cutout first.

 **WARNING:**

If the power cords use overhead routing, then put them in a location where they will not fall or hit anyone.

4. For replacing PDU-B when there is no side panel access, and the rack does not have an InfiniBand cable harness:

 **Note:**

Do not unstrap any cables attached to the cable arms.

- a. Unscrew the T25 screws holding the square cable arms to the rack.
- b. Move the InfiniBand cables to the middle, out of the way.
5. Unplug all power cables that connect the servers and switches to the PDU. Keep the power cables together in group bundles.
6. Remove the T30 screws from the top and bottom of the L-bracket, and note where the screws are used.
7. Note where the PDU sits in the rack frame.
The PDU is typically an inch back from the rack frame, to allow access to the breaker switches.
8. Angle and maneuver the PDU out of the rack.
9. Hold the PDU or put it down, if there is enough room, while maneuvering the AC power cords through the rack. You might need to cut the cable ties that hold the AC cord flush with the bottom side of the PDU.
10. Pull the cords as near to the bottom or top of the rack as possible. There is more room between the servers to guide the outlet plug through the routing hole.
11. Remove the smaller Torx T25 screws, and loosen the nut on the top and bottom to remove the PDU from the L-bracket. You do not need to remove the nut.
12. Attach the L-bracket to the new PDU.
13. Put the new PDU next to the rack.
14. Route the AC cords through the rack to the outlets.

 **Note:**

Do not cable-tie the AC cord to the new PDU.

15. Place the new PDU in the rack by angling and maneuvering it until the L-brackets rest on the top and bottom rails.
16. Line up the holes and slots so that the PDU is about an inch back from the rack frame.
17. Attach the power cords, using the labels on the cords as a guide. For example, G5-0 indicates PDU group 5 outlet 0 on the PDU.
18. Attach the InfiniBand cable holders, if you removed them in step 4 (page 12-10). Oracle recommends that you first screw in the holders by hand to avoid stripping the screws.
19. Attach the AC power cords to the outlets.
20. Turn on the breakers.
21. Cable and program the PDU monitor for the network, as needed.

 **See Also:**

Oracle Sun Rack II Power Distribution Units User's Guide for information about programming the PDU monitor at

<http://docs.oracle.com/cd/E19844-01/html/E23956/index.html>

12.6 Resetting a Non-Responsive Oracle ILOM

The Oracle Integrated Lights Out Manager (Oracle ILOM) might become unresponsive. If this happens, then you must manually reset the Service Processor (SP) on Oracle ILOM.

The following procedures describe how to reset Oracle ILOM:

- [Resetting Oracle ILOM Using SSH](#) (page 12-12)
- [Resetting Oracle ILOM Using the Remote Console](#) (page 12-12)
- [Resetting Oracle ILOM Using IPMITool](#) (page 12-12)
- [Resetting Oracle ILOM By Removing Power](#) (page 12-12)

 **See Also:**

Oracle Integrated Lights Out Manager (ILOM) 3.0 documentation at

<http://docs.oracle.com/cd/E19860-01/E21549/bbgiedhj.html#z4000b491400243>

12.6.1 Resetting Oracle ILOM Using SSH

To reset Oracle ILOM using SSH:

1. Connect to Oracle ILOM using SSH from another system.
2. Enter the following command at the ILOM prompt:

```
reset /SP
```

12.6.2 Resetting Oracle ILOM Using the Remote Console

If you cannot connect to Oracle ILOM using SSH, then log in to the remote console.

To reset Oracle ILOM using the remote console:

1. Log in to the Oracle ILOM remote console.
2. Select **Reset SP** from the Maintenance tab.
3. Click **Reset SP**.

12.6.3 Resetting Oracle ILOM Using IPMItool

If you cannot connect to Oracle ILOM using either SSH or the remote console, then use IPMItool.

To reset Oracle ILOM using IPMItool:

1. Log in to the local host or another host on the management network.
2. Use the following IPMItool command:

- On the local host:

```
$ ipmitool mc reset cold  
Sent cold reset command to MC
```

- On another host:

```
$ ipmitool -H ILOM_host_name -U ILOM_user mc reset cold  
Sent cold reset command to MC
```

In the preceding command, *ILOM_host_name* is the host name being used, and *ILOM_user* is the user name for Oracle ILOM.

12.6.4 Resetting Oracle ILOM By Removing Power

If you cannot reset Oracle ILOM using the preceding options:

1. Unplug the server from the power supply.
2. Plug the server back into the power supply.

This action power cycles the server and Oracle ILOM.

12.7 Maintaining the Compute Servers

You do not need to shut down a compute server in Recovery Appliance to repair the physical disks. No downtime of the rack is required; however, individual servers might require downtime, and you might need to take them out of the cluster temporarily.

An LSI MegaRAID SAS 9261-8i disk controller manages the disk drives in each compute server. The disks have a RAID-5 configuration. Each compute server has four disk drives. One virtual drive comprises the RAID set.

See Also:

- "[LED Status Descriptions](#) (page D-2)" for information about the LEDs
- "[Parts for Compute Servers](#) (page D-6)" for the repair procedures

12.7.1 Verifying the RAID Status of a Compute Server

Oracle recommends that you periodically verify the status of the compute server RAID devices. The impact is minimal. In contrast, the impact of corrective action varies depending on the specific issue uncovered, and can range from simple reconfiguration to an outage.

Log in to each compute server as `root` and perform the following procedure.

To verify the RAID status:

1. Check the current disk controller configuration:

```
# /opt/MegaRAID/MegaCli/MegaCli64 -AdpAllInfo -aALL | grep "Device Present" -A 8
```

```

                Device Present
                =====
Virtual Drives   : 1
  Degraded       : 0
  Offline        : 0
Physical Devices : 5
  Disks          : 4
  Critical Disks : 0
  Failed Disks   : 0

```

Verify that the output shows one virtual drive, none degraded or offline, five physical devices (one controller + four disks), four disks, and no critical or failed disks.

If the output is different, then investigate and correct the problem. Degraded virtual drives usually indicate absent or failed physical disks. Replace critical disks and failed disks immediately. Otherwise, you risk data loss if the number of working disks in the server is less than the number required to sustain normal operation.

2. Check the current virtual drive configuration:

```
# /opt/MegaRAID/MegaCli/MegaCli64 CfgDsply -aALL | grep "Virtual Drive:"; \
/opt/MegaRAID/MegaCli/MegaCli64 CfgDsply -aALL | grep "Number Of Drives"; \
/opt/MegaRAID/MegaCli/MegaCli64 CfgDsply -aALL | grep "^State"
```

```
Virtual Drive          : 0 (Target Id: 0)
Number Of Drives      : 4
State                  : Optimal
```

Verify that virtual device 0 has four drives, and the state is Optimal. If the output is different, then investigate and correct the problem.

3. Check the current physical drive configuration:

```
# /opt/MegaRAID/MegaCli/MegaCli64 -PDList -aALL | grep "Firmware state"
Firmware state: Online, Spun Up
```

Ensure that all drives are Online, Spun Up. If the output is different, then investigate and correct the problem.

If the output is different, then investigate and correct the problem. Degraded virtual drives usually indicate absent or failed physical disks. Replace critical disks and failed disks immediately. Otherwise, you risk data loss if the number of working disks in the server is less than the number required to sustain normal operation.

12.8 Reimaging a Compute Server

If a compute server is irretrievably damaged, then you must replace it and reimage the replacement server. During the reimaging procedure, the other compute servers in the cluster are available. When adding the new server to the cluster, you copy the software from a working compute server to the new server.

The following tasks describe how to reimage a compute server:

- ["Contacting Oracle Support Services \(page 12-14\)"](#)
- ["Downloading the Latest Release of the Cluster Verification Utility \(page 12-15\)"](#)
- ["Removing the Failed Compute Server from the Cluster \(page 12-15\)"](#)
- ["Preparing the USB Flash Drive for Imaging \(page 12-16\)"](#)
- ["Copying the Image to the New Compute Server \(page 12-17\)"](#)
- ["Configuring the Replacement Compute Server \(page 12-17\)"](#)
- ["Preparing the Replacement Compute Server for the Cluster \(page 12-18\)"](#)
- ["Applying Patch Bundles to a Replacement Compute Server \(page 12-20\)"](#)
- ["Cloning the Oracle Grid Infrastructure \(page 12-20\)"](#)
- ["Clone Oracle Database Homes to the Replacement Compute Server \(page 12-22\)"](#)

12.8.1 Contacting Oracle Support Services

Open a support request with Oracle Support Services. The support engineer identifies the failed server and sends you a replacement. The support engineer also asks for the output from the `imagehistory` command, run from a working compute server. The output provides a link to the `computeImageMaker` file that was used to image the original compute server, and is used to restore the system.

12.8.2 Downloading the Latest Release of the Cluster Verification Utility

The latest release of the cluster verification utility (`cluvfy`) is available from My Oracle Support Doc ID 316817.1.

12.8.3 Removing the Failed Compute Server from the Cluster

You must remove the failed compute server from Oracle Real Application Clusters (Oracle RAC).

In these steps, *working_server* is a working compute server in the cluster, *failed_server* is the compute server being replaced, and *replacement_server* is the new server.

To remove a failed compute server from the Oracle RAC cluster:

1. Log in to *working_server* as the `oracle` user.

2. Disable the listener that runs on the failed server:

```
$ srvctl disable listener -n failed_server
$ srvctl stop listener -n failed_server
```

3. Delete the Oracle home directory from the inventory:

```
$ cd $ORACLE_HOME/oui/bin
$ ./runInstaller -updateNodeList ORACLE_HOME= \
/u01/app/oracle/product/12.1.0/dbhome_1 "CLUSTER_NODES=list_of_working_servers"
```

In the preceding command, *list_of_working_servers* is a list of the compute servers that are still working in the cluster, such as `ra01db02`, `ra01db03`, and so on.

4. Verify that the failed server was deleted—that is, unpinned—from the cluster:

```
$ olsnodes -s -t
```

```
ra01db01      Inactive      Unpinned
ra01db02      Active        Unpinned
```

5. Stop and delete the virtual IP (VIP) resources for the failed compute server:

```
# srvctl stop vip -i failed_server-vip
PRCC-1016 : failed_server-vip.example.com was already stopped

# srvctl remove vip -i failed_server-vip
Please confirm that you intend to remove the VIPs failed_server-vip (y/[n]) y
```

6. Delete the compute server from the cluster:

```
# crsctl delete node -n failed_server
CRS-4661: Node failed_server successfully deleted.
```

If you receive an error message similar to the following, then relocate the voting disks.

```
CRS-4662: Error while trying to delete node ra01db01.
CRS-4000: Command Delete failed, or completed with errors.
```

To relocate the voting disks:

- a. Determine the current location of the voting disks. The sample output shows that the current location is DBFS_DG.

```
# crsctl query css votedisk

## STATE      File Universal Id      File Name               Disk group
--  -
1. ONLINE    123456789abab (o/192.168.73.102/DATA_CD_00_ra01cel07) [DBFS_DG]
2. ONLINE    123456789cdcd (o/192.168.73.103/DATA_CD_00_ra01cel08) [DBFS_DG]
3. ONLINE    123456789efef (o/192.168.73.100/DATA_CD_00_ra01cel05) [DBFS_DG]
Located 3 voting disk(s).
```

- b. Move the voting disks to another disk group:

```
# ./crsctl replace votedisk +DATA

Successful addition of voting disk 2345667aabdd.
...
CRS-4266: Voting file(s) successfully replaced
```

- c. Return the voting disks to the original location. This example returns them to DBFS_DG:

```
# ./crsctl replace votedisk +DBFS_DG
```

- d. Repeat the `crsctl` command to delete the server from the cluster.

7. Update the Oracle inventory:

```
$ cd $ORACLE_HOME/oui/bin
$ ./runInstaller -updateNodeList ORACLE_HOME=/u01/app/12.1.0/grid \
  "CLUSTER_NODES=list_of_working_servers" CRS=TRUE
```

8. Verify that the server was deleted successfully:

```
$ cluvfy stage -post nodedel -n failed_server -verbose

Performing post-checks for node removal
Checking CRS integrity...
The Oracle clusterware is healthy on node "ra0ldb02"
CRS integrity check passed
Result:
Node removal check passed
Post-check for node removal was successful.
```

See Also:

Oracle Real Application Clusters Administration and Deployment Guide for information about deleting a compute server from a cluster

12.8.4 Preparing the USB Flash Drive for Imaging

Use a USB flash drive to copy the image to the new compute server.

To prepare the USB flash drive for use:

1. Insert a blank USB flash drive into a working compute server in the cluster.
2. Log in as the `root` user.
3. Unzip the `computeImage` file:

```
# unzip computeImageMaker_release_LINUX.X64_release_date.platform.tar.zip
# tar -xvf computeImageMaker_release_LINUX.X64_release_date.platform.tar
```

4. Load the image onto the USB flash drive:

```
# cd dl360
# ./makeImageMedia.sh -dualboot no
```

The `makeImageMedia.sh` script prompts for information.

5. Remove the USB flash drive from the compute server.

6. Remove the unzipped `dl360` directory and the `computeImageMaker` file from the working compute server. The directory and file require about 2 GB of disk space.

12.8.5 Copying the Image to the New Compute Server

Before you perform the following procedure, replace the failed compute server with the new server. See [Expanding a Recovery Appliance Rack with Additional Storage Servers](#) (page 7-1).

To load the image onto the replacement server:

1. Insert the USB flash drive into the USB port on the replacement server.
2. Log in to the console through the service processor to monitor progress.
3. Power on the compute server either by physically pressing the power button or by using Oracle ILOM.
4. If you replaced the motherboard:
 - a. Press F2 during BIOS
 - b. Select **BIOS Setup**
 - c. Set the USB flash drive first, and then the RAID controller.

Otherwise, press F8 during BIOS, select the one-time boot selection menu, and choose the USB flash drive.

5. Allow the system to start.

As the system starts, it detects the CELLUSBINSTALL media. The imaging process has two phases. Let both phases complete before proceeding to the next step.

The first phase of the imaging process identifies any BIOS or firmware that is out of date, and upgrades the components to the expected level for the image. If any components are upgraded or downgraded, then the system automatically restarts.

The second phase of the imaging process installs the factory image on the replacement compute server.

6. Remove the USB flash drive when the system prompts you.
7. Press Enter to power off the server.

12.8.6 Configuring the Replacement Compute Server

The replacement compute server does not have a host names, IP addresses, DNS, or NTP settings. This task describes how to configure the replacement compute server.

The information must be the same on all compute servers in Recovery Appliance. You can obtain the IP addresses from the DNS. You should also have a copy of the Installation Template from the initial installation.

To configure the replacement compute server:

1. Assemble the following information:
 - Name servers
 - Time zone, such as Americas/Chicago
 - NTP servers
 - IP address information for the management network
 - IP address information for the client access network
 - IP address information for the InfiniBand network
 - Canonical host name
 - Default gateway
2. Power on the replacement compute server. When the system starts, it automatically runs the configuration script and prompts for information.
3. Enter the information when prompted, and confirm the settings. The startup process then continues.

 **Note:**

- If the compute server does not use all network interfaces, then the configuration process stops with a warning that some network interfaces are disconnected. It prompts whether to retry the discovery process. Respond with `yes` or `no`, as appropriate for the environment.
- If bonding is used for the ingest network, then it is now set in the default active-passive mode.

12.8.7 Preparing the Replacement Compute Server for the Cluster

The initial installation of Recovery Appliance modified various files.

To modify the files on the replacement compute server:

1. Replicate the contents of the following files from a working compute server in the cluster:
 - a. Copy the `/etc/security/limits.conf` file.
 - b. Merge the contents of the `/etc/hosts` files.
 - c. Copy the `/etc/oracle/cell/network-config/cellinit.ora` file.
 - d. Update the IP address with the IP address of the BONDIB0 interface on the replacement compute server.
 - e. Copy the `/etc/oracle/cell/network-config/cellip.ora` file.
 - f. Configure additional network requirements, such as 10 GbE.

- g. Copy the `/etc/modprobe.conf` file.
 - h. Copy the `/etc/sysctl.conf` file.
 - i. Restart the compute server, so the network changes take effect.
2. Set up the Oracle software owner on the replacement compute server by adding the user name to one or more groups. The owner is usually the `oracle` user.
 - a. Obtain the current group information from a working compute server:


```
# id oracle
uid=1000(oracle) gid=1001(oinstall) groups=1001(oinstall),1002(dba),
1003(oper),1004(asmdba)
```
 - b. Use the `groupadd` command to add the group information to the replacement compute server. This example adds the groups identified in the previous step:


```
# groupadd -g 1001 oinstall
# groupadd -g 1002 dba
# groupadd -g 1003 oper
# groupadd -g 1004 asmdba
```
 - c. Obtain the current user information from a working compute server:


```
# id oracle uid=1000(oracle) gid=1001(oinstall) \
groups=1001(oinstall),1002(dba),1003(oper),1004(asmdba)
```
 - d. Add the user information to the replacement compute server. This example adds the group IDs from the previous step to the `oracle` user ID:


```
# useradd -u 1000 -g 1001 -G 1001,1002,1003,1004 -m -d /home/oracle -s \
/bin/bash oracle
```
 - e. Create the `ORACLE_BASE` and Grid Infrastructure directories. This example creates `/u01/app/oracle` and `/u01/app/12.1.0/grid`:


```
# mkdir -p /u01/app/oracle
# mkdir -p /u01/app/12.1.0/grid
# chown -R oracle:oinstall /u01/app
```
 - f. Change the ownership of the `cellip.ora` and `cellinit.ora` files. The owner is typically `oracle:dba`.


```
# chown -R oracle:dba /etc/oracle/cell/network-config
```
 - g. Secure the restored compute server:


```
$ chmod u+x /opt/oracle.SupportTools/harden_passwords_reset_root_ssh
$ /opt/oracle.SupportTools/harden_passwords_reset_root_ssh
```

The compute server restarts.
 - h. Log in as the `root` user. When you are prompted for a new password, set it to match the `root` password of the other compute servers.
 - i. Set the password for the Oracle software owner. The owner is typically `oracle`.


```
# passwd oracle
```
 3. Set up SSH for the `oracle` account:
 - a. Change to the `oracle` account on the replacement compute server:


```
# su - oracle
```
 - b. Create the `dcli` group file on the replacement compute server, listing the servers in the Oracle cluster.

- c. Run the `setssh-Linux.sh` script on the replacement compute server. This example runs the script interactively:

```
$ /opt/oracle.SupportTools/onecommand/setssh-Linux.sh -s
```

The script prompts for the `oracle` password on the servers. The `-s` option causes the script to run in silent mode.

- d. Change to the `oracle` user on the replacement compute server:

```
# su - oracle
```

- e. Verify SSH equivalency:

```
$ dcli -g dbs_group -l oracle date
```

4. Set up or copy any custom login scripts from the working compute server to the replacement compute server:

```
$ scp .bash* oracle@replacement_server:.
```

In the preceding command, `replacement_server` is the name of the new server, such as `ra01db01`.

12.8.8 Applying Patch Bundles to a Replacement Compute Server

Oracle periodically releases software patch bundles for Recovery Appliance. If the working compute server has a patch bundle that is later than the release of the `computeImageMaker` file, then you must apply the patch bundle to the replacement compute server.

To determine if a patch bundle was applied, use the `imagehistory` command. Compare information on the replacement compute server to information on the working compute server. If the working database has a later release, then apply the storage server patch bundle to the replacement compute server.

12.8.9 Cloning the Oracle Grid Infrastructure

The following procedure describes how to clone the Oracle Grid infrastructure onto the replacement compute server. In the commands, `working_server` is a working compute server, and `replacement_server` is the replacement compute server.

To clone the Oracle Grid infrastructure:

1. Log in as `root` to a working compute server in the cluster.
2. Verify the hardware and operating system installation using the cluster verification utility (`cluvfy`):

```
$ cluvfy stage -post hwos -n replacement_server,working_server -verbose
```

The phrase `Post-check for hardware and operating system setup was successful` should appear at the end of the report.

3. Verify peer compatibility:

```
$ cluvfy comp peer -refnode working_server -n replacement_server \
  -orainv oinstall -osdba dba | grep -B 3 -A 2 mismatched
```

The following is an example of the output:

```

Compatibility check: Available memory [reference node: ra01db02]
Node Name Status Ref. node status Comment
-----
ra01db01 31.02GB (3.2527572E7KB) 29.26GB (3.0681252E7KB) mismatched
Available memory check failed
Compatibility check: Free disk space for "/tmp" [reference node: ra01db02]
Node Name Status Ref. node status Comment
-----
ra01db01 55.52GB (5.8217472E7KB) 51.82GB (5.4340608E7KB) mismatched
Free disk space check failed

```

If the only failed components are related to the physical memory, swap space, and disk space, then it is safe for you to continue.

4. Perform the requisite checks for adding the server:
 - a. Ensure that the `GRID_HOME/network/admin/samples` directory has permissions set to 750.
 - b. Validate the addition of the compute server:

```

$ cluvfy stage -ignorePrereq -pre nodeadd -n replacement_server \
  -fixup -fixupdir /home/oracle/fixup.d

```

If the only failed component is related to swap space, then it is safe for you to continue.

You might get an error about a voting disk similar to the following:

```

ERROR:
PRVF-5449 : Check of Voting Disk location "o/192.168.73.102/ \
DATA_CD_00_ra01cel07(o/192.168.73.102/DATA_CD_00_ra01cel07)" \
failed on the following nodes:
Check failed on nodes:
    ra01db01
    ra01db01:No such file or directory
...
PRVF-5431 : Oracle Cluster Voting Disk configuration check failed

```

If this error occurs, then use the `-ignorePrereq` option when running the `addnode` script in the next step.

5. Add the replacement compute server to the cluster:

```

$ cd /u01/app/12.1.0/grid/addnode/
$ ./addnode.sh -silent "CLUSTER_NEW_NODES={replacement_server}" \
  "CLUSTER_NEW_VIRTUAL_HOSTNAMES={replacement_server-vip}" [-ignorePrereq]

```

The `addnode` script causes Oracle Universal Installer to copy the Oracle Clusterware software to the replacement compute server. A message like the following is displayed:

```

WARNING: A new inventory has been created on one or more nodes in this session.
However, it has not yet been registered as the central inventory of this
system. To register the new inventory please run the script at
'/u01/app/oraInventory/orainstRoot.sh' with root privileges on nodes
'ra01db01'. If you do not register the inventory, you may not be able to
update or patch the products you installed.

```

The following configuration scripts need to be executed as the "root" user in each cluster node:

```
/u01/app/oraInventory/orainstRoot.sh #On nodes ra01db01

/u01/app/12.1.0/grid/root.sh #On nodes ra01db01
```

6. Run the configuration scripts:
 - a. Open a terminal window.
 - b. Log in as the `root` user.
 - c. Run the scripts on each cluster server.

After the scripts are run, the following message is displayed:

```
The Cluster Node Addition of /u01/app/12.1.0/grid was successful.
Please check '/tmp/silentInstall.log' for more details.
```

7. Run the `orainstRoot.sh` and `root.sh` scripts:

```
# /u01/app/oraInventory/orainstRoot.sh
Creating the Oracle inventory pointer file (/etc/oraInst.loc)
Changing permissions of /u01/app/oraInventory.
Adding read,write permissions for group.
Removing read,write,execute permissions for world.
Changing groupname of /u01/app/oraInventory to oinstall.
The execution of the script is complete.

# /u01/app/12.1.0/grid/root.sh
```

Check the log files in `/u01/app/12.1.0/grid/install/` for the output of the `root.sh` script. The output file reports that the listener resource on the replaced compute server failed to start. This is an example of the expected output:

```
/u01/app/12.1.0/grid/bin/srvctl start listener -n ra01db01 \
...Failed
/u01/app/12.1.0/grid/perl/bin/perl \
-I/u01/app/12.1.0/grid/perl/lib \
-I/u01/app/12.1.0/grid/crs/install \
/u01/app/12.1.0/grid/crs/install/rootcrs.pl execution failed
```

8. Reenable the listener resource that you stopped in "[Removing the Failed Compute Server from the Cluster](#) (page 12-15)".

```
# GRID_HOME/grid/bin/srvctl enable listener -l LISTENER \
-n replacement_server

# GRID_HOME/grid/bin/srvctl start listener -l LISTENER \
-n replacement_server
```

See Also:

Oracle Real Application Clusters Administration and Deployment Guide for information about cloning

12.8.10 Clone Oracle Database Homes to the Replacement Compute Server

To clone the Oracle Database homes to the replacement server:

1. Add Oracle Database ORACLE_HOME to the replacement compute server:

```
$ cd /u01/app/oracle/product/12.1.0/db_home/addnode/
$ ./addnode.sh -silent "CLUSTER_NEW_NODES={replacement_server}" -ignorePrereq
```

The `addnode` script causes Oracle Universal Installer to copy the Oracle Database software to the replacement compute server.

WARNING: The following configuration scripts need to be executed as the "root" user in each cluster node.

```
/u01/app/oracle/product/12.1.0/dbhome_1/root.sh #On nodes ra01db01
To execute the configuration scripts:
Open a terminal window.
Log in as root.
Run the scripts on each cluster node.
```

After the scripts are finished, the following messages appear:

```
The Cluster Node Addition of /u01/app/oracle/product/12.1.0/dbhome_1 was
successful.
Please check '/tmp/silentInstall.log' for more details.
```

2. Run the `root.sh` script on the replacement compute server:

```
# /u01/app/oracle/product/12.1.0/dbhome_1/root.sh
```

Check the `/u01/app/oracle/product/12.1.0/dbhome_1/install/root_replacement_server.company.com_date.log` file for the output of the script.

3. Ensure that the instance parameters are set for the replaced database instance. The following is an example for the `CLUSTER_INTERCONNECTS` parameter.

```
SQL> SHOW PARAMETER cluster_interconnects
```

NAME	TYPE	VALUE
-----	-----	-----
cluster_interconnects	string	

```
SQL> ALTER SYSTEM SET cluster_interconnects='192.168.73.90' SCOPE=spfile
SID='dbm1';
```

4. Validate the configuration files and correct them as necessary:

- The `ORACLE_HOME/dbs/initSID.ora` file points to server parameter file (SPFILE) in the Oracle ASM shared storage.
- The password file that is copied in the `ORACLE_HOME/dbs` directory has been changed to `orapwSID`.

5. Restart the database instance. **See Also:**

Oracle Real Application Clusters Administration and Deployment Guide

12.9 Maintaining the Storage Servers

This section describes how to perform maintenance on the storage servers. It contains the following topics:

- [Shutting Down a Storage Server](#) (page 12-24)
- [Enabling Network Connectivity Using the Diagnostics ISO](#) (page 12-25)

12.9.1 Shutting Down a Storage Server

When performing maintenance on a storage server, you might need to power down or restart the server. Before shutting down a storage server, verify that taking a server offline does not impact Oracle ASM disk group and database availability. Continued database availability depends on the level of Oracle ASM redundancy used on the affected disk groups, and the current status of disks in other storage servers that have mirror copies of the same data.

Caution:

- If a disk in a different cell fails while the cell undergoing maintenance is not completely back in service on the Recovery Appliance, a double disk failure can occur. If the Recovery Appliance is deployed with `NORMAL` redundancy for the `DELTA` disk group and if this disk failure is permanent, you will lose all backups on the Recovery Appliance.
- Ensure that the cell undergoing maintenance is not offline for an extended period of time. Otherwise, a rebalance operation will occur and this will cause issues because of insufficient space for the operation to complete. By default, the rebalance operation begins 24 hours after the cell goes offline.

To power down a storage server:

1. Log in to the storage server as `root`.
2. (Optional) Keep the grid disks offline after restarting the storage server:

```
CellCLI> ALTER GRIDDISK ALL INACTIVE
```

Use this command when doing multiple restarts, or to control when the cell becomes active again. For example, so you can verify the planned maintenance activity was successful before the server is used.

3. Stop the cell services:

```
CellCLI> ALTER CELL SHUTDOWN SERVICES ALL
```

The preceding command checks if any disks are offline, in predictive failure status, or must be copied to its mirror. If Oracle ASM redundancy is intact, then the command takes the grid disks offline in Oracle ASM, and stops the services.

The following error indicates that stopping the services might cause redundancy problems and force a disk group to dismount:

```
Stopping the RS, CELLSRV, and MS services...
The SHUTDOWN of ALL services was not successful.
CELL-01548: Unable to shut down CELLSRV because disk group DATA, RECO may be
forced to dismount due to reduced redundancy.
Getting the state of CELLSRV services... running
Getting the state of MS services... running
Getting the state of RS services... running
```

If this error occurs, then restore Oracle ASM disk group redundancy. Retry the command when the status is normal for all disks.

4. Shut down the server. See "[Powering Down the Servers](#) (page 12-4)".
5. After you complete the maintenance procedure, power up the server. The services start automatically. During startup, all grid disks are automatically online in Oracle ASM.
6. Verify that all grid disks are online:

```
CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatus
```

Wait until `asmmodestatus` shows `ONLINE` or `UNUSED` for all grid disks.

7. If you inactivated the grid disks in step 2 (page 12-24), then reactivate them:

```
CellCLI> ALTER GRIDDISK ALL ACTIVE
```

If you skipped step 2 (page 12-24), then the grid disks are activated automatically.

See Also:

My Oracle Support Doc ID 1188080.1, "Steps to shut down or reboot an Exadata storage cell without affecting ASM."

12.9.2 Enabling Network Connectivity Using the Diagnostics ISO

You might need to use the diagnostics ISO to access a storage server that fails to restart normally. After starting the server, you can copy files from the ISO to the server, replacing the corrupt files.

The ISO is located on all Recovery Appliance servers at `/opt/oracle.SupportTools/diagnostics.iso`.

Caution:

Use the diagnostics ISO only after other restart methods, such as using the USB drive, have failed. Contact Oracle Support for advise and guidance before starting this procedure.

To use the diagnostics ISO:

1. Enable a one-time CD-ROM boot in the service processor, using either the web interface or a serial console, such as Telnet or puTTY. For example, use this command from a serial console:

```
set boot_device=cdrom
```

2. Mount a local copy of `diagnostics.iso` as a CD-ROM, using the service processor interface.
3. Use the `reboot` command to restart the server.
4. Log in to the server as the `root` user with the diagnostics ISO password.

5. To avoid pings:

```
alias ping="ping -c"
```

6. Make a directory named `/etc/network`.
7. Make a directory named `/etc/network/if-pre-up.d`.
8. Add the following settings to the `/etc/network/interfaces` file, entering the actual IP address and netmask of the server, and the IP address of the gateway:

```
iface eth0 inet static
address IP address of server
netmask netmask of server
gateway gateway IP address of server
```

9. Start the `eth0` interface:

```
# ifup eth0
```

Ignore any warning messages.

10. Use either FTP or the `wget` command to retrieve the files needed to repair the server.

12.10 Maintaining the Physical Disks of Storage Servers

This section contains the following topics:

- [About System Disks and Data Disks](#) (page 12-27)
- [Monitoring the Status of Physical Disks](#) (page 12-27)
- [What Happens When Disk Errors Occur?](#) (page 12-28)
- [About Detecting Underperforming Disks](#) (page 12-28)
- [About Rebalancing the Data](#) (page 12-29)
- [Monitoring Hard Disk Controller Write-Through Cache Mode](#) (page 12-30)
- [Replacing a Failed Physical Disk](#) (page 12-31)
- [Replacing a Faulty Physical Disk](#) (page 12-32)
- [Removing an Underperforming Physical Disk](#) (page 12-34)
- [Moving All Drives from One Storage Server to Another](#) (page 12-34)
- [Removing and Replacing the Same Physical Disk](#) (page 12-36)
- [Reenabling a Rejected Physical Disk](#) (page 12-36)

 **See Also:**

Oracle Maximum Availability Architecture (MAA) website at <http://www.oracle.com/goto/maa> for additional information about maintenance best practices

12.10.1 About System Disks and Data Disks

The first two disks of storage servers are system disks. Storage server software system software resides on a portion of each of the system disks. These portions on both system disks are referred to as the system area. The nonsystem area of the system disks, referred to as data partitions, is used for normal data storage. All other disks in a storage server are called data disks.

12.10.2 Monitoring the Status of Physical Disks

You can monitor a physical disk by checking its attributes with the `CellCLI LIST PHYSICALDISK` command. For example, a physical disk with a status of `failed` or `warning - predictive failure` is having problems and probably must be replaced. The disk firmware maintains the error counters, and marks a drive with `Predictive Failure` when internal thresholds are exceeded. The drive, not the server software, determines if it needs replacement.

The following list identifies the storage server physical disk statuses.

Physical Disk Status for Storage Servers

Physical Disk Status

- normal
- normal - dropped for replacement
- normal - confinedOnline
- normal - confinedOnline - dropped for replacement
- not present
- failed
- failed - dropped for replacement
- failed - rejected due to incorrect disk model
- failed - rejected due to incorrect disk model - dropped for replacement
- failed - rejected due to wrong slot
- failed - rejected due to wrong slot - dropped for replacement
- warning - confinedOnline
- warning - confinedOnline - dropped for replacement
- warning - peer failure
- warning - poor performance
- warning - poor performance - dropped for replacement
- warning - poor performance, write-through caching
- warning - predictive failure, poor performance
- warning - predictive failure, poor performance - dropped for replacement
- warning - predictive failure, write-through caching
- warning - predictive failure

warning - predictive failure - dropped for replacement
warning - predictive failure, poor performance, write-through caching
warning - write-through caching

12.10.3 What Happens When Disk Errors Occur?

Oracle ASM performs bad extent repair for read errors caused by hardware errors. The disks stay online, and no alerts are sent.

When a disk fails:

- The Oracle ASM disks associated with it are dropped automatically with the `FORCE` option, and then an Oracle ASM rebalance restores data redundancy.
- The blue LED and the amber LED are turned on for the drive, indicating that disk replacement can proceed. The drive LED stays on solid. See "[LED Status Descriptions](#) (page D-2)" for information about LED status lights during predictive failure and poor performance.
- The server generates an alert, which includes specific instructions for replacing the disk. If you configured the system for alert notifications, then the alert is sent by email to the designated address.

When a disk has a faulty status:

- The Oracle ASM disks associated with the grid disks on the physical drive are dropped automatically.
- An Oracle ASM rebalance relocates the data from the predictively failed disk to other disks.
- The blue LED is turned on for the drive, indicating that disk replacement can proceed.

When Oracle ASM gets a read error on a physically-addressed metadata block, it does not have mirroring for the blocks:

- Oracle ASM takes the disk offline.
- Oracle ASM drops the disk with the `FORCE` option.
- The storage server software sends an alert stating that the disk can be replaced.

See Also:

- "[Replacing a Failed Physical Disk](#) (page 12-31)"
- "[Replacing a Faulty Physical Disk](#) (page 12-32)"

12.10.4 About Detecting Underperforming Disks

ASR automatically identifies and removes a poorly performing disk from the active configuration. Recovery Appliance then runs a set of performance tests. When `CELLSRV` detects poor disk performance, the cell disk status changes to `normal - confinedOnline`, and the physical disk status changes to `warning - confinedOnline`. [Table 12-2](#) (page 12-29) describes the conditions that trigger disk confinement:

Table 12-2 Alerts Indicating Poor Disk Performance

Alert Code	Cause
CD_PERF_HANG	Disk stopped responding
CD_PERF_SLOW_ABS	High service time threshold (slow disk)
CD_PERF_SLOW_RLTV	High relative service time threshold (slow disk)
CD_PERF_SLOW_LAT_WT	High latency on writes
CD_PERF_SLOW_LAT_RD	High latency on reads
CD_PERF_SLOW_LAT_RW	High latency on reads and writes
CD_PERF_SLOW_LAT_ERR	Frequent very high absolute latency on individual I/Os
CD_PERF_IOERR	I/O errors

If the problem is temporary and the disk passes the tests, then it is brought back into the configuration. If the disk does not pass the tests, then it is marked `poor performance`, and ASR submits a service request to replace the disk. If possible, Oracle ASM takes the grid disks offline for testing. Otherwise, the cell disk status stays at `normal - confinedOnline` until the disks can be taken offline safely. See "[Removing an Underperforming Physical Disk](#) (page 12-34)".

The disk status change is recorded in the server alert history:

```
MESSAGE ID date_time info "Hard disk entered confinement status. The LUN
n_m changed status to warning - confinedOnline. CellDisk changed status to normal
- confinedOnline. Status: WARNING - CONFINEDONLINE Manufacturer: name Model
Number: model Size: size Serial Number: serial_number Firmware: fw_release
Slot Number: m Cell Disk: cell_disk_name Grid Disk: grid disk 1, grid disk 2
.
.
Reason for confinement: threshold for service time exceeded"
```

These messages are entered in the storage cell alert log:

```
CDHS: Mark cd health state change cell_disk_name with newState HEALTH_BAD_
ONLINE pending HEALTH_BAD_ONLINE ongoing INVALID cur HEALTH_GOOD
Celldisk entering CONFINE ACTIVE state with cause CD_PERF_SLOW_ABS activeForced: 0
inactiveForced: 0 trigger HistoryFail: 0, forceTestOutcome: 0 testFail: 0
global conf related state: numHDsConf: 1 numFDsConf: 0 numHDsHung: 0 numFDsHung: 0
.
.
.
```

12.10.5 About Rebalancing the Data

After you replace the physical disk, you must re-create the grid disks and cell disks that existed on the previous disk in that slot. If those grid disks were part of an Oracle ASM group, then add them back to the disk group, and rebalance the data, based on the disk group redundancy and the `ASM_POWER_LIMIT` parameter.

Oracle ASM rebalance occurs when dropping or adding a disk. To check the status of the rebalance:

- Did the rebalance operation run successfully?

Check the Oracle ASM alert logs.

- Is the rebalance operation currently running?

Check the `GV$ASM_OPERATION` view.

- Did the rebalance operation fail?

Check the `V$ASM_OPERATION.ERROR` view.

You can perform rebalance operations from multiple disk groups on different Oracle ASM instances in the same cluster, if the failed physical disk contained ASM disks from multiple disk groups. One Oracle ASM instance can run one rebalance operation at a time. If all Oracle ASM instances are busy, then the rebalance operations are queued.

12.10.6 Monitoring Hard Disk Controller Write-Through Cache Mode

The hard disk controller on each storage server periodically performs a discharge and charge of the controller battery. During the operation, the write cache policy changes from write-back caching to write-through caching. Write-through cache mode is slower than write-back cache mode. However, write-back cache mode risks data loss if the storage server loses power or fails. The operation occurs every three months, for example, at 01:00 on the 17th day of January, April, July and October.

This example shows an informational alert that a storage server generates about the status of the caching mode for its logical drives:

```
HDD disk controller battery on disk controller at adapter 0 is going into a learn cycle. This is a normal maintenance activity that occurs quarterly and runs for approximately 1 to 12 hours. The disk controller cache might go into WriteThrough caching mode during the learn cycle. Disk write throughput might be temporarily lower during this time. The message is informational only, no action is required.
```

Use the following commands to manage changes to the periodical write cache policy:

- To change the start time for the learn cycle, use a command like the following example:

```
CellCLI> ALTER CELL bbuLearnCycleTime="2013-01-22T02:00:00-08:00"
```

The time reverts to the default learn cycle time after the cycle completes.

- To see the time for the next learn cycle:

```
CellCLI> LIST CELL ATTRIBUTES bbuLearnCycleTime
```

- To view the status of the battery:

```
# /opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -GetBbuStatus -a0
```

```
BBU status for Adapter: 0
```

```
BatteryType: iBBU08
Voltage: 3721 mV
Current: 541 mA
Temperature: 43 C
```

```
BBU Firmware Status:
```

```
Charging Status : Charging
Voltage : OK
Temperature : OK
Learn Cycle Requested : No
```

```
Learn Cycle Active : No
Learn Cycle Status : OK
Learn Cycle Timeout : No
I2c Errors Detected : No
Battery Pack Missing : No
Battery Replacement required : No
Remaining Capacity Low : Yes
Periodic Learn Required : No
Transparent Learn : No
```

Battery state:

```
GasGaugeStatus:
Fully Discharged : No
Fully Charged : No
Discharging : No
Initialized : No
Remaining Time Alarm : Yes
Remaining Capacity Alarm: No
Discharge Terminated : No
Over Temperature : No
Charging Terminated : No
Over Charged : No
```

```
Relative State of Charge: 7 %
Charger System State: 1
Charger System Ctrl: 0
Charging current: 541 mA
Absolute state of charge: 0 %
Max Error: 0 %
```

Exit Code: 0x00

12.10.7 Replacing a Failed Physical Disk

A physical disk outage can reduce performance and data redundancy. Therefore, you should replace a failed disk with a new disk as soon as possible.

To replace a disk when it fails:

1. Determine which disk failed.

```
CellCLI> LIST PHYSICALDISK WHERE diskType=HardDisk AND status=failed DETAIL
```

```
name:                28:5
deviceId:            21
diskType:            HardDisk
enclosureDeviceId:  28
errMediaCount:      0
errOtherCount:      0
foreignState:        false
luns:                0_5
makeModel:           "SEAGATE ST360057SSUN600G"
physicalFirmware:    0705
physicalInterface:   sas
physicalSerial:      A01BC2
physicalSize:        558.9109999993816G
slotNumber:       5
status:           failed
```

The slot number shows the location of the disk, and the status shows that the disk failed.

2. Ensure that the blue "OK to Remove" LED on the disk is lit, before you remove the disk.
3. Replace the physical disk on the storage server and wait three minutes. The physical disk is hot pluggable, and you can replace it with the power on.
4. Confirm that the disk is online and its status is `NORMAL`:

```
CellCLI> LIST PHYSICALDISK WHERE name=28:5 ATTRIBUTES status
```

When you replace a physical disk, the RAID controller must acknowledge the replacement disk before you can use it. Acknowledgment is quick.

5. Verify that the firmware is correct:

```
ALTER CELL VALIDATE CONFIGURATION
```

You can also check the `ms-odl.trc` file to confirm that the firmware was updated and the logical unit number (LUN) was rebuilt.

6. Re-create the grid disks and cell disks that existed on the previous disk in that slot. See "[About Rebalancing the Data](#) (page 12-29)".

See Also:

- "[Parts for Storage Servers](#) (page D-20)"
- *Oracle Database Reference* about the `V$ASM_OPERATION` view

12.10.8 Replacing a Faulty Physical Disk

You might need to replace a physical disk because its status is `warning - predictive failure`. This status indicates that the physical disk will fail soon, and you should replace it at the earliest opportunity.

If the drive fails before you replace it, then see "[Replacing a Failed Physical Disk](#) (page 12-31)".

To replace a disk before it fails:

1. Identify the faulty disk:

```
CellCLI> LIST PHYSICALDISK WHERE diskType=HardDisk AND status= \
"warning - predictive failure" DETAIL
```

```

name:                28:3
deviceId:            19
diskType:            HardDisk
enclosureDeviceId:  28
errMediaCount:      0
errOtherCount:      0
foreignState:       false
luns:                0_3
makeModel:          "SEAGATE ST360057SSUN600G"
physicalFirmware:   0705
```

```

physicalInterface:    sas
physicalSerial:      E07L8E
physicalSize:        558.9109999993816G
slotNumber:          3
status:              warning - predictive failure

```

In the sample output from the previous command, the slot number shows the location of the disk, and the status shows that the disk is expected to fail.

2. Ensure that the blue "OK to Remove" LED on the disk is lit, before you remove the disk.
3. Wait while the affected Oracle ASM disks are dropped. To check the status, query the `V$ASM_DISK_STAT` view on the Oracle ASM instance.

Caution:

The disks in the first two slots are system disks, which store the operating system and the Recovery Appliance storage server software. One system disk must be in working condition for the server to operate.

Before replacing the other system disk, wait until `ALTER CELL VALIDATE CONFIGURATION` shows no RAID `mdadm` errors. This output indicates that the system disk resynchronization is complete.

See Also:

Oracle Database Reference for information about querying the `V$ASM_DISK_STAT` view

4. Replace the physical disk on the storage server and wait three minutes. The physical disk is hot pluggable, and you can replace it when the power is on.
5. Confirm that the disk is online and its status is `NORMAL`:

```
CellCLI> LIST PHYSICALDISK WHERE name=28:5 ATTRIBUTES status
```

When you replace a physical disk, the RAID controller must acknowledge the replacement disk before you can use it. Acknowledgment is quick.

6. Verify that the firmware is correct:


```
ALTER CELL VALIDATE CONFIGURATION
```
7. Re-create the grid disks and cell disks that existed on the previous disk in that slot. See "[About Rebalancing the Data](#) (page 12-29)".

See Also:

- "[Parts for Storage Servers](#) (page D-20)"
- *Oracle Database Reference* for information about the `V$ASM_OPERATION` view

12.10.9 Removing an Underperforming Physical Disk

A bad physical disk can degrade the performance of other good disks. You should remove the bad disk from the system.

To remove a physical disk after identifying the bad disk:

1. Illuminate the physical drive service LED to identify the drive to be replaced:

```
cellcli -e 'alter physicaldisk disk_name serviceled on'
```

In the preceding command, *disk_name* is the name of the physical disk to be replaced, such as 20:2.

2. Identify all grid disks on the bad disk, and direct Oracle ASM to stop using them:

```
ALTER DISKGROUP diskgroup_name DROP DISK asm_disk_name
```

3. Ensure that the blue "OK to Remove" LED on the disk is lit.
4. Query the `V$ASM_DISK_STAT` view to ensure that the Oracle ASM disks affected by the bad disk were dropped successfully.
5. Remove the bad disk.

An alert is sent when the disk is removed.

6. When a new disk is available, install it in the system. The cell disks and grid disks are created automatically on the new physical disk.
7. Confirm that the disk is online and its status is `NORMAL`:

```
CellCLI> LIST PHYSICALDISK WHERE name=28:5 ATTRIBUTES status
```

When you replace a physical disk, the RAID controller must acknowledge the replacement disk before you can use it. Acknowledgment is quick.



See Also:

["About Detecting Underperforming Disks \(page 12-28\)"](#)

12.10.10 Moving All Drives from One Storage Server to Another

You might need to move all drives from one storage server to another storage server. This situation might occur when a chassis-level component fails, such as a motherboard or Oracle ILOM, or when you are troubleshooting a hardware problem.

To move the drives between storage servers:

1. Back up the files in the following directories:

- `/etc/hosts`
- `/etc/modprobe.conf`
- `/etc/sysconfig/network`
- `/etc/sysconfig/network-scripts`

2. Inactivate all grid disks and shut down the storage server. See "[Shutting Down a Storage Server](#) (page 12-24)".
3. Ensure that the Oracle ASM `disk_repair_time` attribute is set long enough, so that Oracle ASM does not drop the disks before you can activate the grid disks in another storage server.
4. Move the physical disks, flash disks, disk controller, and USB flash drive from the original storage server to the new storage server.

▲ Caution:

- Ensure that the first two disks, which are the system disks, are in the same, first two slots. Otherwise, the storage server will function improperly.
- Ensure that the flash cards are installed in the same PCIe slots as in the original storage server.

5. Power on the new storage server. You can either use the service processor interface or press the power button.
6. Log in to the console using the service processor.
7. Check the files in the following directories. Restore corrupt files from the backups.

- `/etc/hosts`
- `/etc/modprobe.conf`
- `/etc/sysconfig/network`
- `/etc/sysconfig/network-scripts`

8. Use the `ifconfig` command to retrieve the new MAC addresses for `eth0`, `eth1`, `eth2`, and `eth3`. This example shows that the `eth0` MAC address (HWaddr) is `00:14:4F:CA:D9:AE`.

```
# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:14:4F:CA:D9:AE
          inet addr:10.204.74.184  Bcast:10.204.75.255  Mask:255.255.252.0
          inet6 addr: fe80::214:4fff:feca:d9ae/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:141455 errors:0 dropped:0 overruns:0 frame:0
          TX packets:6340 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:9578692 (9.1 MiB)  TX bytes:1042156 (1017.7 KiB)
          Memory:f8c60000-f8c80000
```

9. In the `/etc/sysconfig/network-scripts` directory, edit the following files to change HWADDR to the value returned in step 8 (page 12-35):

```
ifcfg-eth0
ifcfg-eth1
ifcfg-eth2
ifcfg-eth3
```

The following example shows the edited `ifcfg-eth0` file:

```
#### DO NOT REMOVE THESE LINES ####
#### %GENERATED BY CELL% ####
DEVICE=eth0
BOOTPROTO=static
ONBOOT=yes
IPADDR=10.204.74.184
NETMASK=255.255.252.0
NETWORK=10.204.72.0
BROADCAST=10.204.75.255
GATEWAY=10.204.72.1
HOTPLUG=no
IPV6INIT=no
HWADDR=00:14:4F:CA:D9:AE
```

10. Restart the storage server.

11. Activate the grid disks:

```
CellCLI> ALTER GRIDDISK ALL ACTIVE
```

If the Oracle ASM disks were not dropped, then they go online automatically and start being used.

12. Validate the configuration:

```
CellCLI> ALTER CELL VALIDATE CONFIGURATION
```

13. Activate Oracle ILOM for ASR.

12.10.11 Removing and Replacing the Same Physical Disk

If you remove the wrong physical disk and replace it, then Recovery Appliance automatically adds the disk back in the Oracle ASM disk group, and resynchronizes its data.



Note:

When replacing a faulty or failed disk, look for a lit LED on the disk. The LED is lit to help you locate the bad disk.

12.10.12 Reenabling a Rejected Physical Disk

Recovery Appliance rejects a physical disk when it is in the wrong slot.



Caution:

Reenabling a physical disk removes all data stored on it.

- To reenabling a rejected physical disk, replace *hard_disk_name* and *hard_disk_id* with the appropriate values in this command:

```
CellCLI> ALTER PHYSICALDISK hard_disk_name/hard_disk_id reenable force
Physical disk hard_disk_name/hard_disk_id was reenabled.
```

12.11 Maintaining the Flash Disks of Storage Servers

This section describes how to perform maintenance on flash disks. It contains the following topics:

- [About the Flash Disks](#) (page 12-37)
- [Faulty Status Indicators](#) (page 12-38)
- [Identifying Flash Disks in Poor Health](#) (page 12-39)
- [Identifying Underperforming Flash Disks](#) (page 12-39)
- [When Is It Safe to Replace a Faulty Flash Disk?](#) (page 12-40)
- [Replacing a Failed Flash Disk](#) (page 12-40)
- [Replacing a Faulty Flash Disk](#) (page 12-41)
- [Removing an Underperforming Flash Disk](#) (page 12-42)
- [About Write-Back Flash Cache](#) (page 12-43)

12.11.1 About the Flash Disks

Recovery Appliance mirrors data across storage servers, and sends write operations to at least two storage servers. If a flash card in one storage server has problems, then Recovery Appliance services the read and write operations using the mirrored data in another storage server. Service is not interrupted.

If a flash card fails, then the storage server software identifies the data in the flash cache by reading the data from the surviving mirror. It then writes the data to the server with the failed flash card. When the failure occurs, the software saves the location of the data lost in the failed flash cache. Resilvering then replaces the lost data with the mirrored copy. During resilvering, the grid disk status is `ACTIVE -- RESILVERING WORKING`.

Each storage server has four PCIe cards. Each card has four flash disks (FDOMs) for a total of 16 flash disks. The four PCIe cards are located in PCI slot numbers 1, 2, 4, and 5.

To identify a failed flash disk, use the following command:

```
CellCLI> LIST PHYSICALDISK WHERE DISKTYPE=flashdisk AND STATUS=failed DETAIL
```

```
name: FLASH_5_3
diskType: FlashDisk
luns: 5_3
makeModel: "Sun Flash Accelerator F40 PCIe Card"
physicalFirmware: TI35
physicalInsertTime: 2012-07-13T15:40:59-07:00
physicalSerial: 5L002X4P
physicalSize: 93.13225793838501G
slotNumber: "PCI Slot: 5; FDOM: 3"
status: failed
```

The card `name` and `slotNumber` attributes show the PCI slot and the FDOM number.

When the server software detects a failure, it generates an alert that indicates that the flash disk, and the LUN on it, failed. The alert message includes the PCI slot number

of the flash card and the exact FDOM number. These numbers uniquely identify the field replaceable unit (FRU). If you configured the system for alert notification, then the alert is sent to the designated address in an email message.

A flash disk outage can reduce performance and data redundancy. Replace the failed disk at the earliest opportunity. If the flash disk is used for flash cache, then the effective cache size for the server is reduced. If the flash disk is used for flash log, then the flash log is disabled on the disk, thus reducing the effective flash log size. If the flash disk is used for grid disks, then the Oracle ASM disks associated with them are automatically dropped with the `FORCE` option from the Oracle ASM disk group, and an Oracle ASM rebalance starts to restore the data redundancy.

See Also:

- "[Parts for Storage Servers](#) (page D-20)" for part number information and a link to the service guide
- *Oracle Database Reference* for information about the `V$ASM_OPERATION` view
- *Sun Flash Accelerator F80 PCIe Card User's Guide* at http://docs.oracle.com/cd/E41278_01/html/E41251/index.html

12.11.2 Faulty Status Indicators

The following status indicators generate an alert. The alert includes specific instructions for replacing the flash disk. If you configured the system for alert notifications, then the alerts are sent by email message to the designated address.

warning - peer failure

One of the flash disks on the same Sun Flash Accelerator PCIe card failed or has a problem. For example, if `FLASH5_3` fails, then `FLASH5_0`, `FLASH5_1`, and `FLASH5_2` have peer failure status:

```
CellCLI> LIST PHYSICALDISK
36:0          L45F3A          normal
36:1          L45WAE          normal
36:2          L45WQW          normal
.
.
FLASH_5_0     5L0034XM        warning - peer failure
FLASH_5_1     5L0034JE        warning - peer failure
FLASH_5_2     5L002WJH        warning - peer failure
FLASH_5_3     5L002X4P        failed
```

warning - predictive failure

The flash disk will fail soon, and should be replaced at the earliest opportunity. If the flash disk is used for flash cache, then it continues to be used as flash cache. If the flash disk is used for grid disks, then the Oracle ASM disks associated with these grid disks are automatically dropped, and Oracle ASM rebalance relocates the data from the predictively failed disk to other disks.

When one flash disk has predictive failure status, then the data is copied. If the flash disk is used for write back flash cache, then the data is flushed from the flash disks to the grid disks.

warning - poor performance

The flash disk demonstrates extremely poor performance, and should be replaced at the earliest opportunity. If the flash disk is used for flash cache, then flash cache is dropped from this disk, thus reducing the effective flash cache size for the storage server. If the flash disk is used for grid disks, then the Oracle ASM disks associated with the grid disks on this flash disk are automatically dropped with the `FORCE` option, if possible. If `DROP...FORCE` cannot succeed because of offline partners, then the grid disks are dropped normally, and Oracle ASM rebalance relocates the data from the poor performance disk to the other disks.

warning - write-through caching

The capacitors used to support data cache on the PCIe card failed, and the card should be replaced as soon as possible.

12.11.3 Identifying Flash Disks in Poor Health

To identify a flash disk with a particular health status, use the `LIST PHYSICALDISK` command. This example queries for the `warning - predictive failure` status:

```
CellCLI> LIST PHYSICALDISK WHERE DISKTYPE=flashdisk AND STATUS= \
'warning - predictive failure' DETAIL

name:                FLASH_5_3
diskType:            FlashDisk
luns:                5_3
makeModel:           "Sun Flash Accelerator F40 PCIe Card"
physicalFirmware:    TI35
physicalInsertTime:  2012-07-13T15:40:59-07:00
physicalSerial:      5L002X4P
physicalSize:        93.13225793838501G
slotNumber:          "PCI Slot: 1; FDOM: 2"
status:              warning - predictive failure
```

12.11.4 Identifying Underperforming Flash Disks

ASR automatically identifies and removes a poorly performing disk from the active configuration. Recovery Appliance then runs a set of performance tests. When `CELLSRV` detects poor disk performance, the cell disk status changes to `normal - confinedOnline`, and the physical disk status changes to `warning - confinedOnline`. [Table 12-2](#) (page 12-29) describes the conditions that trigger disk confinement. The conditions are the same for both physical and flash disks.

If the problem is temporary and the disk passes the tests, then it is brought back into the configuration. If the disk does not pass the tests, then it is marked `poor performance`, and ASR submits a service request to replace the disk. If possible, Oracle ASM takes the grid disks offline for testing. Otherwise, the cell disk status stays at `normal - confinedOnline` until the disks can be taken offline safely.

The disk status change is recorded in the server alert history:

```
MESSAGE ID date_time info "Hard disk entered confinement status. The LUN
n_m changed status to warning - confinedOnline. CellDisk changed status to normal
```

```
- confinedOnline. Status: WARNING - CONFINEDONLINE Manufacturer: name Model
Number: model Size: size Serial Number: serial_number Firmware: fw_release
Slot Number: m Cell Disk: cell_disk_name Grid Disk: grid disk 1, grid disk 2
... Reason for confinement: threshold for service time exceeded"
```

These messages are entered in the storage cell alert log:

```
CDHS: Mark cd health state change cell_disk_name with newState HEALTH_BAD_
ONLINE pending HEALTH_BAD_ONLINE ongoing INVALID cur HEALTH_GOOD
Celldisk entering CONFINE ACTIVE state with cause CD_PERF_SLOW_ABS activeForced: 0
inactiveForced: 0 trigger HistoryFail: 0, forceTestOutcome: 0 testFail: 0
global conf related state: numHdsConf: 1 numFDsConf: 0 numHdsHung: 0 numFDsHung: 0
.
.
.
```

12.11.5 When Is It Safe to Replace a Faulty Flash Disk?

When the server software detects a predictive or peer failure in a flash disk used for write back flash cache, and only one FDOM is bad, then the server software resilvers the data on the bad FDOM, and flushes the data on the other three FDOMs. If there are valid grid disks, then the server software initiates an Oracle ASM rebalance of the disks. You cannot replace the bad disk until the tasks are completed and an alert indicates that the disk is ready.

An alert is sent when the Oracle ASM disks are dropped, and you can safely replace the flash disk. If the flash disk is used for write-back flash cache, then wait until none of the grid disks are cached by the flash disk.

12.11.6 Replacing a Failed Flash Disk

Caution:

The PCIe cards are not hot pluggable; you must power down a storage server before replacing the flash disks or cards.

Before you perform the following procedure, shut down the server. See "[Shutting Down a Storage Server](#) (page 12-24)".

To replace a failed flash disk:

1. Replace the failed flash disk. Use the PCI number and FDOM number to locate the failed disk. A white cell LED is lit to help you locate the affected server.
2. Power up the server. The services start automatically. As part of the server startup, all grid disks are automatically online in Oracle ASM.
3. Verify that all grid disks are online:

```
CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatus
```

Wait until `asmmodestatus` shows `ONLINE` or `UNUSED` for all grid disks.

 **See Also:**

- "[Parts for Storage Servers](#) (page D-20)" for part numbers and links to the service guide
- *Oracle Database Reference* for information about the `V$ASM_OPERATION` view
- *Sun Flash Accelerator F80 PCIe Card User's Guide* at http://docs.oracle.com/cd/E41278_01/html/E41251/index.html

12.11.7 Replacing a Faulty Flash Disk

 **Caution:**

The PCIe cards are not hot pluggable; you must power down a storage server before replacing the flash disks or cards.

Before you perform the following procedure, review the "[When Is It Safe to Replace a Faulty Flash Disk?](#) (page 12-40)" topic.

To replace a faulty flash disk:

1. Use the following command to check the `cachedBy` attribute of all grid disks.

```
CellCLI> LIST GRIDDISK ATTRIBUTES name, cachedBy
```

The cell disk on the flash disk should not appear in any grid disk `cachedBy` attribute. If the flash disk is used for both grid disks and flash cache, then wait until receiving the alert, and the cell disk is not shown in any grid disk `cachedBy` attribute.

2. Stop all services:

```
CellCLI> ALTER CELL SHUTDOWN SERVICES ALL
```

The preceding command checks if any disks are offline, in predictive failure status, or must be copied to a mirror. If Oracle ASM redundancy is intact, then the command takes the grid disks offline in Oracle ASM, and then stops the services.

The following error indicates that it might be unsafe to stop the services, because stopping them might force a disk group to dismount:

```
Stopping the RS, CELLSRV, and MS services...
The SHUTDOWN of ALL services was not successful.
CELL-01548: Unable to shut down CELLSRV because disk group DATA, RECO may be
forced to dismount due to reduced redundancy.
Getting the state of CELLSRV services... running
Getting the state of MS services... running
Getting the state of RS services... running
```

If this error occurs, then restore Oracle ASM disk group redundancy, and retry the command when the disk status is normal for all disks.

3. Shut down the server.

See "[Shutting Down a Storage Server](#) (page 12-24)".

4. Replace the failed flash disk. Use the PCI number and FDOM number to locate the failed disk. A white cell LED is lit to help you locate the affected server.
5. Power up the server. The services start automatically. As part of the server startup, all grid disks are automatically online in Oracle ASM.
6. Verify that all grid disks are online:

```
CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatus
```

Wait until `asmmodestatus` shows `ONLINE` or `UNUSED` for all grid disks.

The system automatically uses the new flash disk, as follows:

- If the flash disk is used for flash cache, then the effective cache size increases.
- If the flash disk is used for grid disks, then the grid disks are re-created on the new flash disk.
- If the grid disks were part of an Oracle ASM disk group, then they are added back to the disk group. The data is rebalanced on them, based on the disk group redundancy and the `ASM_POWER_LIMIT` parameter.

12.11.8 Removing an Underperforming Flash Disk

A bad flash disk can degrade the performance of other good flash disks. You should remove a bad flash disk. See "[Identifying Underperforming Flash Disks](#) (page 12-39)".

To remove an underperforming flash drive:

1. If the flash disk is used for flash cache:
 - a. Ensure that data not synchronized with the disk (dirty data) is flushed from flash cache to the grid disks:


```
CellCLI> ALTER FLASHCACHE ... FLUSH
```
 - b. Disable the flash cache and create a new one. Do not include the bad flash disk when creating the flash cache.


```
CellCLI > DROP FLASHCACHE
CellCLI > CREATE FLASHCACHE CELLDISK='fd1,fd2,fd3,fd4, ...'
```
2. If the flash disk is used for grid disks, then direct Oracle ASM to stop using the bad disk immediately:

```
SQL> ALTER DISKGROUP diskgroup_name DROP DISK asm_disk_name FORCE
```

Offline partners might cause the `DROP` command with the `FORCE` option to fail. If the previous command fails, do one of the following:

- Restore Oracle ASM data redundancy by correcting the other server or disk failures. Then retry the `DROP...FORCE` command.
- Direct Oracle ASM to rebalance the data off the bad disk:

```
SQL> ALTER DISKGROUP diskgroup_name DROP DISK asm_disk_name NOFORCE
```

3. Wait until the Oracle ASM disks associated with the bad flash disk are dropped successfully. The storage server software automatically sends an alert when it is safe to replace the flash disk.
4. Stop the services:

```
CellCLI> ALTER CELL SHUTDOWN SERVICES ALL
```

The preceding command checks if any disks are offline, in predictive failure status, or must be copied to its mirror. If Oracle ASM redundancy is intact, then the command takes the grid disks offline in Oracle ASM, and stops the services.

The following error indicates that stopping the services might cause redundancy problems and force a disk group to dismount:

```
Stopping the RS, CELLSRV, and MS services...
The SHUTDOWN of ALL services was not successful.
CELL-01548: Unable to shut down CELLSRV because disk group DATA, RECO may be
forced to dismount due to reduced redundancy.
Getting the state of CELLSRV services... running
Getting the state of MS services... running
Getting the state of RS services... running
```

If this error occurs, then restore Oracle ASM disk group redundancy. Retry the command when the status is normal for all disks.

5. Shut down the server. See "[Shutting Down a Storage Server](#) (page 12-24)".
6. Remove the bad flash disk, and replace it with a new flash disk.
7. Power up the server. The services are started automatically. As part of the server startup, all grid disks are automatically online in Oracle ASM.
8. Add the new flash disk to flash cache:

```
CellCLI> DROP FLASHCACHE
CellCLI> CREATE FLASHCACHE ALL
```

9. Verify that all grid disks are online:

```
CellCLI> LIST GRIDDISK ATTRIBUTES asmmodestatus
```

Wait until `asmmodestatus` shows `ONLINE` or `UNUSED` for all grid disks.

The flash disks are added as follows:

- If the flash disk is used for grid disks, then the grid disks are re-created on the new flash disk.
- If these grid disks were part of an Oracle ASM disk group and `DROP...FORCE` was used in Step 2 (page 12-42), then they are added back to the disk group and the data is rebalanced on based on disk group redundancy and the `ASM_POWER_LIMIT` parameter.
- If `DROP...NOFORCE` was used in Step 2 (page 12-42), then you must manually add the grid disks back to the Oracle ASM disk group.

12.11.9 About Write-Back Flash Cache

You cannot modify the write-back flash cache settings on Recovery Appliance.

12.12 Replacing a Disk Controller Battery Backup Unit

The disk controller battery backup unit (disk controller BBU) resides on a drive tray in the compute and storage servers. You can replace the disk controller BBU without downtime. The following procedures describe how to replace the disk controller BBU:

- [Replacing a Disk Controller BBU on a Compute Server](#) (page 12-44)

- [Replacing a Disk Controller BBU on a Storage Server \(page 12-45\)](#)



Note:

The procedures in this section do not apply to on-controller battery backup units. Replacement of those units require a system shutdown, because the system must be opened to access the controller card.

12.12.1 Replacing a Disk Controller BBU on a Compute Server

The following procedure describes how to replace a disk controller BBU on a compute server:

1. Drop the disk controller BBU for replacement:

```
# /opt/oracle.cellos/compmon/exadata_mon_hw_asr.pl -drop_bbu_for_replacement
```

2. Verify that the disk controller BBU has been dropped for replacement:

```
# /opt/oracle.cellos/compmon/exadata_mon_hw_asr.pl -list_bbu_status
```

```
BBU status: dropped for replacement.
```

3. Replace the disk controller BBU by releasing the drive caddy and slowly pulling out the tray, and then sliding the replacement tray into the slot. The disk controller BBU is located in slot 7.

4. Verify that the new disk controller BBU has been detected. It may take several minutes.

```
# /opt/oracle.cellos/compmon/exadata_mon_hw_asr.pl -list_bbu_status
```

```
BBU status: present
```

5. Verify that the current logical disk drive cache policy uses `writeback` mode.

```
# /opt/MegaRAID/MegaCli/MegaCli64 -ldinfo -lall -a0 | egrep \
'Default Cache|Current Cache'
Default Cache Policy: WriteBack, ReadAheadNone, Direct, No Write Cache if
Bad BBU
Current Cache Policy: WriteBack, ReadAheadNone, Direct, No Write Cache if
Bad BBU
```

If the cache policy is not `writeback`, then go to step 6 (page 12-44). Otherwise, go to step 7 (page 12-45).

6. Verify that the battery state is `Operational`. This step is required only when the cache policy output from step 5 (page 12-44) is not `writeback`.

```
# /opt/MegaRAID/MegaCli/MegaCli64 -adpbucmd -getbbustatus -a0|grep Battery
BatteryType: iBBU08
Battery State : Operational
Battery Pack Missing : No
Battery Replacement required : No
```

If the battery state is not `Operational`, then investigate and correct the problem.

7. Perform battery checks as described in My Oracle Support Doc ID 1274318.1. If the checks return unexpected results, then refer to the note for additional information and instructions.
8. (Optional) Use the `exachk` tool to verify the health of the system. See My Oracle Support Doc ID 1070954.1.

12.12.2 Replacing a Disk Controller BBU on a Storage Server

To replace a disk controller BBU on a storage server:

1. Drop the disk controller BBU for replacement using the following command:

```
# cellcli -e alter cell bbu drop for replacement
```

2. Verify that the disk controller BBU has been dropped for replacement using the following command:

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

```
BBU status: dropped for replacement.
```

3. Replace the disk controller BBU by releasing the drive caddy and slowly pulling out the tray, and then sliding the replacement tray into the slot. The disk controller BBU is located in rear slot 1 of the server.
4. Verify that the disk controller BBU battery state is operational.

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

```
BBU status: normal
```

5. Perform battery checks as described in My Oracle Support Doc ID 1274318.1. If the checks return unexpected results, then refer to the note for additional information and instructions.
6. (Optional) Use the `exachk` tool to verify the health of the system. The tool is available in My Oracle Support Doc ID 1070954.1.

12.13 Using the Storage Server Rescue Procedure

Each storage server maintains a copy of the software on the USB stick. Whenever the system configuration changes, the server updates the USB stick. You can use this USB stick to recover the server after a hardware replacement or a software failure. You restore the system when the system disks fail, the operating system has a corrupt file system, or the boot area is damaged. You can replace the disks, cards, CPU, memory, and so forth, and recover the server. You can insert the USB stick in a different server, and it will duplicate the old server.

If only one system disk fails, then use CellCLI commands to recover. In the rare event that both system disks fail simultaneously, then use the rescue functionality provided on the storage server CELLBOOT USB flash drive.

This section contains the following topics:

- [First Steps Before Rescuing the Storage Server](#) (page 12-46)
- [About the Rescue Procedure](#) (page 12-46)
- [Rescuing a Server Using the CELLBOOT USB Flash Drive](#) (page 12-47)
- [Reconfiguring the Rescued Storage Server](#) (page 12-48)

- [Recreating a Damaged CELLBOOT USB Flash Drive](#) (page 12-49)

12.13.1 First Steps Before Rescuing the Storage Server

Before rescuing a storage server, you must take steps to protect the data that is stored on it. Those steps depend on whether the system is set up with normal redundancy or high redundancy.

12.13.1.1 If the Server Has Normal Redundancy

If you are using normal redundancy, then the server has one mirror copy. The data could be irrecoverably lost, if that single mirror also fails during the rescue procedure.

Oracle recommends that you duplicate the mirror copy:

1. Make a complete backup of the data in the mirror copy.
2. Take the mirror copy server offline immediately, to prevent any new data changes to it before attempting a rescue.

This procedure ensures that all data residing on the grid disks on the failed server and its mirror copy is inaccessible during the rescue procedure.

The Oracle ASM disk repair timer has a default repair time of 3.6 hours. If you know that you cannot perform the rescue procedure within that time frame, then use the Oracle ASM rebalance procedure to rebalance the disks until you can do the rescue procedure.



See Also:

Oracle Exadata Storage Server Software User's Guide for information about resetting the timer

12.13.1.2 If the Server Has High Redundancy

When the server has high redundancy disk groups, so that Oracle ASM has multiple mirror copies for all the grid disks of the failed server, then take the failed cell offline. After Oracle ASM times out, it automatically drops the grid disks on the failed server, and starts rebalancing the data using mirror copies.

The default time out is two hours. If the server rescue takes more than two hours, then you must re-create the grid disks on the rescued cells in Oracle ASM.

12.13.2 About the Rescue Procedure

Note the following before using the rescue procedure:

- The rescue procedure can rewrite some or all of the disks in the cell. If this happens, then you might lose all the content of those disks without the possibility of recovery. Ensure that you complete the appropriate preliminary steps before starting the rescue. See "[If the Server Has Normal Redundancy](#) (page 12-46)" or "[If the Server Has High Redundancy](#) (page 12-46)".

- Use extreme caution when using this procedure, and pay attention to the prompts. Ideally, use the rescue procedure only with assistance from Oracle Support Services, and when you can afford to lose the data on some or all of the disks.
- The rescue procedure does not destroy the contents of the data disks or the contents of the data partitions on the system disks, unless you explicitly choose to do so during the rescue procedure.
- The rescue procedure restores the storage server software to the same release, including any patches that existed on the server during the last successful boot.
- The rescue procedure does *not* restore these configuration settings:
 - Server configurations, such as alert configurations, SMTP information, administrator email address
 - ILOM configuration. However, ILOM configurations typically remain undamaged even when the server software fails.
- The recovery procedure *does* restore these configuration settings:
 - The network configuration that existed at the end of last successful run of the `/usr/local/bin/ipconf` utility
 - The SSH identities for the server
 - The `root`, `celladmin` and `cellmonitor` users

 **See Also:**

Oracle Exadata Storage Server Software User's Guide for information about the `ALTER CELL` command

- The rescue procedure does not examine or reconstruct data disks or data partitions on the system disks. If there is data corruption on the grid disks, then do not use this rescue procedure. Instead, use the rescue procedures for Oracle Database and Oracle ASM.

After a successful rescue, you must reconfigure the server. If you want to preserve the data, then import the cell disks. Otherwise, you must create new cell disks and grid disks.

 **See Also:**

Oracle Exadata Storage Server Software User's Guide for information on configuring cells, cell disks, and grid disks using the `CellCLI` utility

12.13.3 Rescuing a Server Using the CELLBOOT USB Flash Drive

 **Caution:**

Follow the rescue procedure with care to avoid data loss.

To rescue a server using the CELLBOOT USB flash drive:

1. Connect to the Oracle ILOM service processor (SP) of the rescued server. You can use either HTTPS or SSH.
2. Start the server. As soon as you see the splash screen, press any key on the keyboard. The splash screen is visible for only 5 seconds.
3. In the displayed list of boot options, select the last option, `CELL_USB_BOOT_CELLBOOT_usb_in_rescue_mode`, and press Enter.
4. Select the rescue option, and proceed with the rescue.
5. At the end of the first phase of the rescue, choose the option to enter the shell. Do not restart the system.
6. Log in to the shell using the rescue `root` password.
7. Use the `reboot` command from the shell.
8. Press F8 as the server restarts and before the splash screen appears. Pressing F8 accesses the boot device selection menu.
9. Select the RAID controller as the boot device. This causes the server to boot from the hard disks.

 **Note:**

Additional options might be available that allow you to enter a rescue mode Linux login shell with limited functionality. Then you can log in to the shell as the `root` user with the password supplied by Oracle Support Services, and manually run additional diagnostics and repairs on the server. For complete details, contact your Oracle Support Services representative.

12.13.4 Reconfiguring the Rescued Storage Server

After a successful rescue, you must configure the server. If the data partitions were preserved, then the cell disks are imported automatically during the rescue procedure.

1. For any replaced servers, re-create the cell disks and grid disks.
2. Log in to the Oracle ASM instance, and set the disks to `ONLINE` using the following command for each disk group:

```
SQL> ALTER DISKGROUP disk_group_name ONLINE DISKS IN FAILGROUP \  
cell_name WAIT;
```

3. Reconfigure the cell using the `ALTER CELL` command. The following example shows the most common parameters:

```
CellCLI> ALTER CELL  
smtpServer='my_mail.example.com', -  
smtpFromAddr='john.doe@example.com', -  
smtpFromPwd=email_address_password, -  
smtpToAddr='jane.smith@example.com', -  
notificationPolicy='critical,warning,clear', -  
notificationMethod='mail,snmp'
```

4. Re-create the I/O Resource Management (IORM) plan.

5. Re-create the metric thresholds.

 **See Also:**

Oracle Exadata Storage Server Software User's Guide for information about IORM plans and metric thresholds

12.13.5 Recreating a Damaged CELLBOOT USB Flash Drive

If the CELLBOOT USB flash drive is lost or damaged, then you can create another one.

To create a CELLBOOT flash drive:

1. Log in to the server as the `root` user.
2. Attach a new USB flash drive with a capacity of 1 to 8 GB.
3. Remove any other USB flash drives from the system.
4. Change directories:

```
cd /opt/oracle.SupportTools
```

5. Copy the server software to the flash drive:

```
./make_cellboot_usb -verbose -force
```

13

Reconfiguring the Hardware

This chapter describes how to make basic adjustments. It contains the following topics:

- [Changing Component Passwords](#) (page 13-1)
- [Changing the DNS Servers](#) (page 13-5)
- [Changing the NTP Servers](#) (page 13-8)
- [Changing the Time Zone Settings](#) (page 13-11)

13.1 Changing Component Passwords

Oracle recommends that you change the passwords for the various rack components after the initial hardware and software configuration:

- [Changing the Compute Server Passwords](#) (page 13-1)
- [Changing the Storage Server Passwords](#) (page 13-2)
- [Creating New Users and Changing Passwords on a PDU](#) (page 13-2)
- [Changing the Oracle ILOM Password](#) (page 13-3)
- [Changing an InfiniBand Switch Password](#) (page 13-3)
- [Changing the Cisco Ethernet Switch Password](#) (page 13-4)
- [Changing the Recovery Appliance Database User Password](#) (page 13-4)
- [Changing OSB Tape Backup Passwords](#) (page 13-5)

See Also:

[Default User Accounts](#) (page B-1) for a list of default users and passwords of the various Recovery Appliance components.

13.1.1 Changing the Compute Server Passwords

You can change the user accounts and GRUB passwords on the compute servers. default user accounts on a compute server are `root` and the software owner account. Typically the software owner account is `oracle`.

13.1.1.1 Changing the User Account Password

A compute server has two default user accounts: `root` and the software owner (typically `oracle`).

To change a user account password on a compute server:

1. Log in to the compute server as `root`.
2. Use the Linux `passwd` command. This example changes the password of the `oracle` account:

```
# passwd oracle
Changing password for user oracle.
Changing password for oracle
(current) UNIX password: old_password
New UNIX password: new_password
Retype new UNIX password: new_password
```

13.1.1.2 Changing the GRUB Account Password

To change the GRUB account password on a compute server:

1. Log in to the compute server as `root`.
2. Use the following command:

```
# grub-md5-crypt
```

The system prompts for a new password twice.

3. Copy the string generated by `grub-md5-crypt` to the copy buffer.
4. Open the `/boot/grub/grub.conf` file in a text editor and locate the password setting. It looks similar to the following:

```
password --md5 hashed_string
```

5. Replace the existing hashed string with the one copied from the `grub-md5-crypt` command output.
6. Save the file.

13.1.2 Changing the Storage Server Passwords

The default user accounts on storage servers are `root`, `celladmin`, and `cellmonitor`.

To change a user account password on a storage server:

1. Log in to the compute server as `root`.
2. Use the Linux `passwd` command. This example changes the password of the `celladmin` account:

```
# passwd celladmin
Changing password for user celladmin.
Changing password for celladmin
(current) UNIX password: old_password
New UNIX password: new_password
Retype new UNIX password: new_password
```

13.1.3 Creating New Users and Changing Passwords on a PDU

The default account user for the power distribution unit (PDU) is `admin`. You can create up to five users and set their passwords.

To create users and passwords on a PDU:

1. Open a web browser and enter the IP address of the PDU in the address field.

The Current Measurement page appears.

2. Click **Network Configuration** in the upper left of the page.
3. Log in as the `admin` user on the PDU metering unit.
4. Locate the Admin/User fields. Only letters and numbers are allowed for user names and passwords.
5. Enter up to five user names and passwords in the Admin/Users fields.
6. Designate each user as either an administrator or a user.
7. Click **Submit** to set the user names and passwords.

13.1.4 Changing the Oracle ILOM Password

The default user account for Oracle Integrated Lights Out Manager (ILOM) is `root`.

To change the password for an Oracle ILOM user:

1. Connect to Oracle ILOM using SSH, and log in as `root`.
2. Use the following command to change the password:

```
set /SP/users/user_name password
```

This example changes the password for `user1`:

```
set /SP/users/user1 password
```

```
Changing password for user /SP/users/user1/password...
```

```
Enter new password:*****
```

```
Enter new password again:*****
```

```
New password was successfully set for user /SP/users/user1
```

13.1.5 Changing an InfiniBand Switch Password

The default user accounts on the InfiniBand switches are `root`, `ilom-admin`, `ilom-user`, `ilom-operator`, and `nm2user`.

To change a password for a user on an InfiniBand switch:

1. Connect to the InfiniBand switch using SSH and log in:

```
ssh user_name@switch_name
```

In the preceding command, `user_name` is the name of the user, and `switch_name` is the name of the InfiniBand switch.

2. Check the firmware version of the switch.
3. Use Oracle ILOM to change the password:

```
ssh -l ilom-admin switch_name
```

```
set /SP/users/user_name password
```

**See Also:**

Sun Datacenter InfiniBand Switch 36 User's Guide at

http://docs.oracle.com/cd/E36265_01/index.html

13.1.6 Changing the Cisco Ethernet Switch Password

There are no user accounts on the Cisco Ethernet switch. The password, which is called the "enable" password, is not specific to a user account.

To change the password for the Cisco Ethernet switch:

1. Change to enable mode:

```
Switch> enable
```

2. Set the password:

```
rasw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
rasw-ip(config)# enable password password
rasw-ip(config)# enable secret password
rasw-ip(config)# end
rasw-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 ]
```

3. Save the current configuration:

```
rasw-ip# copy running-config startup-config
```

4. Exit from the session:

```
rasw-ip# exit
```

13.1.7 Changing the Recovery Appliance Database User Password

Oracle recommends that you change the password of the Recovery Appliance default database user account.

To change the Recovery Appliance database user password:

1. Log in to the compute server as `oracle`.
2. Start SQL*Plus and connect to the Recovery Appliance database with the `sysdba` privilege:

```
$ sqlplus / as sysdba
```

3. Run the `alter user` command to set the new password for the `RASYS` account:

```
SQL> ALTER USER RASYS IDENTIFIED BY *****;
```

 **See Also:**

[Default User Accounts](#) (page B-1) for information about the default Recovery Appliance database accounts and passwords.

13.1.8 Changing OSB Tape Backup Passwords

The OSB tape backup default accounts include the application users `admin` and `oracle`, and the encryption key wallet account.

To change an OSB tape backup application user password:

1. Log in to the compute server as `root`.
2. Use the `obtool chuser` command.

This example changes the password of the `admin` account:

```
# obtool chuser -q admin
Password:*****
Password (again):*****
#
```

To change the password for the OSB tape backup encryption key wallet:

1. Log in to the compute server as `root`.
2. Run the following commands:

```
# export OBCONFIG=/osbcat/etc/obconfig
# obcm chpass --keywallet
Old wallet password:*****
New wallet password:*****
New wallet password (again):*****
Wallet password has been reset
#
```

 **See Also:**

[Default User Accounts](#) (page B-1) for information about the default OSB tape backup accounts and passwords for Recovery Appliance.

13.2 Changing the DNS Servers

You can change the configuration settings for the Domain Name System (DNS) servers after the initial setup. All servers and switches in Recovery Appliance must reference the same DNS servers. This section explains how to set the servers and switches to the same DNS servers. Oracle recommends changing the servers one at a time.

- ["Changing the DNS Server Address on an InfiniBand Switch \(page 13-6\)"](#)
- ["Changing the DNS Server Address on the Cisco Ethernet Switch \(page 13-6\)"](#)
- ["Changing the DNS Server Address on the Compute Server \(page 13-7\)"](#)

- ["Changing the DNS Server on Storage Servers \(page 13-7\)"](#)

13.2.1 Changing the DNS Server Address on an InfiniBand Switch

To change the DNS server address on an InfiniBand switch, connect to Oracle ILOM and log in as the `ilom-admin` user. You can choose from one of the following methods:

- **Oracle ILOM web interface**

Select the **Configuration** tab and set the DNS server addresses.

- **Oracle ILOM CLI**

Use the following command to set the DNS:

```
set /SP/clients/dns nameserver=dns_ip
```

In the preceding command, `dns_ip` is the IP address of the DNS server. For multiple servers, enter a comma-separated list. For example:

```
set /SP/clients/dns nameserver=dns_ip1,dns_ip2,dns_ip3
```

13.2.2 Changing the DNS Server Address on the Cisco Ethernet Switch

To change the DNS server address on the Cisco Ethernet switch:

1. Access the Ethernet switch using SSH, and log in as the `admin` user with the administrator password.

If SSH is not configured, then use Telnet to access the switch.

2. Change to enable mode:

```
Switch> enable
```

3. List the DNS IP addresses currently in use:

```
Switch# show running-config
```

4. Erase the current DNS server addresses and any invalid IP addresses. This example erases three IP addresses:

```
Switch# configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch(config)# no ip name-server 10.7.7.2  
Switch(config)# no ip name-server 129.148.5.4  
Switch(config)# no ip name-server 10.8.160.2  
Switch(config)# end  
Switch# write memory  
Building configuration...  
Compressed configuration from 2603 bytes to 1158 bytes [OK ]
```

5. Configure up to three DNS servers. The following example configures three new DNS servers:

```
Switch# configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
Switch(config)# ip name-server 10.7.7.3  
Switch(config)# ip name-server 129.148.5.5  
Switch(config)# ip name-server 10.8.160.1  
Switch(config)# write memory
```

```
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes [OK ]
```

6. Verify the changes:

```
Switch# show running-config
Building configuration...
...
ip domain-name example.com
ip name-server 192.168.10.2
ip name-server 192.168.10.3
ip name-server 192.168.10.4
...
```

7. Save the running configuration as the startup configuration:

```
Switch# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
Compressed configuration from 14343 bytes to 3986 bytes[OK]
```

8. End the session:

```
Switch# exit
```

13.2.3 Changing the DNS Server Address on the Compute Server

To change the DNS server address on the compute servers:

1. Log in as the `root` user.
2. Open `/etc/resolv.conf` file in a text editor.
3. Locate the name server setting for the DNS server and change it to the new value. For example:

```
search      example.com
nameserver  10.7.7.3
```

4. Save the file and close the editor.
5. Set the DNS server in the server ILOM:

```
ipmitool sunoem cli 'set /SP/clients/dns nameserver=dns_ip'
```

In the preceding command, `dns_ip` is the IP address of the DNS server. For multiple DNS servers, enter a comma-separated list. For example:

```
set /SP/clients/dns nameserver=dns_ip1,dns_ip2,dns_ip3
```

6. Repeat this procedure for each compute server.

13.2.4 Changing the DNS Server on Storage Servers

To change the DNS server or NTP settings on a storage server:

1. Log in to the storage server as the `root` user.
2. Follow steps 1 through 5 of My Oracle Support Doc ID 1188080.1 to take the server offline.
3. Shut down the cell services and `ocrvottargetd` service:

```
# cellcli -e alter cell shutdown services all
# service ocrvottargetd stop
```

4. Use the `ipconf` utility to change the DNS settings:

```
# /usr/local/bin/ipconf
```
5. Restart the server:

```
# shutdown -r now
```
6. Follow steps 7 through 9 of My Oracle Support note 1188080.1 to put the server back online.
7. Repeat this procedure for each storage server.

13.3 Changing the NTP Servers

You can change the configuration settings for the Network Time Protocol (NTP) servers after the initial setup. All servers and switches in Recovery Appliance must reference the same NTP servers, so that they are synchronized to the same time. This section explains how to set the Recovery Appliance servers and switches to the same NTP server addresses. Oracle recommends changing each server one at a time.

This section contains the following tasks:

- ["Setting the NTP Server Address on the Compute Servers \(page 13-8\)"](#)
- ["Setting the NTP Server Address on the InfiniBand Switches \(page 13-9\)"](#)
- ["Set the NTP Server Address on the Cisco Ethernet Switch \(page 13-9\)"](#)
- ["Setting the NTP Server on the Storage Servers \(page 13-10\)"](#)
- ["Restart Recovery Appliance \(page 13-10\)"](#)



Note:

- You must restart Recovery Appliance after completing these tasks.
- You can configure one or two NTP servers for use with Recovery Appliance.

13.3.1 Setting the NTP Server Address on the Compute Servers

To set the NTP server address on a compute server:

1. Log in to a compute server as `root`.
2. Stop the NTP services on the server:

```
service ntpd stop
```
3. Open `/etc/ntp.conf` in a text editor and enter the IP address of the new NTP server.
4. Restart the NTP services on the compute server:

```
service ntpd start
```
5. Repeat this task for the second compute server.

13.3.2 Setting the NTP Server Address on the InfiniBand Switches

 **Note:**

Do not manually edit the files on the InfiniBand switches.

To set the NTP server address on a Sun Datacenter InfiniBand Switch 36:

1. Log in to the switch as the `ilom-admin` user.
2. Set the date, time zone, and Network Time Protocol (NTP), using one of the following methods:
 - The Configuration page of the ILOM graphical interface.
 - The following CLI commands:

```
set /SP/clock timezone=preferred_tz
set /SP/clients/ntp/server/1 address=ntp_ip1
set /SP/clients/ntp/server/2 address=ntp_ip2
set /SP/clock usntpserver=enabled
```

In the preceding commands, *preferred_tz* is the preferred time zone, and *ntp_ip1* and *ntp_ip2* are the NTP server IP addresses. You must configure at least one NTP server; the second one is optional.

3. Repeat this task for the other InfiniBand switches.

13.3.3 Set the NTP Server Address on the Cisco Ethernet Switch

To set the NTP server on the Cisco Ethernet switch:

1. Open a connection to the Ethernet switch using SSH, and log in as `admin`. If SSH is not configured, then use Telnet.
2. Change to enable mode.

```
Switch> enable
```

3. Review the current configuration:

```
Switch# show running-config
```

4. Erase the current NTP server configuration and any invalid NTP addresses. In the following example, the current IP addresses are 10.10.10.1 and 10.8.8.1.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# no ntp server 10.10.10.1
Switch(config)# no ntp server 10.8.8.1
Switch(config)# end
Switch# write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes [OK ]
```

5. Configure up to two NTP servers. The following example uses 10.7.7.1 and 10.9.9.1 as the new IP addresses:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ntp server 10.7.7.1 prefer
Switch(config)# ntp server 10.9.9.1
Switch(config)# end
Switch# write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes [OK ]
```

6. Verify the changes:

```
Switch# show running-config

Building configuration...
...
ntp server 192.168.10.10 prefer
...
```

7. Save the configuration:

```
Switch# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
Compressed configuration from 14343 bytes to 3986 bytes[OK]
```

8. Exit from the session:

```
Switch# exit
```

13.3.4 Setting the NTP Server on the Storage Servers

To set the NTP server on a storage server:

1. Log in to the storage server as the `root` user.
2. Follow steps 1 through 5 of My Oracle Support note 1188080.1, to take the server offline.
3. Shut down the cell services and `ocrvottargetd` service:

```
# cellcli -e alter cell shutdown services all
# service ocrvottargetd stop
```
4. Use the `ipconf` utility to change the NTP settings:

```
# /usr/local/bin/ipconf
```
5. Restart the server:

```
# shutdown -r now
```
6. Follow steps 7 through 9 of My Oracle Support note 1188080.1, to bring the server back online.

13.3.5 Restart Recovery Appliance

After changing the servers and switches, restart Recovery Appliance.



See Also:

["Powering On and Off a Recovery Appliance Rack \(page 12-2\)"](#)

13.4 Changing the Time Zone Settings

This section provides information about changing the time zones on Recovery Appliance after the initial configuration and deployment. You must modify the following components when changing the time zone settings:

- Storage servers
- Compute servers
- Sun Datacenter InfiniBand Switch 36 switches
- Cisco switch

Note:

Before changing the time zone settings, you must stop the cell services on the storage servers and the Oracle Clusterware Services.

The following tasks describe how to change the time zone settings on the components:

- ["Changing the Time Zone Settings on the Storage Servers \(page 13-11\)"](#)
- ["Change Time Zone Settings on the Compute Servers \(page 13-13\)"](#)
- ["Change the Time Zone Settings on the InfiniBand Switches \(page 13-13\)"](#)
- ["Change the Time Zone Settings on the Cisco Switch \(page 13-13\)"](#)

13.4.1 Changing the Time Zone Settings on the Storage Servers

Complete the setting changes to all storage servers before changing the settings on the compute servers.

To change the time zone setting on a storage server:

1. Log in as the `root` user on the storage server.
2. Stop the processes on the cells:

```
# cellcli -e alter cell shutdown services all
```

3. Run the `ipconf` script:

```
# /opt/oracle.cellos/ipconf
```

4. Proceed through the script prompts until you see the time zone prompts. Do not change any other settings.

The following example changes the time zone from McMurdo, Antarctica, to New York City, USA. The number for the United States is 230.

```
The current timezone: Antarctica/McMurdo
Do you want to change it (y/n) [n]: y
```

```
Setting up local time...
```

```
1) Andorra
2) United Arab Emirates
3) Afghanistan
.
.
.
15) Aruba
16) Aaland Islands
Select country by number, [n]ext, [l]ast: 230

Selected country: United States (US). Now choose a zone

1) America/New_York
2) America/Detroit
3) America/Kentucky/Louisville
.
.
.
15) America/North_Dakota/New_Salem
16) America/Denver
Select zone by number, [n]ext: 1

Selected timezone: America/New_York
Is this correct (y/n) [y]:
```

5. Proceed through the remaining prompts, but do not change any other values.
6. Ensure that the time zone changes appear in the following files. Examples of the changes are shown for the files.

- /opt/oracle/cellos/cell.conf

```
$VAR1 = {
    'Hostname' => 'xdserver.us.example.com',
    'Ntp servers' => [
        '10.141.138.1'
    ],
    'Timezone' => 'America/New_York',
```

- /etc/sysconfig/clock

```
ZONE="America/New_York"
UTC=false
ARC=false
#ZONE="Antarctica/McMurdo"
#ZONE="America/New_York"
#ZONE="America/Los_Angeles"
```

The uncommented value (not preceded by #) is the current setting.

- /etc/localtime

Use the `strings /etc/localtime` command to see the printable characters in the file. The last line includes the time zone setting.

```
EST5EDT,M3.2.0,M11.1.0
```

7. Restart the server.
8. Use the `date` command to see the current time zone. The following is an example of the output from the command:

```
# date
Thu Jun 12 16:43:47 EDT 2014
```

9. Review the `$ADR_BASE/diag/asm/cell/host_name/alert.log` file. Verify that the time that you restarted the processes matches the current and correct time.

13.4.2 Change Time Zone Settings on the Compute Servers

To change the time zone setting on a compute server:

1. Log in as the `root` user on the compute server.

2. Shut down Oracle Clusterware Services:

```
# GI/bin/crsctl stop crs
```

3. Copy the `clock` file from one of the storage servers:

```
# scp root@storage_cell:/etc/sysconfig/clock /etc/sysconfig/clock
```

4. Disable Cluster Ready Services (CRS):

```
# GI/bin/crsctl disable crs
```

This step prevents Cluster Ready Services (CRS) from restarting automatically when you restart the server.

5. Restart the compute server.
6. Use the `date` command verify the change for the time zone.
7. Enable CRS:

```
# GI/bin/crsctl enable crs
```

This step restores CRS to its previous configuration setting.

8. Start CRS:

```
# GI/bin/crsctl start crs
```

13.4.3 Change the Time Zone Settings on the InfiniBand Switches

The following procedure describes how to change the time zone setting on a Sun Datacenter InfiniBand Switch 36:

1. Connect to the switch using SSH, and log in as the `root` user.
2. Log in to Oracle ILOM using the web address:

```
http://switch_alias
```
3. Select the **Configuration** tab.
4. Select the **Clock** tab.
5. Ensure the Synchronize Time Using NTP field is enabled.
6. Enter the correct IP address for the servers.
7. Click **Save**.

13.4.4 Change the Time Zone Settings on the Cisco Switch

To change the time zone setting on the Cisco switch:

1. Use Telnet to connect to the Cisco switch as the `root` user.

2. Use the `enable` command to enter privileged mode.
3. Use the `configure terminal` command to begin configuration.
4. Set the clock using the following command:

```
clock timezone zone hours_offset [minutes_offset]
```

In the preceding command, *zone* is the name of the time zone to be displayed when standard time is in effect, *hours_offset* is the hours offset from UTC, and *minutes_offset* are the minutes offset from UTC. The default time zone is UTC.

To set summer time (daylight savings time), use the following command:

```
clock summer-time zone recurring [week day month hh:mm week day month \
hh:mm [offset]]
```

In the preceding command, *week* is the week of the month, between 1 to 5, *day* is the day of the week, such as Sunday or Monday, *month* is the month, such as January or June, *hh:mm* is the time in 24-hour format, and *offset* is the number of minutes to add during summer time. The default for *hh:mm* is 60. Summer time is disabled by default.

The following is an example of setting the time zone to US Eastern time with summer time enabled:

```
$ telnet raeth-ip
Connected to raeth-ip
Escape Character is 'CTRL+]'

User Access Verification

Password: password
raeth-ip> enable
Password: password
raeth-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
raeth-ip(config)# clock timezone EST -5
raeth-ip(config)# clock summer-time EDT recurring
raeth-ip(config)# end
raeth-ip# write memory
Building configuration...
Compressed configuration from 6421 bytes to 2041 bytes[OK]
raeth-ip# show clock
12:03:43.516 EDT Wed May 12 2012
```

Part IV

Appendixes

Zero Data Loss Recovery Appliance Owner's Guide appendixes are:

- [RACLI Overview](#) (page A-1)
- [Default User Accounts](#) (page B-1)
- [Error Messages](#) (page C-1)
- [Replacement Units](#) (page D-1)
- [In-Rack Cabling Tables for Recovery Appliance X7, X6, and X5](#) (page E-1)
- [In-Rack Cabling Tables for Recovery Appliance X4](#) (page F-1)
- [Multirack Cabling Tables](#) (page G-1)

A

RACLI Overview

This topic provides a summary of all the RACLI commands.

NAME

racli - Recovery Appliance Command Line Interface

DESCRIPTION

The racli is a simplified and controlled environment for users to manage their Recovery Appliance via the command line.

SYNOPSIS

racli <COMMAND> <OBJECT>

HELP

Displays information about racli commands and their objects.

Usage: racli help [<COMMAND> [<OBJECT>]]

Options:

`racli help`

General usage information for the racli help command.

`racli help <COMMAND>`

Detailed information specific to the COMMAND requested.

`racli help <COMMAND> <OBJECT>`

Specific detailed information on the COMMAND and OBJECT relationship.

Command	Description
racli add (page A-4)	Adds an object to the Recovery Appliance.
racli add acfs (page A-4)	Specifies that the Automatic Storage Management Cluster (ACFS) should be added to the Recovery Appliance schema.
racli add cloud_storage (page A-4)	Adds a cloud storage object to the Recovery Appliance.
racli add credential (page A-5)	Adds a user credential to the wallet of the Recovery Appliance.
racli add credential_wallet (page A-6)	Adds a wallet object to the Recovery Appliance.
racli add keystore (page A-7)	Adds a keystore to the Recovery Appliance.
racli add nfs (page A-7)	Specifies that the Network File System (NFS) should be added to the Recovery Appliance schema.
racli add storage_cells (page A-8)	Prepares new Recovery Appliance storage cells for storage location expansion.

Command	Description
racli add storage_disks (page A-8)	Adds disks with the newly added storage cells to the ASM disk groups for storage location expansion.
racli add vpc_user (page A-9)	Creates a new Virtual Private Catalog (VPC) user account in the Recovery Appliance schema.
racli alter (page A-10)	Modifies the properties on objects supported in the Recovery Appliance schema.
racli alter cloud_storage (page A-10)	Modifies the properties of cloud storage supported on the Recovery Appliance.
racli alter credential_wallet (page A-11)	Changes the password on a wallet object in the Recovery Appliance.
racli alter keystore (page A-11)	Creates a TDE master key for the Recovery Appliance, generates new keys, or changes the password for the keystore on the Recovery Appliance.
racli alter passphrase (page A-12)	Modifies password used to access the Recovery Appliance.
racli alter rasys_user (page A-12)	Modifies the Recovery Appliance rasys user account password and management of the password.
racli alter vpc_user (page A-13)	Modifies an existing Recovery Appliance user account and changes its password.
racli disable (page A-13)	Disables an object from operations performed by the Recovery Appliance.
racli disable cloud_storage (page A-13)	Disables a cloud storage object from operations performed by the Recovery Appliance.
racli disable keystore (page A-14)	Disables the keystore.
racli disable nfs (page A-14)	Disables the specified Network File System (NFS).
racli enable (page A-14)	Enables an object for operations performed by the Recovery Appliance.
racli enable cloud_storage (page A-14)	Enables a cloud storage object for operations performed by the Recovery Appliance.
racli enable keystore (page A-15)	Opens the keystore so that it can be used.
racli enable nfs (page A-15)	Enables the specified Network File System (NFS).
racli export (page A-16)	Used for exporting data.
racli export appliance (page A-16)	Creates an export bundle for the Recovery Appliance.
racli help (page A-16)	Refer to the HELP section above.
racli install (page A-16)	Installs software.
racli install appliance (page A-16)	Installs software to the Recovery Appliance software.
racli install okv_endpoint (page A-17)	Installs the software required for Oracle Key Vault (OKV) end points on the Recovery Appliance.
racli list (page A-18)	Lists information about objects given in the Recovery Appliance schema.

Command	Description
racli list cloud_storage (page A-18)	Lists details about the cloud storage known to the Recovery Appliance.
racli list nfs (page A-18)	Lists details about one or more Network File System mounts.
racli list okv_endpoint (page A-19)	Lists details about the OKV end points known to the Recovery Appliance.
racli patch (page A-19)	The action to patch a specified object.
racli patch appliance (page A-19)	Patch the Recovery Appliance Software.
racli remove (page A-20)	Removes objects from the Recovery Appliance schema.
racli remove acfs (page A-20)	Removes the specified Automatic Storage Management Cluster (ACFS) mount from the Recovery Appliance schema.
racli remove cloud_storage (page A-20)	Removes the specified cloud storage from the Recovery Appliance.
racli remove keystore (page A-21)	Removes the specified keystore from the Recovery Appliance schema.
racli remove nfs (page A-21)	Removes the specified Network File System (NFS) mount from the Recovery Appliance schema.
racli remove vpc_user (page A-21)	Removes the specified user account from the Recovery Appliance metadata schema.
racli start (page A-22)	The action to start a specified service (object).
racli start nfs (page A-22)	Start NFS services of the Recovery Appliance.
racli status (page A-22)	Check status of objects associated with the Recovery Appliance.
racli status appliance (page A-22)	Returns a summary of the status of the appliance.
racli status cloud_storage (page A-22)	Returns the status of cloud storage.
racli status disk_group (page A-23)	Verifies that the rebalance operation completed after storage cells or storage disks are added to the ASM disk group.
racli status keystore (page A-23)	Returns the status of the keystore in use by the Recovery Appliance.
racli status nfs (page A-23)	Returns the status of the Network File Service (NFS).
racli status okv_endpoint (page A-23)	Returns the status of the OKV_ENDPOINTs associated with the Recovery Appliance.
racli stop (page A-24)	The action to stop a specified service (object).
racli stop nfs (page A-24)	Stop NFS services of the Recovery Appliance.
racli uninstall (page A-24)	Backs out a portion of the installation.
racli uninstall appliance (page A-24)	Backout a step during install.
racli uninstall okv_endpoint (page A-25)	Uninstalls the specified OKV end point.
racli update (page A-26)	Update a component of the Recovery Appliance.
racli update storage_location (page A-26)	Update RA pool data file with the most current number of storage cells.
racli update support_tools (page A-26)	Updates the TFA and Exachk versions using the staged binaries zip files.

Command	Description
racli upgrade (page A-26)	The action to upgrade a specified object to the next version.
racli upgrade appliance (page A-26)	Upgrade the Recovery Appliance Software Versions.
racli version (page A-27)	Provides version information for software running on the Recovery Appliance.

A.1 racli add

Adds an object to the Recovery Appliance.

Options

A.2 racli add acfs

Specifies that the Automatic Storage Management Cluster (ACFS) should be added to the Recovery Appliance schema.

Syntax

```
racli add acfs --name=myAcfsName --disk_group=myDGName --volume_name=myVolumeName --
mount_point=myPath --size=SIZE
```

Options

--name

Name to be given to the ACFS mount object

--disk_group

Name of the ASM disk group to use for this ACFS mount.

--volume_name

Name of the ASM VOLUME to mount the ACFS on.

--mount_point

Path to mount the ACFS object.

--size

Size of the ACFS object.

A.3 racli add cloud_storage

Adds a cloud storage object to the Recovery Appliance.

Syntax

```
racli add cloud_storage --opc_id=OPC_USER_ID --opc_url='https://
OPC_STORAGE_LOCATION' --container_name='STORAGE_CONTAINER_NAME' --
proxy_host=PROXY_NAME --proxy_port=PROXYPORTNUM --storage_class=[Tiering | Standard]
--streams=NUM_STREAMS
```

Options

--opc_id

The Oracle Private Cloud (OPC) user identification included with the credentials associated with the cloud archive service.

--opc_url

The OPC storage location included with the credentials associated with the cloud archive service.

--container_name

Name for new container within the OPC storage location that will house the copy-to-cloud backups. This container name will also serve as the `sbt_library` name in the `sbt_job_template` creation.

--proxy_host

The name of the proxy host.

--proxy_port

The port number for the proxy.

--storage_class

The type or class of storage to be designated to the newly created container. *Standard* uses the Object Storage only; *Tiering* uses the Archive Storage. For archive to cloud support 'Tiering' class should be used.

--streams

The number of streams to use. If not specified, two (2) streams are created by default.

Details

This is equivalent to the `sbt_library` used for copy-to-tape but is specific to copy-to-cloud. The user's cloud credentials are required.

Examples

This example adds a new cloud storage object named "tango" to the Recovery Appliance schema.

```
racli add cloud_storage --opc_id=user@oracle.com --opc_url='https://
storage.oraclecorp.com' --container_name='tango' --proxy_host='myHost' --
proxy_port=8080 --storage_class=Tiering
```

When prompted, enter the password for `user@oracle.com`.

A.4 racli add credential

Adds a user credential to the wallet of the Recovery Appliance.

Syntax

```
racli add credential --user_name=OS_USER_ID --alias=ALIAS_ID --
wallet_path=PATH_TO_WALLET
```

Options

--user_name

The OS user name associated with the credential to be added.

--alias

The alias you want to associate to the OS username that corresponds to the credential to be added.

--wallet_path

The path to the wallet where you want to store the credential.

Examples

This example adds a credential for OS user *charlie* with an alias of *sierra* into the designated wallet.

```
racli add credential --wallet_path=/raacfs/raadmin/cloud/test --user_name=charlie
--alias=sierra
```

When prompted, enter the password for the OS user for which the credential is to be added.

```
Enter Password [charlie]:
Confirm Password [charlie]:
```

When prompted, enter the password for the wallet the credential is to be added.

```
Enter Wallet Password:
Confirm Wallet Password:
```

A.5 racli add credential_wallet

Adds a wallet object to the Recovery Appliance.

Syntax

```
racli add credential_wallet --wallet_path=PATH_TO_WALLET
```

Options

--wallet_path

Path the credential wallet will reside.

Details

You are prompted to enter new passwords for the keystore and then the wallet. The credentials to access the Recovery Appliance encryption keystore are saved in this wallet.

If no wallet path is specified, the default path is used `/raacfs/raadmin/ra_wallet/`.

Examples

This example adds a credential wallet in the default path.

```
racli add credential_wallet
```

When prompted, enter the password for the new credential wallet.

```
Enter Wallet Password:  
Confirm Wallet Password:
```

A.6 racli add keystore

Adds a keystore to the Recovery Appliance.

Syntax

```
racli add keystore --type=[HSM|SW] --wrl=PATH_TO_WALLET
```

Options

--type

Specifies the type of keystore to be created. If not specified, default is HSM.

--wrl

Specifies the path to the wallet that will be used for the keystore.

Details

The keystore contains one or more TDE master keys for each Recovery Appliance client database, plus the Recovery Appliance's TDE master key. The per-client TDE master keys are used to encrypt backup pieces that are copied to the cloud.

ATTENTION: The Recovery Appliance database is restarted to activate the keystore; plan for a short outage.

Examples

This example adds a keystore of type HSM to the Recovery Appliance.

```
racli add keystore --type=hsm
```

A.7 racli add nfs

Specifies that the Network File System (NFS) should be added to the Recovery Appliance schema.

Syntax

```
racli add nfs --nfs_server=myServerName --nfs_export=myExportPath --  
name=myNfsMountName --mount_options='NFS_OPTIONS'
```

Options

--name

Name to be given to the NFS mount object.

--nfs_server

Server name where the `myExportPath` resides.

```
--nfs_export  
Path on the myServerName intending to mount.  
  
--mount_options  
NFS mounting options. For example: 'rw,tcp,vers=3,timeo=600,actimeo=0'
```

A.8 racli add storage_cells

Prepares new Recovery Appliance storage cells for storage location expansion.

Syntax

```
racli add storage_cells --cells=cell1,cell2,...celln
```

Options

```
--cells  
The comma-separated list of storage cell names to prepare.
```

Details

- The new storage cells must be the same Exadata image version as the existing storage cells.
- The new storage cells must not have any associated databases. To verify if this is the case, run:

```
cellcli -e list database
```

Examples

This example adds the storage cells named *cellhost1* and *cellhost2* to the Recovery Appliance.

```
racli add storage_cells --cell=cellhost1,cellhost2
```

A.9 racli add storage_disks

Adds disks with the newly added storage cells to the ASM disk groups for storage location expansion.

Syntax

```
racli add storage_disks --cell=cell1,cell2,...celln --priority=high
```

Options

```
--cell  
The comma-separated list of names of the newly added storage cells for the new disk groups.
```

```
--priority  
Indicates the priority level for the rebalance. The default is low. Low = 4. High = 32.  
This command should only be ran after rebalance is complete.
```

Details

You must have previously run `racli add storage_cells` to prepare the cells before adding the storage cells to the existing disk groups.

Run this command only if prompted to do so while running the `racli add storage_cells` command.

You can execute this command while the system is running, but the operation can take longer to complete and can negatively impact performance. Avoid running other I/O intensive tasks such as patching or maintenance during this time.

Examples

The following example adds storage cells named `cellhost1` and `cellhost2` to the storage disks.

```
racli add storage_disks --cell=cellhost1,cellhost2
```

A.10 racli add vpc_user

Creates a new Virtual Private Catalog (VPC) user account in the Recovery Appliance schema.

Syntax

```
racli add vpc_user --user_name=myVpcUserName
```

Options

user_name

The name of the Recovery Appliance catalog user account to add.

Details

Each Recovery Appliance user account is the owner of the virtual private catalog.

- The password is not provided on the command line for security reasons. When prompted, enter a password for the Recovery Appliance user account. The password must contain eight or more characters. Strong password authentication is used, so simple passwords are not allowed.
- The Recovery Appliance user account being added must not already exist in the metadata database.

Because strong password authentication is used, simple passwords can result in ORA-280003 or ORA-20006 errors.

When a `vpc_user` is created, the account is locked. The `DEFAULT` database profile has the following limits.

- Failed attempts is set to 10.
- Password life time is 180 days.
- Password lock time is 1 day.
- Password grace time is 7 days.

Examples

This example adds the user name *ravpc1* to the Recovery Appliance schema.

```
racli add vpc_user --user_name=ravpc1
```

When prompted, enter the password for *ravpc1*.

```
ravpc1 New Password:
```

A.11 racli alter

Modifies the properties on objects supported in the Recovery Appliance schema.

Options

A.12 racli alter cloud_storage

Modifies the properties of cloud storage supported on the Recovery Appliance.

Syntax

```
racli alter cloud_storage --opc_id=OPC_USER_ID --opc_url='https://  
OPC_STORAGE_LOCATION' --container_name='STORAGE_CONTAINER_NAME' --  
storage_class=Tiering
```

Options

--opc_id

The OPC User ID included with the credentials associated with the cloud archive service.

--opc_url

The OPC Storage Location included with the credentials associated with the cloud archive service.

--container_name

Name for new container within the OPC Storage Location that will house the archive to cloud backups. This *container_name* will also serve as the *sbt_library* name in the *sbt_job_template* creation.

--storage_class

The type or class of storage to be designated to the newly created container. For archive to cloud support 'Tiering' class should be used.

Details

This is equivalent to the *sbt_library* used for copy-to-tape. This *sbt_library* is specific to copy-to-cloud.

A.13 racli alter credential_wallet

Changes the password on a wallet object in the Recovery Appliance.

Syntax

```
racli alter credential_wallet --wallet_path=PATH_TO_WALLET
```

Options

--wallet_path

Path the credential wallet will reside.

Details

If no wallet_path is specified the default path is used /raacfs/raadmin/ra_wallet/.

Examples

This example alters a credential wallet in the default path

```
racli add credential_wallet
```

When prompted, enter the password for the new credential wallet.

```
Enter Wallet Password:  
Re Enter Password:
```

A.14 racli alter keystore

Creates a TDE master key for the Recovery Appliance, generates new keys, or changes the password for the keystore on the Recovery Appliance.

Syntax

```
racli alter keystore { --initialize_key | --restart_db | --password | --rekey }
```

Options

--rekey

Generates new TDE master encryption key (rekey) for the protected database.

--password

Specifies a new password for the keystore.

--initialize_key

Generates the initial TDE master key for the keystore.

--restart_db

Restarts the database stack after changes to the keystore to ensure they are honored at service restart.

Examples

The following example restarts the database stack following a keystore change.

```
racli alter keystore --restart_db
```

A.15 racli alter passphrase

Modifies password used to access the Recovery Appliance.

Syntax

```
racli alter passphrase
```

Options

Details

- The password specified for the Recovery Appliance user account must contain eight or more characters. Strong password authentication is used, so simple passwords are not allowed.

A.16 racli alter rasys_user

Modifies the Recovery Appliance rasys user account password and management of the password.

Syntax

```
racli alter rasys_user { --init | --password }
```

Options

Details

--init

Establishes management of the *rasys* user. This will be run for you during patch or upgrade in most cases.

In the event the *rasys* wallet goes away, this can be used to re establish *rasys* password management.

--password

Set a new *rasys* user password

Examples

The following example changes the password for the *ravpc1* user account.

```
racli alter rasys_user --password
```

When prompted as shown, enter the new password.

```
Enter rasys Password:  
Confirm new rasys Password:
```

The password specified for the Recovery Appliance user account must contain eight or more characters. Strong password authentication is used, so simple passwords are not allowed.

A.17 racli alter vpc_user

Modifies an existing Recovery Appliance user account and changes its password.

Syntax

```
racli alter vpc_user --user_name=myVpcUserName
```

Options

Details

--user_name

The Virtual Private Catalog user account for the Recovery Appliance

The `user_name` specified in the command must be an existing Recovery Appliance user account.

The password specified for the Recovery Appliance user account must contain eight or more characters. Strong password authentication is used.

Examples

The following example changes the password for the `ravpc1` user account.

```
racli alter vpc_user --user_name=ravpc1
```

When prompted as shown, enter the new password.

```
ravpc1 New Password:
```

A.18 racli disable

Disables an object from operations performed by the Recovery Appliance.

Options

Details

Object remains known to the Recovery Appliance for a quick enable later.

A.19 racli disable cloud_storage

Disables a cloud storage object from operations performed by the Recovery Appliance.

Syntax

```
racli disable cloud_storage { --container_name=CONTAINER_NAME | --all }
```

Options

--container_name

Specifies the specific container name associated with the cloud storage to be disabled.

--all

Disables all cloud storage known to the Recovery Appliance.

Examples

The following example disables all of the known cloud storage objects.

```
racli disable cloud_storage --all
```

A.20 racli disable keystore

Disables the keystore.

Syntax

```
racli disable keystore
```

Options

A.21 racli disable nfs

Disables the specified Network File System (NFS).

Syntax

```
racli disable nfs --name=myMountName
```

Options

--name

Name of the NFS mount object.

A.22 racli enable

Enables an object for operations performed by the Recovery Appliance.

Options

A.23 racli enable cloud_storage

Enables a cloud storage object for operations performed by the Recovery Appliance.

Syntax

```
racli enable cloud_storage { --container_name=STORAGE_CONTAINER_NAME | --all }
```

Options

--container_name
Specifies a specific container name associated with the cloud storage to enable.

--all
Enables all cloud storage known to the Recovery Appliance.

Examples

The following example enables cloud storage objects associated with the container named *Lima*.

```
racli enable cloud_storage --container_name=lima
```

A.24 racli enable keystore

Opens the keystore so that it can be used.

Syntax

```
racli enable keystore
```

Options

Details

This is the OKV endpoint keystore, also known as the OKV shared wallet.

A.25 racli enable nfs

Enables the specified Network File System (NFS).

Syntax

```
racli enable nfs --name=myMountName
```

Options

--name
Name of the NFS mount object.

A.26 racli export

Used for exporting data.

Options

A.27 racli export appliance

Creates an export bundle for the Recovery Appliance.

Syntax

```
racli export appliance
```

Options

A.28 racli help

Refer to the HELP section above.

Options

A.29 racli install

Installs software.

Options

A.30 racli install appliance

Installs software to the Recovery Appliance software.

Syntax

```
racli install appliance --step={ STEP_NUMBER | ALL } [--import=IMPORT_BUNDLE_PATH]
```

Options

--import

Specifies the location of the RA export bundle to use for the installation.

--step

Indicate which step number of install appliance to run. Value range is 1-7 or "all" for all steps.

Step Numbers:

1. Validates and configures the hardware and platform.
2. Configures the operating system, including setting up file system properties, and users and groups.

3. Configures the `oracle` operating system user.
4. Creates the database file system (DBFS) mounts, configures TNS aliases, and sets up the required wallets for communication.
5. If a tape backup system is detected, validates and configures the appliance using details that were specified on the Tape Library Connectivity Page in the Oracle Exadata Deployment Assistant.
6. Configures backup of the Recovery Appliance catalog database and creates credentials for performing the backups. The catalog is backed up to a Recovery Appliance disk on an automated schedule.
7. Starts all Recovery Appliance services installed in the previous steps.

Details

Before you run `racli install appliance`:

- Ensure that all current installation patches are staged in the `/u01/ra_install/linux-x64/WorkDir` directory.
- Ensure that the `ra_preinstall.pl` and `install.sh` scripts were run and successfully completed.

If a step fails, contact Oracle Support Services.

Examples

The following example runs only step 2 of the software installation.

```
racli install appliance --step=2
```

The following example runs all of the steps necessary to install the software based on the system and configuration settings.

```
racli install appliance
```

A.31 racli install okv_endpoint

Installs the software required for Oracle Key Vault (OKV) end points on the Recovery Appliance.

Syntax

```
racli install okv_endpoint
```

Options

Details

This is run after the a keystore has been activated and is only run once. It adds all OKV endpoints associated with the Recovery Appliance; it applies to all nodes.

Examples

This example installs the `okv_endpoints` associated with each compute node of the Recovery Appliance.

```
racli install okv_endpoint
```

```
Enter New Keystore Password:
Confirm New Keystore Password:
Enter RA Wallet Password:
Wed December 23 20:14:40 2018: Start: Install OKV End Point [compute_node1]
Wed December 23 20:14:43 2018: End: Install OKV End Point [compute_node1]
Wed December 23 20:14:43 2018: Start: Install OKV End Point [compute_node2]
Wed December 23 20:14:45 2018: End: Install OKV End Point [compute_node2]
```

A.32 racli list

Lists information about objects given in the Recovery Appliance schema.

Options

A.33 racli list cloud_storage

Lists details about the cloud storage known to the Recovery Appliance.

Syntax

```
racli list cloud_storage { --container_name=STORAGE_CONTAINER_NAME | --all } [--long]
```

Options

--container_name

Specifies a container name associated with the cloud storage to list details.

--all

Lists details about all cloud storage known to the Recovery Appliance.

--long

Lists extended details about one or all containers.

Examples

The following example lists all known cloud_storage objects.

```
racli list cloud_storage --all
```

A.34 racli list nfs

Lists details about one or more Network File System mounts.

Syntax

```
racli list nfs { --mount_name=myName | --all }
```

Options

--mount_name

Name of NFS mount to be listed.

--all

Specifies that all NFS mounts should be listed.

A.35 racli list okv_endpoint

Lists details about the OKV end points known to the Recovery Appliance.

Syntax

```
racli list okv_endpoint [--all] [--long]
```

Options

--all

Lists all of the okv_endpoint(s) known to the Recovery Appliance.

--long

Lists extended details about the OKV end points known to the Recovery Appliance.

Examples

The following example lists all of the OKV end points with extended details.

```
racli list okv_endpoint --all --long
```

A.36 racli patch

The action to patch a specified object.

Options

A.37 racli patch appliance

Patch the Recovery Appliance Software.

Syntax

```
racli patch appliance --step=STEP_NUMBER [--ignore_incidents]
```

Options

--step

Indicate which step number of upgrade appliance to run. Value range is 1-2.

Step Numbers:

1. Runs the prechecks prior to upgrade. No downtime required. Does not change the software.
2. Performs the Patch. A full outage is required and is followed by a postcheck.

--ignore_incidents

Ignore active incidents (not recommended).

A.38 racli remove

Removes objects from the Recovery Appliance schema.

Options

Details

Removed objects become unknown to the Recovery Appliance and cannot be quickly enabled later.

A.39 racli remove acfs

Removes the specified Automatic Storage Management Cluster (ACFS) mount from the Recovery Appliance schema.

Syntax

```
racli remove acfs --mount_name=myName
```

Options

--mount_name

Name of ACFS mount to be removed.

A.40 racli remove cloud_storage

Removes the specified cloud storage from the Recovery Appliance.

Syntax

```
racli remove cloud_storage --container_name=STORAGE_CONTAINER_NAME [--force]
```

Options

--container_name

Specifies the <STORAGE_CONTAINER_NAME> that correlates to the cloud_storage to be removed.

--force

Removes the cloud_storage regardless of the state/contents of the underlying container.

Examples

This example removes the cloud_storage object associated with the storage container name gorilla regardless of container contents.

```
racli remove cloud_storage --container_name=gorilla --force
```

A.41 racli remove keystore

Removes the specified keystore from the Recovery Appliance schema.

Syntax

```
racli remove keystore
```

Options

A.42 racli remove nfs

Removes the specified Network File System (NFS) mount from the Recovery Appliance schema.

Syntax

```
racli remove nfs --mount_name=myName
```

Options

--mount_name

Name of NFS mount to be removed.

A.43 racli remove vpc_user

Removes the specified user account from the Recovery Appliance metadata schema.

Syntax

```
racli remove vpc_user --user_name=myVPCUserName
```

Options

--user_name

The Virtual Private Catalog user for the Recovery Appliance.

Details

- The specified user account must exist in the Recovery Appliance schema.
- The specified user account must not be associated with a protected database.

Examples

The following example removes the user account `ravpc1` from the metadata repository.

```
racli remove vpc_user --user_name=ravpc1
```

A.44 racli start

The action to start a specified service (object).

Options

A.45 racli start nfs

Start NFS services of the Recovery Appliance.

Syntax

```
racli start nfs --name=myMountName
```

Options

--name

Name of the NFS mount to start.

A.46 racli status

Check status of objects associated with the Recovery Appliance.

Options

A.47 racli status appliance

Returns a summary of the status of the appliance.

Syntax

```
racli status appliance
```

Options

A.48 racli status cloud_storage

Returns the status of cloud storage.

Syntax

```
racli status cloud_storage [--long]
```

Options

--long

Lists extended details about cloud storage known to the Recovery Appliance.

A.49 racli status disk_group

Verifies that the rebalance operation completed after storage cells or storage disks are added to the ASM disk group.

Syntax

```
racli status disk_group
```

Options

Details

Run this command before running the `racli update storage_location` command to ensure that the ASM rebalance completed.

All rebalance must complete before Griddisks and Storage Spaces are added to a Recovery Appliance.

A.50 racli status keystore

Returns the status of the keystore in use by the Recovery Appliance.

Syntax

```
racli status keystore
```

Options

A.51 racli status nfs

Returns the status of the Network File Service (NFS).

Syntax

```
racli status nfs --name=myMountName
```

Options

A.52 racli status okv_endpoint

Returns the status of the OKV_ENDPOINTs associated with the Recovery Appliance.

Syntax

```
racli status okv_endpoint
```

Options

A.53 racli stop

The action to stop a specified service (object).

Options

A.54 racli stop nfs

Stop NFS services of the Recovery Appliance.

Syntax

```
racli stop nfs --name=myMountName
```

Options**Details****--name**

Name of the ACFS mount to start.

A.55 racli uninstall

Backs out a portion of the installation.

Options

A.56 racli uninstall appliance

Backout a step during install.

Syntax

```
racli uninstall appliance --step={ STEP_NUMBER | ALL }
```

Options**Details****--step**

Indicate which step number of uninstall appliance to run. Value range is 1-7 or "all" for all steps.

Step Numbers:

1. Unconfigures the hardware and platform.
2. Unconfigures the operating system, including setting up file system properties, and users and groups.
3. Unconfigures the Oracle operating system user.

4. Removes file system mounts, unconfigures TNS aliases, and invalidates wallets for communication.
5. If a tape backup system is detected, unconfigures the resources used for tape backup.
6. Unconfigures backup of the Recovery Appliance catalog database and removes credentials for performing the backups.
7. Stops all Recovery Appliance services installed.

Generally used if you found an error and have to redo a step after correcting. You must run the uninstallation in the reverse numerical order.

 **Note:**

If a failed installation step is the reason that you plan to run the uninstall command, it is highly recommended that you contact Oracle Support Services.

Examples

The following example undoes only step 3 of the installation procedure.

```
racli uninstall appliance --step=3
```

The following undoes all of the steps of the installation procedure in reverse sequence.

```
racli uninstall appliance
```

A.57 racli uninstall okv_endpoint

Uninstalls the specified OKV end point.

Syntax

```
racli uninstall okv_endpoint [--force]
```

Options

--force

When specified, forces the uninstallation of the OKV end point.

Otherwise, a check is done for active storage locations utilizing the OKV end point. If any are in use, a warning is issued and uninstall fails.

Examples

This example forces uninstallation of the OKV end points associated with the Recovery Appliance.

```
racli uninstall okv_endpoint --force
```

```
Thu August 17 06:03:55 2018: Start: Remove OKV End Point  
Thu August 17 06:03:55 2018:
```

```
WARNING:
  Found active cloud locations.
  Removing the OKV endpoint will mean this Recovery Appliance can't recover
  backups at cloud locations without an endpoint configured.Thu August 17
06:03:55 2018:   Start: Remove OKV End Point [node1]
Thu August 17 06:03:55 2018:   End: Remove OKV End Point [node1]
Thu August 17 06:03:55 2018:   Start: Remove OKV End Point [node2]
Thu August 17 06:03:55 2018:   End: Remove OKV End Point [node2]
Thu August 17 06:03:55 2018:   End: Remove OKV End Point
```

A.58 racli update

Update a component of the Recovery Appliance.

Options

A.59 racli update storage_location

Update RA pool data file with the most current number of storage cells.

Syntax

```
racli update storage_location --cells=cell1,cell2,...celln
```

Options

A.60 racli update support_tools

Updates the TFA and Exachk versions using the staged binaries zip files.

Syntax

```
racli update support_tools
```

Options

A.61 racli upgrade

The action to upgrade a specified object to the next version.

Options

A.62 racli upgrade appliance

Upgrade the Recovery Appliance Software Versions.

Syntax

```
racli upgrade appliance --step=STEP_NUMBER [--ignore_incidents]
```

Options

--step

Indicate which step number of upgrade appliance to run. Value range is 1-4.

Step Numbers:

1. Runs the prechecks prior to upgrade. No downtime required. Does not change the software.
2. Stages the new software stack. No downtime required. Does not change the software.
3. Performs the upgrade using the staged items. A full outage is required; it is followed by a postcheck and cleanup of the old files.
4. Migrate DBFS to ACFS. All files residing on DBFS will be migrated to new ACFS.

--ignore_incidents

Ignore active incidents (not recommended).

A.63 racli version

Provides version information for software running on the Recovery Appliance.

Syntax

```
racli version
```

Options

Details

The versioning information provided includes:

rarpm version

The ra_automation rpm currently installed on the Recovery Appliance.

rdbms version

The database version currently installed on the Recovery Appliance.

transaction

The specific transaction associated with the version running on the Recovery Appliance.

zdlra version

The base software version installed and running on the Recovery Appliance.

B

Default User Accounts

The following table lists the default users and passwords for the Zero Data Loss Recovery Appliance components. All default passwords should be changed after installation of the Recovery Appliance.

Table B-1 Default Users and Passwords

Component	User Name and Password
Compute servers	<p>OS users:</p> <ul style="list-style-type: none"> • root/welcome1 • oracle/Welcome\$ • dbmadmin/welcome • dbmmonitor/welcome • raext/(locked and blocked from SSH access) • railm/(locked and blocked from SSH access) • Password for the GRUB boot loader: sos1Exadata <p>Database users:</p> <ul style="list-style-type: none"> • SYS/Welcome\$ • SYSTEM/Welcome\$ • raext/(externally authenticated) • ralim/(externally authenticated) • rasys/change^Me2 <p>OSB tape backup application users:</p> <ul style="list-style-type: none"> • admin/welcome1 • oracle/welcome1 • encryption key wallet/welcome1
Storage servers	<ul style="list-style-type: none"> • root/welcome1 • celladmin/welcome • cellmonitor/welcome • CELLDIAG <p>CELLDIAG is an Exadata storage software user, not an operating system user.</p> <p>The password of the CELLDIAG user is reset to a random password during the "Apply Security Fixes" step of Oracle Exadata Deployment Assistant. If this step is not run, then the default password is Welcome12345.</p> <ul style="list-style-type: none"> • Password for the GRUB boot loader: sos1Exadata
InfiniBand switches	<ul style="list-style-type: none"> • root/welcome1 • nm2user/changeme • ilom-admin/ilom-admin • ilom-operator/ilom-operator
Ethernet switches	<p>admin/welcome1</p> <p>Note: Secure the <code>enable mode password</code> and <code>secret</code> values for the admin user.</p>

Table B-1 (Cont.) Default Users and Passwords

Component	User Name and Password
Power distribution units (PDUs)	<ul style="list-style-type: none"> admin/welcome1 <p>The password for the admin user is admin if you reset the PDU to factory default settings.</p>
Compute server ILOMs	<ul style="list-style-type: none"> root/welcome1 MSUser <p>Management Server (MS) uses this account to manage ILOM and reset it if it detects a hang.</p> <p>Do not modify this account. This account is to be used by MS only.</p> <p>Each time MS starts up, it deletes the previous MSUser account and re-creates the account with a randomly generated password.</p> <p>MS communicates with ILOM using MSUser through the lanplus interface. It uses the IPMI v2.0 RMCP+ protocol for authentication. RMCP+ (remote management control protocol) is a UDP-based protocol with stronger authentication than RMCP.</p> <p>The MSUser password is not persisted anywhere. If you need to change account passwords regularly, you can restart MS to change the password of the MSUser account.</p>
Storage server ILOMs	<ul style="list-style-type: none"> root/welcome1 MSUser <p>See the description above for details about this user.</p>
InfiniBand ILOMs	<ul style="list-style-type: none"> ilom-admin/ilom-admin ilom-operator/ilom-operator root/welcome1

 **Note:**

After the Recovery Appliance has been deployed, the installation process disables all root SSH keys and expires all user passwords as a security measure for your system. If you do not want the SSH keys disabled or the passwords expired, advise the installation engineer before the deployment.

 **See Also:**

["Changing Component Passwords"](#) (page 13-1) to learn how to change the passwords for the Recovery Appliance components.

C

Error Messages

This appendix contains the hardware messages for Recovery Appliance. They are identical to the messages for Oracle Exadata Database Machine.

This appendix contains the following sections:

- [Contacting Oracle Support Services](#) (page C-1)
- [Generic Alert Messages: HALRT-00001 to HALRT-00500](#) (page C-2)
- [Temperature Alert Messages: HALRT-00513 to HALRT-00600](#) (page C-5)
- [Fan Alert Messages: HALRT-00616 to HALRT-00700](#) (page C-6)
- [Power Supply/Converters Alert Messages: HALRT-00717 to HALRT-00800](#) (page C-6)
- [Network Alert Messages: HALRT-00912 to HALRT-01000](#) (page C-7)
- [Hard Disk and Flash Disk Alert Messages: HALRT-01015 to HALRT-01200](#) (page C-8)
- [Accelerator/Controller Alert Messages: HALRT-01209 to HALRT-01300](#) (page C-28)
- [Voltage Alert Messages: HALRT-01301 to HALRT-01400](#) (page C-32)
- [Physical Disk Alert Messages: HALRT-01401 to HALRT-01500](#) (page C-32)
- [Auto Service Request Alert Messages for Storage Servers: HALRT-02001 to HALRT-02013](#) (page C-34)
- [HDD Controller Battery Alert Messages: HALRT-02101 to HALRT-02200](#) (page C-36)
- [Generic Actions for Alert Messages: HALRT-10001 to HALRT-10500](#) (page C-39)
- [Generic Actions for Alert Messages: HALRT-10503 to HALRT-11000](#) (page C-44)
- [Generic Actions for Alert Messages: HALRT-11001 to HALRT-11100](#) (page C-45)

See Also:

Oracle Database Error Messages for information about Oracle Database SNMP error messages

C.1 Contacting Oracle Support Services

Some messages recommend contacting Oracle Support Services to report a problem. To report a problem, sign in to My Oracle Support at <http://support.oracle.com> and log a service request. You might also want to contact Oracle Support when Auto Service Request submits a service request.

When you contact Oracle Support Services, have the following information available:

- The hardware, operating system, and release number of the operating system running Oracle Database.
- The complete release number of Oracle Database, such as release 12.1.0.2.
- All Oracle programs (with release numbers) in use when the error occurred.
- If you encountered one or more error codes or messages, then the exact code numbers and message text, in the order in which they appeared.
- The problem severity, according to the following codes:
 - 1: Program not usable. Critical effect on operations.
 - 2: Program usable. Operations severely restricted.
 - 3: Program usable with limited functions. Not critical to overall operations.
 - 4: Problem circumvented by customer. Minimal effect, if any, on operations.

You will also be expected to provide the following:

- The name of your organization
- Your name
- Your Oracle Support ID number
- Your telephone number
- Rack master serial number

See "[Related Documentation](#) (page xxii)" for links to My Oracle Support notes that are related to Recovery Appliance.

C.2 Generic Alert Messages: HALRT-00001 to HALRT-00500

HALRT-00002, Chassis Intrusion - Physical Security Violation

Cause: Detected a Physical Security Violation on the chassis.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00003, Chassis Intrusion (Physical Security Violation) Event Cleared

Cause: Detected a clear event of Physical Security Violation on the chassis.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00004, CPU Internal Error

Cause: Detected an Internal CPU error.

Action: See "[HALRT-10025](#) (page C-40)".

HALRT-00005, CPU Internal Error Cleared

Cause: Detected a clear event for CPU internal error.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00006, CPU Configuration Error

Cause: Detected a CPU configuration error.

Action: See "[HALRT-10025](#) (page C-40)".

HALRT-00007, CPU Configuration Error Cleared

Cause: Detected a clear event for a CPU configuration error.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00008, CPU Presence (Processor Presence detected)

Cause: Detected presence event for a new CPU.

Action: See "[HALRT-10025](#) (page C-40)".

HALRT-00009, CPU Not Present (Processor Not Present)

Cause: Detected a CPU is not present.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00010, System Event Log (SEL) Full (Logging Disabled)

Cause: System event log detected to be full.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00011, ASR (Automatic System Recovery) Timer Expired

Cause: Automatic System Recovery timer expired.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00012, ASR (Automatic System Recovery) Reset Occurred

Cause: Automatic System Recovery reset occurred.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00013, ASR (Automatic System Recovery) Power Down Occurred

Cause: Automatic System Recovery was powered down.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00014, ASR (Automatic System Recovery) Power Cycle Occurred

Cause: Automatic System Recovery power was cycled.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00015, {0} Sensor Threshold state change: Sensor Component Name: {1} Sensor State: {2} Sensor Threshold Type: {3} Sensor Threshold Severity: {4} Sensor Value: {5} Sensor Threshold Value: {6} Sensor Additional Info: {7}

Cause: The sensor threshold state has changed.

Action: See "[HALRT-10029, A generic sensor has detected an error.](#) (page C-40)".

HALRT-00016, A {0} component is suspected of causing a fault. Component Name: {1} Fault class: {2} Fault certainty: {3} Fault message: {4} Trap Additional Info: {5}

Level Critical

Cause: A component might have caused a fault to occur.

Action: See "[HALRT-10030](#) (page C-40)".

HALRT-00017, A {0} component fault has been cleared. Component Name: {1} Trap Additional Info: {2}

Level Clear

Cause: The component fault has cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00018, A {0} sensor has detected an error. Component Name: {1} Trap Additional Info: {2}

Cause: An error was detected by a sensor.

Action: See "[HALRT-10029, A generic sensor has detected an error.](#) (page C-40)".

HALRT-00019, A {0} sensor has returned to its normal state. Component Name: {1} Trap Additional Info: {2}

Cause: The sensor has returned to its normal state.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00020, A field replaceable unit at {0} has been inserted into the system.

Cause: A unit was inserted into the system.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00021, A field replaceable unit at {0} has been removed from the system.

Cause: A unit was removed from the system.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00022, The component operational state has changed or is changing to disabled. Component Name: {0} Disable Reason: {1} Trap Additional Info: {2}

Cause: The operational state for a component has changed.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00023, The component operational state has changed or is changing to enabled. Component Name: {0} Trap Additional Info: {1}

Cause: The operational state for a component has changed.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00024, An error has been detected in the period before the operating system has taken control of the system.

Cause: An error was detected.

Action: See "[HALRT-10032](#) (page C-40)".

HALRT-00025, The high availability component identified by {0} has changed state. Old HA State: {1} New HA State: {2} Trap Additional Info: {3}

Cause: The state for a high availability component has changed.

Action: See "[HALRT-10030](#) (page C-40)".

HALRT-00026, An intrusion sensor has detected that someone may have physically tampered with the system.

Cause: A sensor has detected a change to the system.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00027, The hard drive has changed state. Component Name: {0} New State: {1}

Cause: The state of the hard drive has changed.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00028, This is a test trap.

Cause: User generated a test trap, and an alert was created.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00030, A {0} component fault has been cleared.

Cause: A component fault was cleared.

Action: This is an informational message, and no action is required.

C.3 Temperature Alert Messages: HALRT-00513 to HALRT-00600

HALRT-00513, Under-Temperature Warning (Lower non-critical, going low).

Cause: The temperature was below the specified range.

Action: See "[HALRT-10015](#) (page C-40)"

HALRT-00514, Under-Temperature Warning Cleared.

Cause: The low temperature warning was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00515, Under-Temperature Problem (Lower Critical - going low)

Cause: A low temperature problem was detected.

Action: See "[HALRT-10004](#) (page C-39)".

HALRT-00516, Under-Temperature Problem Cleared

Cause: The low temperature problem was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00517, Over-Temperature warning (Upper non-critical, going high)

Cause: The temperature was above the specified range.

Action: "[HALRT-10015](#) (page C-40)".

HALRT-00518, Over-Temperature warning Cleared

Cause: The high temperature warning was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00519, Over-Temperature Problem (Upper Critical - going high)

Cause: A high temperature was detected.

Action: "[HALRT-10004](#) (page C-39)".

HALRT-00520, Over-Temperature Problem Cleared

Cause: The high temperature problem was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00521, CPU Thermal Trip (Over Temperature Shutdown)

Cause: The CPU high temperature shut down was tripped.

Action: See "[HALRT-10004](#) (page C-39)".

HALRT-00522, CPU Thermal Trip (Over Temperature Shutdown) Cleared

Cause: The CPU high temperature shut down was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

C.4 Fan Alert Messages: HALRT-00616 to HALRT-00700

HALRT-00616, Generic Critical Fan Failure

Cause: A fan has failed.

Action: See "[HALRT-10005](#) (page C-40)".

HALRT-00617, Generic Critical Fan Failure Cleared

Cause: A fan failure was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00618, Fan Redundancy Lost

Cause: A redundant fan has stopped functioning.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00619, Fan redundancy Has Returned to Normal

Cause: A redundant fan has started to function normally.

Action: See "[HALRT-10001](#) (page C-39)".

C.5 Power Supply/Converters Alert Messages: HALRT-00717 to HALRT-00800

HALRT-00717, Power Supply Redundancy Lost

Cause: A redundant power supply has stopped functioning.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00718, Power Supply Redundancy Has Returned to Normal

Cause: A redundant power supply has started to function normally.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00719, Power Supply Inserted

Cause: A power supply was inserted into the system.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00720, Power Supply Removed

Cause: A power supply was removed from the system.

Action: "[HALRT-10007](#) (page C-40)".

HALRT-00721, Power Supply Failure

Cause: A power supply has failed.

Action: "[HALRT-10007](#) (page C-40)".

HALRT-00722, Power Supply Failure Cleared

Cause: A power supply failure was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00723, Power Supply Warning

Cause: A power supply problem has set off a warning.

Action: "[HALRT-10007](#) (page C-40)".

HALRT-00724, Power Supply Warning Cleared

Cause: A power supply warning was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00725, Power Supply AC Lost

Cause: A power supply has lost AC power.

Action: Examine the power supply. Replace the power supply if necessary.

HALRT-00726, Power Supply AC Restored

Cause: A power supply has had its AC power restored.

Action: See "[HALRT-10001](#) (page C-39)".

C.6 Network Alert Messages: HALRT-00912 to HALRT-01000

HALRT-00912, Infiniband port {0} indicates invalid state: {1}

Cause: InfiniBand device indicated an invalid state.

Action: See "[HALRT-11001](#) (page C-45)".

HALRT-00913, InfiniBand port {0} indicates invalid rate: {1} Gps

Cause: InfiniBand device indicated invalid rate.

Action: See "[HALRT-11003](#) (page C-46)".

HALRT-00914, InfiniBand port status is OK

Cause: InfiniBand device status was all right.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-00915, InfiniBand port {0} is showing non-zero error counts

Cause: The error counts for `RcvErrors` and `SymbolErrors` on the specified port were non-zero. A non-zero error count can be caused by a bad cable, or a cable that is not properly seated, or by a bad port on the HCA or switch. Manually disabling and enabling InfiniBand ports or resetting InfiniBand switches can also lead to these errors.

Action: See "[HALRT-11003](#) (page C-46)".

C.7 Hard Disk and Flash Disk Alert Messages: HALRT-01015 to HALRT-01200

These alert messages are in the range HALRT-01015 through HALRT-01200.

HALRT-01015, Physical drive inserted on Adapter: {0} Device Id: {1}, Enclosure ID: {2}, Slot ID: {3}

Cause: The physical drive was inserted.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01016, Physical drive removed on Adapter: {0} Device ID: {1}, Enclosure ID: {2}, Slot ID: {3}

Cause: The physical drive was removed.

Action: See "[HALRT-10040](#) (page C-41)".

HALRT-01017, Physical drive state changed from {0} to {1}. Adapter: {2} Device ID: {3}, Enclosure ID: {4}, Slot ID: {5}

Cause: The physical drive status has changed.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01018, FRU changed. Adapter: {0} Device ID: {1}, Enclosure ID: {2}, Slot ID: {3}, Enclosure ID: {4} FRU {5}

Cause: The field replaceable unit (FRU) was changed on the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01020, Physical drive not supported. Adapter: {0} Device ID: {1}, Enclosure ID: {2}, Slot ID: {3}

Cause: The physical drive is not supported on the adapter.

Action: See "[HALRT-10040](#) (page C-41)".

HALRT-01021, Physical drive not certified. Adapter: {0} Device ID: {1}, Enclosure ID: {2}, Slot ID: {3}

Cause: The physical drive was not certified.

Action: See "[HALRT-10040](#) (page C-41)".

HALRT-01024, Physical {0} disk error count is increasing. Disk: {1}, Serial Number: {2}, Slot Number: {3}

Level Critical

Cause: The disk error count increased during the periodic disk checks.

Action: Monitor the physical disk errors and replace the disk necessary.

HALRT-01025, Physical {0} disk error count is stable. Disk: {1}, Serial Number: {2}, Slot Number: {3}

Level Clear

Cause: The disk error count did not increase after the last disk check.

Action: This is an informational message, and no action is required.

HALRT-01027, Flash celldisk {0} write performance is degraded. The supercap module on the Flash card may have failed. Flash disk name : {1} Disk serial number: {2} Slot number: {3}

Cause: The write performance for the flash cell disk was degraded.

Action: Replace the flash disk, flash PCI card, or Energy Storage Module (ESM). If all disks on the PCI card show degraded performance, then check the ESM LED on the PCI card. It is visible from the back of the cell chassis. If this LED is amber, then replace the ESM. If the disks being replaced have persistent data, then rebalance the data out of these disks. Next, replace the flash disk. If persistent data was on the flash disks, then rebalance then data back into the disks. The flash cache portion provided by this flash disk will be disabled automatically.

HALRT-01028, Flash celldisk {0} write performance problem has been resolved

Cause: The performance problem for the flash cell disk was resolved.

Action: This is an informational message, and no action is required.

HALRT-01029, Physical disk {0} could not be set to ONLINE automatically. Physical disk {0} firmware could not be upgraded.

Level Critical

Cause: The physical disk firmware could not be upgraded.

Action: Upgrade the firmware on the physical disks. Restart MS to put the disk online.

HALRT-01032, Physical disk {0} firmware successfully upgraded

Cause: The physical disk firmware was upgraded.

Action: This is an informational message, and no action is required.

HALRT-01035, Physical disks {0} firmware could not be upgraded. Inserted physical disks will be accepted because MS_ACCEPT_DISKFW_DIFFERENT is set to true.

Level Informational

Cause: The firmware upgrade failed.

Action: This is an informational message.

HALRT-01036, Oracle Exadata discovered the following problems in the device utilities: {0}

Cause: A problem has been discovered. The problem is one of the following:

- Output has duplicate name.
- Output has duplicate serial number.
- Output has duplicate OS device name.
- Cannot find device name for the disk with target ID.

Action: The device monitoring utilities had unexpected outputs. A resolved alert will be sent when the problem is resolved. If the problem persists, then contact Oracle Support Services.

HALRT-01037, {0} output has duplicate name: {1}

Cause: The output had a duplicate name.

Action: See "[HALRT-10073](#) (page C-42)".

HALRT-01038, {0} output has duplicate serial number: {1}

Cause: The output had a duplicate serial number.

Action: See "[HALRT-10073](#) (page C-42)".

HALRT-01039, {0} output has duplicate OS device name: {1}

Cause: The output had a duplicate operating system device name.

Action: "[HALRT-10073](#) (page C-42)".

HALRT-01040, Cannot find device name for the disk with target ID: {0}

Cause: The device was not found.

Action: "[HALRT-10073](#) (page C-42)".

HALRT-01041, The problems in the device utilities have been resolved.

Cause: The problem was resolved.

Action: This is an informational message, and no action is required.

HALRT-01042, {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}

Level Warning

Cause: A disk was replaced.

Action: The grid disks and cell disk are re-created on the new physical disk. If the grid disks were part of an Oracle ASM disk group, then they will be added back to the disk group, and the data will be rebalanced on them, based on the disk group redundancy and the ASM_POWER_LIMIT parameter.

HALRT-01043, {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10} Flash Cache: {11} Flash Log {12}

Level Warning

Cause: The flash disk was replaced.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then flash cache will be reenabled automatically.
- If the flash disk in this slot was previously used for flash log, then flash log will be reenabled automatically.
- If the flash disk in this slot was previously used for grid disks, then Oracle ASM will re-create the grid disks and rebalance the data automatically.

HALRT-01044, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}

Level Critical

Cause: A system disk failed.

Action: Locate the affected drive and replace it as soon as possible. A white LED is lit on the affected storage server, and both an amber service action LED and a blue "OK to Remove" LED are lit on the drive. Oracle ASM rebalance will automatically restore the data redundancy.

HALRT-01045, {0} entered predictive failure status. {1} was replaced. Status: {2} Manufacturer : {3} Model Number : {4} Size: {5} Serial Number : {6} Firmware: {7} Slot Number : {8} Cell Disc: {9} Grid Disk: {10}

Level Critical

Cause: The system hard disk entered predictive failure status.

Action: Locate the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically rebalance the data from the failing disk to other disks. Another alert is sent and a blue OK to Remove LED is lit on the drive when rebalancing completes. Replace the disk after the rebalancing is complete.

HALRT-01046, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}

Level Critical

Cause: The hard disk failed.

Action: Locate the failed disk and replace it as soon as possible. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive.

HALRT-01047, {0} entered predictive failure status. {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}

Level Critical

Cause: The data hard disk entered predictive failure status.

Action: Locate the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically rebalance the data from the failing disk to other disks. Another alert is sent and a blue OK to Remove LED is lit on the drive when rebalancing completes. Replace the disk after the rebalancing is complete.

HALRT-01048, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}

Level Critical

Cause: The flash disk failed.

Action: Replace the flash disk. A white LED is lit on the affected server. If the flash disk is used for flash cache, then flash cache is disabled on this disk, thus reducing the effective flash cache size. If the flash disk is used for flash log, then flash log is disabled on this disk, thus reducing the effective flash log size. If the flash disk is used for grid disks, then Oracle ASM rebalance automatically restores the data redundancy.

HALRT-01049, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}

Level Critical

Cause: The system hard disk failed.

Action: Locate the failed drive, and replace it as soon as possible. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance automatically restores the data redundancy.

HALRT-01050, {0} entered predictive failure status. {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}

Cause: The system hard disk entered predictive failure.

Action: Locate the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically rebalance the data from the failing disk to other disks. Another alert is sent when rebalancing completes. Replace the disk after the rebalancing is complete.

HALRT-01051, {0} failed. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}.

Level Critical

Cause: The data hard disk failed.

Action: Replace the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically restore the data redundancy.

HALRT-01052, {0} entered predictive failure status. {0} was replaced. Status: {1} Manufacturer : {2} Model Number : {3} Size: {4} Serial Number : {5} Firmware: {6} Slot Number : {7} Cell Disc: {9} Grid Disk: {10}

Level Critical

Cause: The data hard disk entered predictive failure.

Action: Locate the affected drive. A white LED is lit on the affected storage server, and an amber service action LED is lit on the drive. Oracle ASM rebalance will automatically rebalance the data from the failing disk to other disks. Another alert is

sent and a blue "OK to Remove" LED is lit on the drive when rebalancing completes. Replace the disk after the rebalancing is complete.

HALRT-01053, The Oracle ASM drop operation on the following disk failed due to ORA-{0}** ASM Disk Name: **{1}** ASM Diskgroup: **{2}** Grid Disk: **{3}** Cell Disk: **{4}** Physical Disk: **{5}** Physical Disk Status: **{6}****

Level Critical

Cause: The grid disk was not dropped from the Oracle ASM disk group, because of the specified reason.

Action: None. The storage server software continues to retry the command, and sends an alert when the command succeeds.

HALRT-01054, The Oracle ASM drop operation on the following disk failed due to ORA-{0}** ASM Disk Name: **{1}** ASM Diskgroup: **{2}** Grid Disk: **{3}** Cell Disk: **{4}** Physical Disk: **{5}** Physical Disk Status: **{6}****

Level Critical

Cause: The Oracle ASM `DROP` command failed, because of the specified reason.

Action: Retry the `DROP` command manually on the Oracle ASM disk.

HALRT-01055, The Oracle ASM rebalance operation on the following disk failed due to ORA-{0}** ASM Disk Name: **{1}** ASM Diskgroup: **{2}** Grid Disk: **{3}** Cell Disk: **{4}** Physical Disk: **{5}** Physical Disk Status: **{6}****

Level Critical

Cause: The Oracle ASM rebalance operation failed, because of the specified reason.

Action: This is an informational message, and no action is required.

HALRT-01056, The Oracle ASM rebalance operation on the following disk resumed ORA-{0}** ASM Disk Name: **{1}** ASM Diskgroup: **{2}** Grid Disk: **{3}** Cell Disk: **{4}** Physical Disk: **{5}** Physical Disk Status: **{6}****

Level Warning

Cause: The Oracle ASM rebalance operation resumed.

Action: This is an informational message, and no action is required.

HALRT-01057, The Oracle ASM add operation on the following disk failed due to ORA-{0}** ASM Disk Name: **{1}** ASM Diskgroup: **{2}** Grid Disk: **{3}** Cell Disk: **{4}** Physical Disk: **{5}** Physical Disk Status: **{6}****

Level Critical

Cause: The Oracle ASM operation to add the disk failed because of the specified reason.

Action: None. The storage server software continues to retry the command, and sends an alert when the command succeeds.

HALRT-01058, The Oracle ASM add operation on the following disk failed due to ORA-{0}** ASM Disk Name: **{1}** ASM Diskgroup: **{2}** Grid Disk: **{3}** Cell Disk: **{4}** Physical Disk Name: **{5}** Physical Disk Status: **{6}****

Level Critical

Cause: The Oracle ASM operation to add the disk failed, because of the specified reason.

Action: Retry the command manually on the Oracle ASM disk.

HALRT-01059, Oracle Exadata completed the rebalance operation triggered by the add operation on the following ASM disks. ASM Disk Name: {0} ASM Diskgroup: {1} Grid Disk: {2} Cell Disk: {3} Physical Disk: {4} Physical Disk Status : {5}

Level Clear

Cause: The rebalance operation completed.

Action: This is an informational message, and no action is required.

HALRT-01060, Oracle Exadata completed the rebalance operation triggered by the drop operation on the following ASM disks. ASM Disk Name: {0} ASM Diskgroup: {1} Grid Disk: {2} Cell Disk: {3} Physical Disk: {4} Physical Disk Status : {5}

Level Warning

Cause: The rebalance operation to the Oracle ASM disk group completed.

Action: This is an informational message, and no action is required.

HALRT-01061, {0} is ready to use. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}

Level Clear

Cause: The flash disk was ready.

Action: This is an informational message, and no action is required.

HALRT-01062, {0} is ready to use. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}

Level Clear

Cause: The hard disk was ready.

Action: This is an informational message, and no action is required.

HALRT-01063, The Oracle ASM rebalance operation on the following disk failed as ASM rebalance operation is disabled. ASM Disk Name: {0} ASM Diskgroup: {1} Grid Disk: {2} Cell Disk: {3} Physical Disk: {4} Physical Disk Status : {5}

Level Critical

Cause: The rebalance operation failed, because the Oracle ASM rebalance operation is disabled.

Action: Enable the rebalance operation, and retry the command.

HALRT-01106, All Logical drives are in WriteThrough caching mode. HDD disk controller battery needs to be replaced. Please contact Oracle Support.

Cause: The logical drive was found.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01107, Disk drive {0} is in WriteThrough caching mode. Write Performance on this drive will be affected. Please contact Oracle Support.

Cause: The disk drive was lost.

Action: See "[HALRT-10054](#) (page C-42)".

HALRT-01108, Logical drive status changed. LUN: {0}. Status: {1}. Physical {2} disk: {3}. Slot Num: {4}. Serial Number: {5}. Cell Disk: {6}. Grid Disks: {7}

Cause: The logical drive status has changed.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01109, Disk controller battery temperature exceeded the threshold of {0} Celsius. All disk drives have been placed in WriteThrough caching mode. The flash drives are not affected. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge {5} Ambient Temperature: {6}

Cause: The temperature exceeded the threshold.

Action: Verify that the ambient temperature is below 32 Celsius and the fans are working.

HALRT-01110, Logical drive found. It was empty. Failed creating cell disk and grid disks on it. LUN: {0}. Status: {1}. Physical: {2}. Slot Num: {3}. Serial Number: {4}

Cause: The logical drive was found.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01115, BGI completed with uncorrectable errors on Adapter {0} VD Target {1}

Cause: BGI completed with errors.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01116, BGI failed on Adapter {0} VD Target {1}

Cause: BGI failed on adapter.

Action: See "[HALRT-10039](#) (page C-41)".

HALRT-01117, LUN state changed on Adapter {0} VD Target {1} from {2} to {3}

Cause: The LUN state was changed.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01118, LUN bad block table cleared on Adapter {0} VD Target {1}

Cause: The LUN bad block table was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01123, {0} is rejected due to incorrect disk model. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}

Level Critical

Cause: The hard disk or LUN was removed. The storage server cannot use this hard disk because the disk model is not supported.

Action: Remove the hard disk and replace it with a supported disk.



See Also:

["Maintaining the Physical Disks of Storage Servers \(page 12-26\)".](#)

HALRT-01124, Hard Disk replaced. Status: {0} Manufacturer: {1} Model Number: {2} Size: {3} Serial Number: {4} Firmware: {5} Slot Number: {6} Cell Disk: {7} Grid Disk: {8}
Level Warning

Cause: The hard disk was replaced.

Action: This is an informational message, and no action is required.

HALRT-01126, The full charge capacity of the disk controller battery has fallen below the minimum threshold of {0} mAh. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be somewhat reduced. The flash drives are not affected. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}

Cause: The full charge capacity of the battery was below the threshold.

Action: Replace the disk controller battery.

HALRT-01127, All disk drives have been placed in WriteThrough caching mode. The flash drives are not affected. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}

Cause: The cause could not be determined.

Action: Contact Oracle Support Services.

HALRT-01128, Cell disk was missing and is recreated on the {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}

Cause: The cell disk was missing, and has been recreated on the disk.

Action: See [HALRT-10001 \(page C-39\)](#)"[HALRT-10001 \(page C-39\)](#)".

HALRT-01130, {0} status changed to poor performance. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}

Cause: The hard disk or flash disk status changed to poor performance.

Action: See ["Maintaining the Physical Disks of Storage Servers \(page 12-26\)".](#)

HALRT-01131, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}

Level Warning

Cause: The hard disk or flash disk status changed to normal.

Action: This is an informational message, and no action is required.

HALRT-01132, {0} removed. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Error Count: {10} Last Failure: {11}
Level Critical

Cause: The hard or flash disk was removed.

Action: This is an informational message, and no action is required.

HALRT-01133, {0} status changed to offline. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}.
Level Critical

Cause: The hard or flash disk status changed to offline.

Action: This is an informational message, and no action is required.

HALRT-01141, Flash Disk was replaced. Status: {0} Manufacturer: {1} Model Number: {2} Size: {3} Serial Number: {4} Firmware: {5} Slot Number: {6} Cell Disk: {7} Grid Disk: {8} Flash Cache: {9} Flash Log {10}
Level Warning

Cause: Flash disk was replaced.

Action: This is an informational message, and no action is required.

HALRT-01143, All disk drives have been placed in WriteThrough caching mode. The controller has preserved cache for a disk drive for more than five minutes. The flash drives are not affected. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}

Cause: The preserved cache was not turned off for more than five minutes.

Action: Contact Oracle Support Services.

HALRT-01144, Cell disk was missing and is recreated on the {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}

Cause: The cell disk was missing and has been recreated on the drive.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01145, Flash Disk status changed to predictive failure. Status: {0} Manufacturer: {1} Model Number: {2} Size: {3} Serial Number: {4} Firmware: {5} Slot Number: {6} Cell Disk: {7} Grid Disk: {8} Flash Cache: {9} Flash Log: {10}
Level Critical

Cause: Flash disk status changed to predictive failure.

Action: The flash disk has entered predictive failure status. A white cell locator LED has been lit to help locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.

- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait until rebalance and flush have completed before replacing the flash disk.



See Also:

"[Maintaining the Flash Disks of Storage Servers](#) (page 12-37)".

HALRT-01146, {0} entered poor performance status. The LUN {1} changed status to {2}. Status: {3} Manufacturer: {4} Model Number: {5} Size: {6} Serial Number: {7} Firmware: {8} Slot Number: {9} Cell Disk: {10} Grid Disk: {11} Flash Cache: {12} Flash Log: {13} Reason for poor performance: {14}

Cause: Flash disk status changed to `poor performance`.

Action: See "[Maintaining the Flash Disks of Storage Servers](#) (page 12-37)".

HALRT-01147, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}

Level Warning

Cause: Flash disk status changed to `normal`.

Action: This is an informational message, and no action is required.

HALRT-01148, Flash Disk removed. Status: {0} Manufacturer: {1} Model Number: {2} Size: {3} Serial Number: {4} Firmware: {5} Slot Number: {6} Cell Disk: {7} Grid Disk: {8} Flash Cache: {9} Flash Log: {10} Error Count: {11} Last Failure: {12}

Level Critical

Cause: Flash disk was removed.

Action: A flash disk has been removed or is no longer visible to Exadata Storage Server.

- If a flash disk has just been replaced, then this alert is informational only and a flash disk replaced CLEAR alert is sent. No action is needed.
- If flash is being used as a cache, then the only effect is that the flash cache size has been reduced on this cell. The flash disk removal does not cause an outage, data loss, or Oracle ASM rebalance. Each flash card contains four flash disks.

Each flash card contains four flash disks. If all four flash disks on a card alert, then the removal is likely for a card. To see whether the flash disks are on the same card, check for identical PCI Slot numbers in the alert outputs.

HALRT-01149, {0} status changed to offline. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}

Level Warning

Cause: Flash disk status changed to *offline*.

Action: This is an informational message, and no action is required.

HALRT-01150, Cell disk was missing and is recreated on the {0}.

Cause: The cell disk was missing, and has been recreated on the disk.

Action: "[HALRT-10001](#) (page C-39)".

HALRT-01153, Disk drive {0} is back in WriteBack caching mode.

Cause: The disk drive entered write back caching mode.

Action: See "[HALRT-10506](#) (page C-44)".

HALRT-01154, Oracle Exadata Storage Server failed to auto-create cell disk and grid disks on the newly inserted physical disk. Physical Disk: {0} Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}

Cause: The server did not auto-create the cell disk or grid disks.

Action: The storage server software continues to retry, and an alert is sent when the auto-creation succeeds.

HALRT-01155, All disk drives have been temporarily placed in WriteThrough caching mode because the disk controller battery was replaced. The flash drives are not affected. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}

Cause: Disk controller battery was replaced.

Action: This is an informational message, and no action is required.

HALRT-01156, Previously missing {0} is found. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}
Level Clear

Cause: The missing flash disk was found.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then the flash cache are automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then the flash log are automatically reenabled.
- If the flash disk in this slot was previously used for grid disks, then the grid disks are automatically recreated and data will be automatically rebalanced by Oracle ASM.

HALRT-01157, Previously missing {0} is found. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}

Level Clear

Cause: The missing hard disk was found.

Action: This is an informational message, and no action is required.

HALRT-01159, {0} import failed. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Warning

Cause: The import procedure failed.

Action: System failed to import this hard disk because preserved cache exists for other missing hard disks. Reseat or replace the hard disks in {0}. Another alert is sent when all disks are successfully imported.

HALRT-01160, Failed to create LUN on {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Informational

Cause: The disk controller did not create a LUN on the physical disk.

Action: Contact Oracle Support Services.

HALRT-01161, All disk drives are back in WriteBack caching mode. Preserved cache has turned off. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}

Cause: Disk drives returned to write back caching mode.

Action: This is an informational message, and no action is required.

HALRT-01162, All disk drives are in WriteBack caching mode. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}

Cause: Disk drives were in write back caching mode.

Action: This is an informational message, and no action is required.

HALRT-01163, All disk drives are back in WriteBack caching mode. Battery has completed charging. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}.

Cause: Battery completed charging, and all disk drives returned to write back caching mode.

Action: This is an informational message, and no action is required.

HALRT-01164, All disk drives are in WriteBack caching mode. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}

Cause: Disk drives were in write back caching mode.

Action: This is an informational message, and no action is required.

HALRT-01165, {0} can be replaced now. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Critical

Cause: Physical disk needed replacement.

Action: Replace the physical disk.

HALRT-01166, {0} can be replaced now. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk {9} Flash Cache {11} Flash Log {12}
Level Critical

Cause: Flash disk needed replacement.

Action: Replace the flash disk.

HALRT-01167, {0} can be replaced now.

Level Critical

Cause: The component needed replacement.

Action: Replace the component.

HALRT-01168, {0} was removed and inserted. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}

Level Informational

Cause: The component was inserted.

Action: This is an informational message, and no action is required. CELLSRV may have reported I/O errors when the physical disk was removed. If any grid disk on this physical disk went offline in Oracle ASM as a result, then Oracle ASM automatically brings the disk online.

HALRT-01169, {0} entered write-through caching mode. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}

Level Critical

Cause: The component entered write through caching mode.

Action: The flash disk has entered predictive failure status. A white cell locator LED has been lit to help locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait until rebalance and flush have completed before replacing the flash disk.

 **See Also:**

["Maintaining the Flash Disks of Storage Servers \(page 12-37\)".](#)

**HALRT-01170, Failed to flush the write back flash cache. Status: {1}
Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6}
Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}**
Level Critical

Cause: Failed to flush the flash cache.

Action: Check the alert log for additional details. After the underlying problem is resolved, flush the write back flash cache on this cell disk is automatically resumed.

HALRT-01171, {0} entered confinement offline status. The LUN {1} changed status to {2}. CellDisk changed status to {3}. All subsequent I/Os on this disk are failed immediately. Confinement tests will be run on the disk to determine if the disk should be dropped. Status: {4} Manufacturer: {5} Model Number: {6} Size: {7} Serial Number: {8} Firmware: {9} Slot Number: {10} Cell Disk: {11} Grid Disk: {12} Reason for confinement: {13}

Cause: The hard disk was performing poorly, and entered confinement.

Action: See "[HALRT-10081](#) (page C-43)".

HALRT-01172, {0} entered confinement offline status. The LUN {1} changed status to {15}. CellDisk changed status to {2}. All subsequent I/Os on this disk are failed immediately. Confinement tests will be run on the disk to determine if the disk should be dropped. Status: {4} Manufacturer: {5} Model Number: {6} Size: {7} Serial Number: {8} Firmware: {9} Slot Number: {10} Cell Disk: {11} Grid Disk: {12} Flash Cache: {13} Flash Log: {14} Reason for confinement: {15}

Cause: The flash disk was performing poorly, and entered confinement.

Action: See "[HALRT-10082](#) (page C-43)".

HALRT-01173, {0} entered confinement status. The LUN {1} changed status to {2}. CellDisk changed status to {3}. Status: {4} Manufacturer: {5} Model Number: {6} Size: {7} Serial Number: {8} Firmware: {9} Slot Number: {10} Cell Disk: {11} Grid Disk: {12} Reason for confinement: {13}

Cause: The system disk entered confinement.

Action: See "[HALRT-10079](#) (page C-43)".

HALRT-01174, {0} entered confinement status. The LUN {1} changed status to {2}. CellDisk changed status to {3}. Status: {4} Manufacturer: {5} Model Number: {6} Size: {7} Serial Number: {8} Firmware: {9} Slot Number: {10} Cell Disk: {11} Grid Disk: {12} Flash Cache: {13} Flash Log: {14} Reason for confinement: {15}

Cause: The disk entered confinement.

Action: See "[HALRT-10080](#) (page C-43)".

HALRT-01175, Partition table is missing in system slot {0}.
Level Warning

Cause: The partition table was not on the drive in the system slot.

Action: Contact Oracle Support Services.

HALRT-01176, The data disk from slot {1} was inserted into system slot {2} of the same Oracle Exadata storage server. Status: {3} Manufacturer: {4} Model Number: {5} Size: {6} Serial Number: {7} Firmware: {8} Slot Number: {9}
Level Warning

Cause: The data disk was inserted into the wrong slot.

Action: Insert the disk into the correct slot.

HALRT-01177, The system disk from slot {0} was inserted into data slot {1} of the same Oracle Exadata storage server. Status: {2} Manufacturer: {3} Model Number: {4} Size: {5} Serial Number: {6} Firmware: {7} Slot Number: {8}
Level Warning

Cause: A system disk from the same Exadata Storage Server was inserted in a data disk slot.

Action: Insert the correct disk into the slot.

HALRT-01178, The system disk from a different Oracle Exadata storage server was inserted in slot {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Warning

Cause: A system disk from a different storage server was inserted in the slot.

Action: Insert the correct disk into the slot.

HALRT-01179, A data disk from a different Oracle Exadata storage server was inserted in system slot {0}. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}.
Level Warning

Cause: A data disk from a different storage server was inserted in a system disk slot.

Action: Insert the correct disk into the slot.

HALRT-01181, {0} entered peer failure status. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Level Critical

Cause: The flash disk entered peer failure status.

Action: The flash disk has entered predictive failure status. A white cell locator LED has been lit to help locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait until rebalance and flush have completed before replacing the flash disk.

HALRT01182, {0} was removed before running the ALTER PHYSICALDISK DROP FOR REPLACEMENT command. Status: {1} Manufacturer: {2} Model

Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk {9}
Level Critical

Cause: The disk was removed before the ALTER PHYSICALDISK DROP FOR REPLACEMENT command was run.

Action: This is an informational message, and no action is required.

HALRT01183, {0} was removed before running the ALTER PHYSICALDISK DROP FOR REPLACEMENT command and inserted. {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}
Level Informational

Cause: The disk was removed before the ALTER PHYSICALDISK DROP FOR REPLACEMENT command was run.

Action: This is an informational message, and no action is required. If an Oracle ASM grid disk was on the physical disk when the disk went offline, then Oracle ASM automatically brings the disk online.

HALRT-01184, {0} automatically failed over to a spare disk. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}
Level Critical

Cause: The cell disk failed over to the spare disk.

Action: This is an informational message, and no action is required.

HALRT-01185, {0} automatically failed over to a spare disk. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Cause: The flash disk failed over to the spare disk.

Action: This is an informational message, and no action is required.

HALRT-01186, {0} was removed. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Level Critical

Cause: A flash disk was removed.

Action: A flash disk was removed or is no longer visible.

- If the flash disk has been replaced, then this alert is informational only, and a flash disk replaced CLEAR alert will be sent. No action is needed.
- If flash is being used as a cache, then the effect is that the flash cache size has reduced on the cell. Flash disk removal does not cause an outage, data loss, or Oracle ASM rebalance.
- Each flash card contains four flash disks. If all four flash disks on a card send an alert, then the removal is likely for a card. To determine if the flash disks are on the same card, check for identical PCI slot numbers in the alert output.

HALRT-01187, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}
Level Warning

Cause: The flash disk was replaced.

Action: This is an informational message, and no action is required.

HALRT--01188, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11}.Board Tracer Number : {12}
Level Warning

Cause: The flash disk was replaced.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then the flash cache is automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then the flash log is automatically reenabled.
- If the flash disk in this slot was previously used for grid disks, then the grid disks are automatically recreated and the data will be automatically rebalanced by Oracle ASM.

HALRT--01189, Oracle Exadata Storage Server failed to auto-create cell disk and grid disks on the newly inserted physical disk. Physical Disk: {0} Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Board Tracer Number: {8}
Level Warning

Cause: The component was replaced.

Action: The storage server software continues to retry, and sends an alert when the auto-create operation succeeds.

HALRT--01190, Previously missing {0} is found. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}
Level Critical

Cause: The flash disk was located.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then the flash cache is automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then the flash log is automatically reenabled.
- If the flash disk in this slot was previously used for grid disks, then the grid disks are automatically recreated, and the data is automatically rebalanced by Oracle ASM.

HALRT--01191, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}

Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}
Level Warning

Cause: The flash disk status changed to normal.

Action: This is an informational message, and no action is required.

HALRT--01192, {0} was removed and inserted. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Board Tracer Number : {9}
Level Informational

Cause: The physical disk was removed and inserted.

Action: This is an informational message, and no action is required. CELLSRV might have reported I/O errors when the disk was removed. If any grid disk on the physical disk went offline in Oracle ASM, then Oracle ASM automatically brings the disk online.

HALRT--01193, {0} failed. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Board Tracer Number : {10}
Level Critical

Cause: A flash disk failed.

Action: Replace the failed flash disk. To identify the affected server, look for a lit white cell locator LED.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

HALRT--01194, {0} entered predictive failure status. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}
Level Critical

Cause: The flash disk entered predictive failure status.

Action: The flash disk has entered predictive failure status. Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait for the second alert before replacing the flash disk. A white cell locator LED has been lit to help you locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.

- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

HALRT--01195, {0} entered peer failure status. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number: {12}

Level Critical

Cause: The flash disk entered peer failure status.

Action: The flash disk has entered peer failure status. Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait for the second alert before replacing the flash disk. A white cell locator LED has been lit to help you locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

HALRT--01196, {0} entered write-through caching mode. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}

Level Critical

Cause: The flash disk entered write-through caching mode.

Action: The flash disk has entered write-through caching mode. Another alert will be sent when all data has been rebalanced and flushed from this flash disk. Wait for the second alert before replacing the flash disk. A white cell locator LED has been lit to help you locate the affected cell.

- If the flash disk is used for flash cache, then flash cache is disabled on this disk, reducing the effective flash cache size.
- If flash cache is in write back mode, then dirty data on this flash disk is automatically flushed to the corresponding grid disks.
- If the flash disk is used for flash log, then flash log is disabled on this disk, reducing the effective flash log size.
- If the flash disk is used for grid disks, then the data from the flash disk is automatically rebalanced by Oracle ASM to other flash disks.

HALRT--01197, {0} sstatus changed to offline. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number: {12}

Level Critical

Cause: The component status changed to offline.

Action: This is an informational message, and no action is required.

HALRT--01198, {0} automatically failed over to a spare disk. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache {10} Flash Log {11} Board Tracer Number : {12}

Level Critical

Cause: The component failed over to a spare disk.

Action: This is an informational message, and no action is required.

C.8 Accelerator/Controller Alert Messages: HALRT-01209 to HALRT-01300

HALRT-01209, Alarm has been enabled by user for Adapter {0}.

Cause: An alarm was enabled for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01210, Background initialization rate changed to {0} for Adapter {1}.

Cause: The background initialization rate was changed for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01211, Controller {0} cache discarded due to memory/battery problems.

Cause: The controller cache was discarded due to memory or battery problems.

Action: See "[HALRT-10033](#) (page C-40)".

HALRT-01212, Unable to recover Cache Data due to configuration mismatch for Adapter {0}.

Cause: Cache data was recovered due to a configuration mismatch.

Action: See "[HALRT-10033](#) (page C-40)".

HALRT-01213, Cache Data Recovered for Adapter {0}.

Cause: Cache data was recovered for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01214, Controller cache discarded due to firmware version incompatibility for Adapter {0}

Cause: Controller cache was discarded because of incompatibility with the firmware version.

Action: See "[HALRT-10033](#) (page C-40)".

HALRT-01215, Consistency Check rate changed to {0} for Adapter {1}.

Cause: The consistency check rate was changed for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01216, Flash downloaded image corrupted for Adapter {0}.

Cause: The flash memory card image was corrupted.

Action: See "[HALRT-10034](#) (page C-40)".

HALRT-01217, Flash erase error for Adapter {0}.

Cause: There was a flash erase error for the adapter.

Action: See "[HALRT-10034](#) (page C-40)".

HALRT-01218, Flash timeout during erase for Adapter {0}.

Cause: The flash timed out during the erase procedure.

Action: See "[HALRT-10034](#) (page C-40)".

HALRT-01219, Flash error for Adapter {0}.

Cause: There was a flash error for the adapter.

Action: See "[HALRT-10034](#) (page C-40)".

HALRT-01220, Flashing image: {0} for Adapter {1}.

Cause: The controller flash was upgraded with a new image.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01221, Flash of new firmware image(s) complete for Adapter {0}.

Cause: The new firmware image was copied to flash.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01222, Flash programming error for Adapter {0}.

Cause: There was a flash programming error for the adapter.

Action: See "[HALRT-10034](#) (page C-40)".

HALRT-01223, Flash timeout during programming for Adapter {0}.

Cause: The flash timed out while programming of the adapter.

Action: See "[HALRT-10034](#) (page C-40)".

HALRT-01224, Fatal error received for Adapter {0}.

Cause: A fatal error occurred.

Action: See "[HALRT-10035](#) (page C-41)".

HALRT-01225, Rebuild rate changed to {0} for Adapter {1}.

Cause: The rebuild rate changed for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01226, Alarm has been disabled by user for Adapter {0}.

Cause: The alarm for the adapter was disabled.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01227, Configuration cleared for Adapter {0}.

Cause: The configuration for the adapter was cleared.

Action: See "[HALRT-10036](#) (page C-41)".

HALRT-01228, Factory defaults restored for Adapter {0}.

Cause: The factory default settings were restored for the adapter.

Action: See "[HALRT-10036](#) (page C-41)".

HALRT-01229, HDD disk controller battery not Present for Adapter {0}.

Cause: The battery was not present for the adapter.

Action: See "[HALRT-10036](#) (page C-41)".

HALRT-01230, New HDD disk controller battery Detected for Adapter {0}.

Cause: A new battery was detected for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01231, HDD disk controller battery has been replaced for Adapter {0}.

Cause: The battery was replaced for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01232, HDD disk controller battery temperature is high for Adapter {0}.

Cause: The battery temperature was high for the adapter.

Action: See "[HALRT-10037](#) (page C-41)".

HALRT-01233, HDD disk controller battery voltage low for Adapter {0}.

Cause: The battery voltage was low for the adapter.

Action: See "[HALRT-10038](#) (page C-41)".

HALRT-01234, HDD disk controller battery temperature is normal for Adapter {0}.

Cause: The battery temperature was normal for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01235, HDD disk controller battery needs replacement, SOH Bad for Adapter {0}.

Cause: The battery needed replacement.

Action: See "[HALRT-10038](#) (page C-41)".

HALRT-01236, HDD disk controller battery removed for Adapter {0}.

Cause: The battery was removed for the adapter.

Action: See "[HALRT-10036](#) (page C-41)".

HALRT-01237, Current capacity of the HDD disk controller battery is below threshold for Adapter {0}.

Cause: The battery capacity was below the threshold for the adapter.

Action: See "[HALRT-10038](#) (page C-41)".

HALRT-01238, Current capacity of the HDD disk controller battery is above threshold for Adapter {0}.

Cause: The battery capacity was above the threshold for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01239, Bbu FRU changed for Adapter {0}.

Cause: The battery backup unit (BBU) field replacement unit (FRU) changed for the adapter.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01240, The disk controller battery is executing a learn cycle and may temporarily enter WriteThrough Caching mode as part of the learn cycle. Disk write throughput might be temporarily lower during this time. The flash drives are not affected. The battery learn cycle is a normal maintenance activity that occurs quarterly and runs for approximately 1 to 12 hours. Note that many learn cycles do not require entering WriteThrough caching mode. When the disk controller cache returns to the normal WriteBack caching mode, an additional informational alert will be sent. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}

Cause: Disk controller battery started the learn cycle.

Action: This is an informational message, and no action is required. A learn cycle is a normal maintenance activity that occurs quarterly and runs for approximately 1 to 12 hours. The disk controller cache might go into write through caching mode during the learn cycle. Disk write throughput might be temporarily lower during this time.

 **See Also:**

["Monitoring Hard Disk Controller Write-Through Cache Mode \(page 12-30\)"](#)

HALRT-01243, All disk drives have been placed in WriteThrough caching mode. The flash drives are not affected. The disk controller battery is executing an unscheduled learn cycle. Battery Serial Number: {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity: {4} Relative Charge: {5} Ambient Temperature: {6}.

Cause: An unscheduled learn cycle started.

Action: Determine the reason for the unscheduled learn cycle.

HALRT-01244, HDD disk controller SAS lanes have encountered errors.

Cause: A problem was encountered for the hard disk drive disk controller SAS lanes.

Action: See "[HALRT-10035](#) (page C-41)".

HALRT-01245, HDD disk controller SAS lanes are back to normal.

Cause: The hard disk drive disk controller SAS lanes were normal.

Action: This is an informational message.

C.9 Voltage Alert Messages: HALRT-01301 to HALRT-01400

HALRT-01301, Under-Voltage Problem (lower critical - going low)

Cause: There was an under-voltage problem.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01302, Under-Voltage Problem Cleared

Cause: The under-voltage problem was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01303, Generic Critical Voltage Problem

Cause: There was a generic voltage problem.

Action: See "[HALRT-10001](#) (page C-39)".

HALRT-01304, Generic Critical Voltage Problem Cleared

Cause: The generic voltage problem was cleared.

Action: See "[HALRT-10001](#) (page C-39)".

C.10 Physical Disk Alert Messages: HALRT-01401 to HALRT-01500

HALRT-01401, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}

Level Clear

Cause: The component was replaced in the rack.

Action: The grid disks are automatically recreated, and the data is automatically rebalanced by Oracle ASM.

HALRT-01402, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}

Level Clear

Cause: The flash disk was replaced in the rack.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then flash cache will be automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then flash log will be automatically reenabled.

- If the flash disk in this slot was previously used for grid disks, then grid disks will be automatically recreated and data will be automatically rebalanced by Oracle ASM.

HALRT-01403, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}

Level Clear

Cause: The component was replaced in the rack.

Action: This is an informational message, and no action is required.

HALRT-01404, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}

Level Clear

Cause: The component was replaced in the rack.

Action: This is an informational message, and no action is required.

HALRT-01405, Oracle Exadata Storage Server failed to auto-create cell disk and grid disks on the newly inserted physical disk. Physical Disk: {0} Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7}

Level Clear

Cause: The component was replaced in the rack.

Action: The storage server software continues to try to create the cell disk and grid disks. An alert is sent when the auto-creation succeeds.

HALRT-01406, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11} Board Tracer Number: {12}

Level Clear

Cause: The component was replaced in the rack.

Action: This is an informational message, and no action is required.

HALRT-01407, {0} was replaced. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11} Board Tracer Number: {12}

Level Clear

Cause: The flash disk was replaced in the rack.

Action: This is an informational message, and no action is required.

- If the flash disk in this slot was previously used for flash cache, then flash cache will be automatically reenabled.
- If the flash disk in this slot was previously used for flash log, then flash log will be automatically reenabled.

- If the flash disk in this slot was previously used for grid disks, then grid disks will be automatically recreated and data will be automatically rebalanced by Oracle ASM.

HALRT-01408, Oracle Exadata Storage Server failed to auto-create cell disk and grid disks on the newly inserted physical disk. Physical Disk: {0} Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Board Tracer Number: {8}

Level Clear

Cause: The component was replaced in the rack.

Action: The storage server software continues to try to create the server disk and grid disks. An alert is sent when the auto-creation succeeds.

HALRT-01409, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9}

Level Clear

Cause: The component status changed to normal.

Action: This is an informational message, and no action is required.

HALRT-01410, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11}

Level Clear

Cause: The component status changed to normal.

Action: This is an informational message, and no action is required.

HALRT-01411, {0} status changed to normal. Status: {1} Manufacturer: {2} Model Number: {3} Size: {4} Serial Number: {5} Firmware: {6} Slot Number: {7} Cell Disk: {8} Grid Disk: {9} Flash Cache: {10} Flash Log: {11} Board Tracer Number: {12}

Level Clear

Cause: The component status changed to normal.

Action: This is an informational message, and no action is required.

C.11 Auto Service Request Alert Messages for Storage Servers: HALRT-02001 to HALRT-02013

HALRT-02001

Cause: The system disk has failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02002

Cause: The system disk had a predictive failure.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02003

Cause: A data disk has failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02004

Cause: A data disk had a `predictive failure`.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02005

Cause: A flash disk has failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02006

Cause: A flash disk had a `predictive failure`.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02007

Cause: A hard disk in a compute server has failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02008

Cause: A hard disk in a compute server had a `predictive failure`.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02009

Cause: System hard disk had poor performance.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02010

Cause: Data hard disk had poor performance.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02011

Cause: Flash disk had poor performance.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02012

Cause: The BBU failed.

Action: An alert has been submitted to Oracle Support Services.

HALRT-02013

Cause: The BBU degraded for unknown reasons.

Action: An alert has been submitted to Oracle Support Services.

C.12 HDD Controller Battery Alert Messages: HALRT-02101 to HALRT-02200

HALRT-02101, The HDD disk controller battery was dropped for replacement. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}
Level Warning

Cause: The HDD disk controller was dropped before replacement.

Action: Replace the HDD disk controller battery or reenable the HDD disk controller battery using the `ALTER CELL BBU REENABLE` command if the HDD disk controller battery was dropped in error.

HALRT-02102, HDD disk controller battery is {0}. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Critical

Cause: The HDD disk controller battery failed.

Action: Replace the HDD disk controller battery.

HALRT-02103, All disk drives have been temporarily placed in WriteThrough caching mode because HDD disk controller battery is performing a learn cycle. Disk write performance may be somewhat reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Informational

Cause: The HDD disk controller battery was performing a learn cycle.

Action: This is an informational alert. Disk drives will be in `writeback` mode when the HDD disk controller battery completes the learn cycle.

HALRT-02104, All disk drives have been temporarily placed in WriteThrough caching mode because HDD disk controller battery is charging. Disk write performance may be somewhat reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Informational

Cause: The HDD disk controller battery was charging.

Action: This is an informational alert. Disk drives will be in `writeback` mode when the HDD disk controller battery completes the charging.

HALRT-02105, The HDD disk controller battery was removed. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected.
Level Informational

Cause: The HDD disk controller battery was removed.

Action: The HDD disk controller battery needs replacement.

HALRT-02106, The HDD disk controller battery capacity was under threshold value of {0} mAh. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Critical

Cause: The HDD disk controller battery capacity went too low.

Action: Replace the HDD disk controller battery or schedule a learn cycle.

HALRT-02107, The HDD disk controller battery temperature exceeded the threshold value of {6} Celsius. All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Critical

Cause: The HDD disk controller battery temperature was too high.

Action: The HDD disk controller battery was overheated. Check if the ambient temperature is below {0} Celsius and if the fans are working.

HALRT-02108, All disk drives have been temporarily placed in WriteThrough caching mode because HDD disk controller battery was replaced. Disk write performance may be somewhat reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Informational

Cause: The HDD disk controller battery was replaced.

Action: This is an informational alert. Disk drives will be in `writeback` mode until the HDD disk controller battery completes charging.

HALRT-02109, All disk drives have been placed in WriteThrough caching mode. Disk write performance may be reduced. The flash drives are not affected. Battery Serial Number : {1} Battery Type: {2} Battery Temperature: {3} Full Charge Capacity : {4} Relative Charge: {5} Ambient Temperature: {6}
Level Critical

Cause: The HDD disk controller battery entered degraded or unknown status.

Action: Replace the HDD disk controller battery or schedule a learn cycle.

HALRT-02110, The HDD disk controller battery was removed before running the ALTER CELL BBU DROP FOR REPLACEMENT command. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}
Level Critical

Cause: The HDD disk controller battery was removed, but the `ALTER CELL BBU DROP FOR REPLACEMENT` command was not run prior to removal of the battery.

Action: Replace the HDD disk controller battery.

HALRT-02111, The HDD disk controller battery was replaced. All disk drives are in WriteBack caching mode. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}

Level Clear

Cause: The HDD disk controller battery was replaced.

Action: This is an informational alert and no user action is specified.

HALRT-02112, All disk drives are in WriteBack caching mode. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}

Level Clear

Cause: The HDD disk controller battery returned to normal.

Action: This is an informational alert and no user action is specified.

HALRT-02113, All disk drives have been placed in WriteThrough caching mode. The controller has preserved cache for a disk drive for more than five minutes. The flash drives are not affected. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}

Level Critical

Cause: All disk drives were placed in `writethrough` caching mode because of preserved cache.

Action: Contact Oracle Support Services.

HALRT-02114, The HDD disk controller battery is performing an unscheduled learn cycle. All disk drives have been placed in WriteThrough caching mode. The flash drives are not affected. Battery Serial Number : {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity : {3} Relative Charge: {4} Ambient Temperature: {5}

Level Informational

Cause: The HDD disk controller battery performed an unscheduled learn cycle.

Action: Determine the cause of the unscheduled learn cycle.

HALRT-02115, Disk drive {0} is in WriteThrough caching mode. Write Performance on this drive will be affected.

Level Critical

Cause: Disk drive {0} entered `writethrough` caching mode.

Action: The HDD disk controller battery is good, but disk drive {0} is in `writethrough` caching mode. Contact Oracle Support Services.

HALRT-02116, Disk drive {0} is back in WriteBack caching mode.

Level Clear

Cause: Disk drive {0} returned to `writeback` caching mode.

Action: This is an informational alert.

HALRT-02117, The HDD disk controller battery is performing a learn cycle. Battery Serial Number: {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity: {3} Relative Charge: {4} Ambient Temperature: {5}
Level Informational

Cause: The hard disk controller battery entered a learn cycle.

Action: This is an informational alert, and no user action is needed.

All disk drives may temporarily enter `writethrough` caching mode as part of the learn cycle. Disk write throughput may be lower during this time. The flash drives are not affected.

The battery learn cycle is a normal maintenance activity that occurs quarterly, and runs for approximately 1 to 12 hours. Many learn cycles do not require entering `writethrough` caching mode. When the disk controller cache returns to `writeback` caching mode, an informational alert will be sent.

HALRT-02118, The HDD disk controller battery was replaced. All disk drives are in WriteBack caching mode. Battery Serial Number: {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity: {3} Relative Charge: {4} Ambient Temperature: {5}
Level Informational

Cause: The hard disk controller battery was replaced.

Action: This is an informational alert, and no user action is needed.

HALRT-02119, The HDD disk controller battery was replaced. Battery Serial Number: {0} Battery Type: {1} Battery Temperature: {2} Full Charge Capacity: {3} Relative Charge: {4} Ambient Temperature: {5}
Level Informational

Cause: The hard disk controller battery was replaced.

Action: This is an informational alert, and no user action is needed.

C.13 Generic Actions for Alert Messages: HALRT-10001 to HALRT-10500

HALRT-10001

Cause: There was a nonerror change for a hardware status, such as a return to normal status.

Action: This is an informational alert and no user action is specified.

HALRT-10003

Cause: An unscheduled learn cycle was started.

Action: Examine how the unscheduled learn cycle was started.

HALRT-10004

Cause: Temperature was not in the acceptable range.

Action: Shut down the system. Determine the cause of thermal failure. If necessary, move the data off of this server.

HALRT-10005

Cause: There was a problem with the fan.

Action: Shut down the system. Determine cause of the fan failure. Replace the fan. If necessary, move the data off of this server.

HALRT-10007

Cause: There was a problem with the power supply.

Action: Examine the power supply. Replace the power supply if necessary.

HALRT-10015

Cause: Temperature was not in the acceptable range.

Action: Examine thermal degradation. Shut down the system and move the data off of this server.

HALRT-10025

Cause: A CPU error was detected.

Action: Shut down the system. Examine CPU error.

HALRT-10029, A generic sensor has detected an error.

Cause: There was a sensor change or error.

Action: Examine the sensor state change or error. If it is unexpected and has not been cleared yet, contact Oracle Support Services.

HALRT-10030

Cause: There was an alert.

Action: Refer to the additional information in the alert message.

HALRT-10031

Cause: The FRU was removed.

Action: If the unit was not intentionally removed, contact Oracle Support Services.

HALRT-10032

Cause: There was an error.

Action: Capture any error message on the console. Restart or power cycle the system and try to bring it back to operation. Contact Oracle Support Services with the error messages.

HALRT-10033

Cause: The cache was lost.

Action: Losing the cache implies that any data on the disks is now stale. Drop force the data on disks attached to this controller from Oracle ASM and then add them back.

HALRT-10034

Cause: The disk controller firmware needed a refresh.

Action: Reflash the disk controller firmware.

HALRT-10035

Cause: The disk controller noticed errors.

Action: Contact Oracle Support Services.

HALRT-10036

Cause: There was a change to the system.

Action: If this change is unintentional, contact Oracle Support Services.

HALRT-10037

Cause: The operating temperature was not within the specified range for the machine.

Action: Ensure the operating temperature is within specified levels for the system.

HALRT-10038

Cause: The battery was not fully charged.

Action: If the battery is not fully charged, then charge it. Otherwise, the battery might be about to fail. Replace disk controller battery.

HALRT-10039

Cause: The LUN on the virtual drive has failed.

Action: Drop the LUN and re-create it.

HALRT-10040

Cause: The physical disk was removed.

Action: If the physical drive was removed unintentionally, replace it immediately.

HALRT-10041

Cause: The disk was not supported or certified for the system.

Action: Replace this disk with a supported and certified disk.

HALRT-10045

Cause: The flash write performance was degraded.

Action: Replace the flash disk, flash PCI card, or energy storage module (ESM). If all disks on the PCI card show degraded performance, then replace the card. Check the ESM LED on the PCI card. The LED is visible from the back of the cell chassis. If this LED is amber, then replace the card.

 **Note:**

This message may include text from message HALRT-10046 if there are grid disks, or from message HALRT-10047 if there are no grid disks but the server disk is used for flash cache.

HALRT-10046

Cause: There was an error with the disk.

Action: If the disks being replaced have persistent data, then rebalance data out of these disks, and replace the flash disk. If there was persistent data on the flash disks, then rebalance data back into the disks.

HALRT-10047

Cause: There was an error with the flash disk.

Action: The flash cache portion provided by the flash disk will be automatically disabled.

HALRT-10054

Cause: The physical disk demonstrated extremely poor performance.

Action: Replace the physical disk at the earliest opportunity.

HALRT-10056

Cause: Auto Service Request was notified with a unique identifier.

Action: This is an informational alert, and no user action is needed.

HALRT-10057

Cause: Detailed information on this problem can be found at My Oracle Support.

Action: Refer to the My Oracle Support note listed in the alert.

HALRT-10066

Cause: The system hard disk demonstrated extremely poor performance.

Action: Replace the system hard disk as soon as possible.

HALRT-10067

Cause: The data hard disk demonstrated extremely poor performance.

Action: Replace the data hard disk as soon as possible.

HALRT-10068

Cause: The flash disk demonstrated extremely poor performance.

Action: Replace the flash disk as soon as possible.

HALRT-10073

Cause: The device monitoring utilities had unexpected output.

Action: A resolved alert will be sent when the problem is resolved. If the problem persists, then Contact Oracle Support Services.

HALRT-10074

Cause: A flash disk was removed or was no longer visible to the storage server.

Action: If the flash disk has just been replaced, then this alert is informational only. If the flash disk is being used as a cache, then the flash cache size has been reduced on this server. Flash disk removal will not cause an outage, data loss, or Oracle ASM rebalance.

HALRT-10075

Cause: The temperature of Recovery Appliance is outside the operating range.

Action: Recovery Appliance has an operating range of 5 to 32 degrees Celsius (41 to 89.6 degrees Fahrenheit). Adjust the data center conditions so that Recovery Appliance is operating within the specified temperature range.

HALRT-10076

Cause: The Oracle ASM disks associated with the grid disks on the physical disk were successfully dropped, and a blue "OK to Remove LED" was lit on the disk to indicate it is safe to remove the disk.

Action: Replace the disk drive at the earliest opportunity.

HALRT-10078

Cause: CELLSRV reported I/O errors when the physical disk was removed.

Action: No action is needed. If a grid disk on this physical disk went offline in Oracle ASM as a result, then Oracle ASM will automatically bring the disk online.

HALRT-10079

Cause: The single system disk in the storage server did not have valid partitions.

Action: Contact Oracle Support Services.

HALRT-10080

Cause: A disk was inserted in the wrong slot.

Action: Insert the disk into the correct slot.

HALRT-10081

Cause: The hard disk demonstrated temporary poor performance. The disk has been confined and is in an offline state.

Action: If the hard disk was used for grid disks, then the Oracle ASM disks associated with the grid disks on this hard disk are offline. Confinement tests run on this disk, and either move the disk permanently to `poor performance` or return the disk back to `normal` state.

HALRT-10082

Cause: The flash disk demonstrated temporary poor performance. The disk has been confined and is in an offline state.

Action: One of the following actions is performed:

- If the flash disk was used for Write-Through flash cache, then flash cache on this disk will be disabled causing a reduction in the effective flash cache size.
- If the flash disk was used for Write-Back flash cache, then flash cache on this disk will be disabled, and read I/O to dirty data in the flash cache will go to the mirror disks on different cells. Write I/Os will go directly to the grid disks on this cell.
- If the flash disk was used for flash log, then flash log is disabled on this disk causing a reduction of the effective flash log size.
- If the flash disk was used for `writethrough` flash cache, then flash cache on this disk will be disabled causing a reduction in the effective flash cache size.

- If the flash disk was used for `writeback` flash cache, then flash cache on this disk will be disabled, and read I/O to dirty data in the flash cache will go to the mirror disks on different cells. Write I/Os will go directly to the grid disks on this cell.
- If the flash disk was used for flash log, then flash log is disabled on this disk causing a reduction of the effective flash log size.
- If the flash disk was used for grid disks, then the Oracle ASM disks associated with the grid disks on this flash disk are offline.

Confinement tests run on this disk, and either move the disk permanently to `poor performance` or return the disk back to `normal` state.

HALRT-10086

Cause: The flash disk has entered `peer failure` status.

Action: Another alert will be sent when all data has been rebalanced and flushed from the flash disk. Wait for the second alert before replacing the flash disk. A white cell locator LED has been lit to help you locate the affected server.

HALRT-10087

Cause: The data on this disk has been successfully rebalanced and flushed.

Action: Replace the drive.

C.14 Generic Actions for Alert Messages: HALRT-10503 to HALRT-11000

These messages are user actions for alert messages. The actions are in the range HALRT-10503 through HALRT-11000.

HALRT-10503

Cause: The HDD disk controller had a problem.

Action: HDD disk controller battery needs replacement. Contact Oracle Support Services.

HALRT-10504

Cause: The battery was in write through caching mode.

Action: HDD disk controller battery is good, but the logical drive is in write through caching mode. Contact Oracle Support

HALRT-10505

Cause: The disk controller battery completed the learn cycle.

Action: Disk controller battery is back to a good state, but the disk drive is in `writethrough` caching mode. Contact Oracle Support Services.

HALRT-10506

Cause: Disk drive entered `writeback` caching mode.

Action: Disk drive is in `writeback` caching mode. No action is needed.

HALRT-10507

Cause: All logical drives were in write through caching mode.

Action: The disk controller battery is good, but all logical drives are in write-through caching mode. Contact Oracle Support Services.

HALRT-10510

Cause: The contents of the disk controller cache have been lost.

Action: Contact Oracle Support Services.

The loss of the disk controller cache potentially leads to some loss of recent write operations. Access to the grid disks on this server has been disabled. The storage server software disk auto management has been disabled for the storage server, to avoid potential conflicts with user-initiated disk offline or drop force during the recovery period. Auto management is not affected on the servers that do not have this problem.

HALRT-10511

Cause: New battery was being charged.

Action: This is an informational message, and no action is required. The new battery is being charged automatically. When the battery completes charging, the disk drives will be in writeback caching mode.

HALRT-10512

Cause: Disk controller battery overheated.

Action: Disk controller battery is over-heated. Check if the ambient temperature is below 32 degrees Celsius and if the fans are working.

HALRT-10513

Cause: Disk controller battery was not functioning.

Action: Replace the disk controller battery.

C.15 Generic Actions for Alert Messages: HALRT-11001 to HALRT-11100

HALRT-11001

Cause: There was a problem with the InfiniBand configuration.

Action: Use `ibstatus` to check the InfiniBand configuration, and verify that both ports show state `ACTIVE` and the expected Gb per second. The alert is cleared when the port status is `ACTIVE`.

A common cause of the InfiniBand link negotiating to a lower link rate is that the InfiniBand cable connectors on the HCA or the switch are loose. Tightening the InfiniBand connector in the HCA and the switch usually solves this problem.

A port may show status `DOWN` after a switch restart. The port is automatically made `ACTIVE` after a few minutes.

 **See Also:**

My Oracle Support note 1612884.1 for additional information

HALRT-11002

Cause: There was a problem with the InfiniBand port.

Action: Use the CellCLI `LIST IBPORT DETAIL` command to check the port. Verify both ports show zero error counts for `rcvErrs` and `symbolErrs`. A non-zero error count can be caused by a bad cable or a cable that is not properly seated. It can also be caused by a bad port on the HCA or switch. Manually disabling and enabling the InfiniBand ports or resetting the InfiniBand switches can also lead to these errors. Correct the problem, and then reset the InfiniBand error counts on both the HCA port and the peer switch port by using the CellCLI `ALTER IBPORT...RESET` command.

HALRT-11003

Cause: There was a problem with the InfiniBand port.

Action: Use the CellCLI `LIST IBPORT DETAIL` command to check the port. Verify that both ports show the expected Gb per second. A common cause for the InfiniBand link negotiating to a lower link rate is that the InfiniBand cable connectors on the HCA or the switch are loose. Tightening the InfiniBand connector in the HCA and in the switch usually solves this problem.

D

Replacement Units

This appendix lists the replacement units for Recovery Appliance. It contains the following topics:

- [About Warranty Coverage](#) (page D-1)
- [Understanding Repair Categories](#) (page D-1)
- [LED Status Descriptions](#) (page D-2)
- [Parts for Compute Servers](#) (page D-6)
- [Parts for Storage Servers](#) (page D-20)
- [Parts for InfiniBand Switches](#) (page D-33)
- [Parts for Cisco Switch](#) (page D-34)
- [Parts for Power Distribution Units](#) (page D-34)
- [Cables and Cable Accessories for Recovery Appliance](#) (page D-35)

D.1 About Warranty Coverage

All replaceable parts are field replaceable units (FRUs) under the Oracle Premier Support for Systems warranty. Oracle field engineers install the replacement parts.

D.2 Understanding Repair Categories

Repair categories describe the potential impact of a repair action. [Table D-1](#) (page D-1) describes the categories and actions.

Table D-1 Recovery Appliance Repair Categories

Category	Description	Examples
Hot Swap (HS)	Repair part is hot swappable, and can be replaced without shutting down the host system. Commands might be required before and after replacement to protect data.	<ul style="list-style-type: none">• Disks• Fans• Power supplies
Infrastructure Repair (IR)	Repairing a connectivity component within the Recovery Appliance rack. No downtime of the rack is required; however, individual components may require downtime.	<ul style="list-style-type: none">• Ethernet switch• External cables• InfiniBand switch

Table D-1 (Cont.) Recovery Appliance Repair Categories

Category	Description	Examples
Compute Server Offline (DBO)	<p>Repairing the part requires shutting down a compute server. No downtime of the rack is required; however, individual servers might require downtime, and they might be taken outside of the cluster temporarily.</p> <p>If the server is currently running, then you must shut it down. See "Powering Down the Servers (page 12-4)".</p> <p>If you monitor the system using Oracle Enterprise Manager Grid Control, then avoid unnecessary alerts by putting the target in a blackout state. Refer to <i>Oracle Enterprise Manager Administration</i> for additional information.</p>	<ul style="list-style-type: none"> • Memory • PCIe cards • PDUs • Processors • System boards
Storage server Offline (EBO)	<p>Repairing the part requires shutting down a storage server within the rack. No downtime of the rack is required; however, individual servers might require downtime, and they might be taken outside of the cluster temporarily. There also might be a temporary performance impact, while the system rebalances to maintain data redundancy.</p> <p>If the server is currently running, then you must shut it down. See "Powering Down the Servers (page 12-4)".</p> <p>If you monitor the system using Oracle Enterprise Manager Grid Control, then avoid unnecessary alerts by putting the target in a blackout state. Refer to <i>Oracle Enterprise Manager Administration</i> for additional information.</p>	<ul style="list-style-type: none"> • Memory • PCIe cards • Processors • System boards

D.3 LED Status Descriptions

This section describes the LEDs.

D.3.1 Oracle Server LEDs

[Table D-2](#) (page D-3) identifies the LED status indicators that can indicate faults on the compute servers and the storage servers.

Table D-2 Oracle Server LED Status Descriptions

Component	LED Status
General Status	<p>The following LEDs are located on the front and back panels of each server:</p> <p>Service Required (amber) Off indicates normal operation.</p> <p>Steady on indicates a fault in the server.</p> <p>Power/OK (green) Steady on indicates that the server is powered on.</p> <p>A slow blink indicates that the server is starting up.</p> <p>A steady blink indicates that standby power is on, but chassis power is off.</p> <p>Off indicates no AC power or an incomplete startup of Oracle ILOM.</p>
Fans	<p>Each fan module has one bicolored LED, which you can see when the top cover fan door is open.</p> <p>Fan Status (amber/green) Amber indicates that the fan module is faulty. The front and rear panel Service Required LEDs are also lit when the system detects a fan module fault.</p> <p>Green indicates that the system is powered on, and the fan module is functioning correctly.</p>
Drives	<p>The following LEDs are located on the storage and boot drives:</p> <p>Service Required (amber) Off indicates normal operation.</p> <p>Steady on indicates a fault in the drive.</p> <p>OK/Activity (green) Steady on indicates normal inactivity.</p> <p>A steady blink indicates normal activity.</p> <p>Off indicates that the power is off, or the system does not recognize the drive.</p>

Table D-2 (Cont.) Oracle Server LED Status Descriptions

Component	LED Status
Power supply	<p>The following LEDs are on each power supply. You can see them from the back of the server.</p> <p>Service Required (amber) Off indicates normal operation.</p> <p>Steady on indicates that the system has detected a problem in the power supply regarding the voltage, current, temperature, or fans.</p> <p>AC OK/DC OK (green) Steady on indicates normal operation for both AC and DC power.</p> <p>A slow blink indicates that input AC power is normal and output DC power is disabled.</p> <p>Off indicates no input AC power.</p>

 **See Also:**

- "About Server and Component Status Indicators" in the *Sun Server X4-2 Service Manual* at
http://docs.oracle.com/cd/E36975_01/html/E38045/z4000be71405942.html#scrolltoc
- "About Server and Component Status Indicators" in the *Sun Server X4-2L Service Manual* at
http://docs.oracle.com/cd/E36974_01/html/E38145/z40003ec1405942.html#scrolltoc

D.3.2 Sun Datacenter InfiniBand Switch 36 LEDs

Table D-3 (page D-5) identifies the LED status indicators that can indicate faults on Sun Datacenter InfiniBand Switch 36.

Table D-3 Sun InfiniBand Switch LED Status Descriptions

Component	LED Status
Chassis	<p>The following LEDs are located on the left side of the rear panel.</p> <p>Attention (amber) Steady on indicates a fault.</p> <p>Off indicates normal operation.</p> <p>Flashing indicates no function.</p> <p>OK (green) Steady on indicates normal operation.</p> <p>Off indicates that the switch is off or initializing.</p> <p>Flashing indicates no function.</p>
Link status	<p>The link status LEDs are located by the data cable connectors on the rear panel.</p> <p>Link (green) Steady on indicates that a link is established.</p> <p>Off indicates that the link is down.</p> <p>Flashing indicates that there are symbol errors.</p>
Network management ports	<p>The following LEDs are located on the network management connector on the rear panel.</p> <p>Link Speed (amber/green): Green steady on indicates a 1000BASE-T link.</p> <p>Amber steady on indicates a 100BASE-T link.</p> <p>Off indicates no link.</p> <p>Flashing indicates no function.</p> <p>Activity (green): Flashing indicates packet activity.</p> <p>On indicates no function.</p> <p>Off indicates no activity.</p>

Table D-3 (Cont.) Sun InfiniBand Switch LED Status Descriptions

Component	LED Status
Power supply	<p>The following LEDs are located on the power supply at the front of the chassis.</p> <p>Attention (amber) Steady on indicates a fault, and 12 VDC is shut down.</p> <p>Off indicates normal operation.</p> <p>Flashing indicates no function.</p> <p>OK (green) Steady on indicates 12 VDC.</p> <p>Off indicates no DC voltage.</p> <p>Flashing indicates no function.</p> <p>AC LED is green Steady on indicates AC power is supplied.</p> <p>Off indicates that AC power is not present.</p> <p>Flashing indicates no function.</p>
Fan	<p>The following LED is located on the air inlet on the front of the chassis.</p> <p>FT On indicates a problem with the fan.</p> <p>Off indicates normal operation</p>



See Also:

"Identifying LEDs" in the *Sun Datacenter InfiniBand Switch 36 Service Manual* at

http://docs.oracle.com/cd/E26698_01/html/E26434/z40000fe1394411.html

D.4 Parts for Compute Servers

The following topics identify the replacement parts for the Recovery Appliance compute servers:

- [Oracle Server X6-2 and X5-2 Parts](#) (page D-13)
- [Sun Server X4-2 Parts](#) (page D-18)

D.4.1 Oracle Server X7 Oracle Database Server Parts

The following image shows the front view of Oracle Server X7 Oracle Database Server.

Figure D-1 Front View of Oracle Server X7 Oracle Database Server



The following image shows the rear view of Oracle Server.

Figure D-2 Rear View of X7 Oracle Database Server



The following table lists the replaceable parts for X7 Oracle Server. Refer to this table only if you have a basic warranty. All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-4 Replacement Parts for Oracle Server X7 Oracle Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
350-1719	Assembly, slim-rail, 1U - 2U, snap in	DBO	CRU	<ol style="list-style-type: none"> 1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services (page 12-3)" for additional information. 2. See "Reinstall the Server Into the Rack" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqrpu.html#scrolltoc
370-5621	3 volt lithium battery	DBO	FRU	See "Servicing the Battery (CRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqviw.html#scrolltoc
7093015	Drive, 600 GB 10K, SAS2, SFF	DBO	FRU	See "Servicing Storage Drives (CRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gquak.html#scrolltoc
7322171	Cable, interlock, 125 mm	DBO	FRU	See "Servicing Other Cables" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvev.html#scrolltoc

Table D-4 (Cont.) Replacement Parts for Oracle Server X7 Oracle Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7046442	Assembly, Dual 40Gb/Sec (4x) QDR InfiniBand Host Channel Adapter Module M3	DBO	CRU	<ol style="list-style-type: none"> <li data-bbox="1224 373 1458 716">1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services (page 12-3)" for additional information. <li data-bbox="1224 730 1458 1024">2. See "Servicing PCIe Cards (CRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvjk.html#scrolltoc <li data-bbox="1224 1039 1458 1150">3. See "Configuring the InfiniBand Switches (page 10-3)" <li data-bbox="1224 1165 1458 1556">4. Verify the new hardware meets the hardware and firmware specifications using the following command: <pre data-bbox="1273 1423 1458 1556">/opt/ oracle.SupportT ools/ CheckHWnFWprof ile</pre> <p data-bbox="1224 1577 1458 1696">See Also: <i>Oracle Exadata System Software User's Guide</i></p>

Table D-4 (Cont.) Replacement Parts for Oracle Server X7 Oracle Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7051223	Assembly, Dual 25 GbE SFP28 LP, X8 PCIe 3.0, RoHS:Y	DBO	CRU	<ol style="list-style-type: none"> 1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services (page 12-3)" for additional information. 2. See "Servicing PCIe Cards (CRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvjk.html#scrolltoc
7328735	Heat sink, CPU 1U	DBO	FRU	See "Servicing Processors (FRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqtmk.html#scrolltoc
7343599	Assembly, riser, 1-slot, 1U	DBO	FRU	See "Servicing PCIe Risers (CRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvft.html#scrolltoc
7330698	DIMM, 32 GB, RDIMM, DDR4, 2666	DBO	CRU	See "Servicing the DIMMs (CRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvkr.html#scrolltoc

Table D-4 (Cont.) Replacement Parts for Oracle Server X7 Oracle Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7333459	Power supply, AC, A266, F, 12V, 1200W	HS	CRU	<ol style="list-style-type: none"> 1. Remove rear left side of the CMA, and then swing it back and out to the right. The telescopic cross member supporting the cables directly behind the server can be retracted to access both power supplies. 2. See "Servicing Power Supplies (CRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqunc.html#scrolltoc 3. Replace the CMA.
7332895	HBA, 12 Gb SAS PCIe 16-port RAID internal with flash module	DBO	FRU	See "Servicing the Internal HBA Card and HBA Super Capacitor (FRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvmq.html#scrolltoc
7086345	Assembly, super cap, 13.5 volt, 6.4F, SAS3 HBA	DBO	FRU	See "Servicing the Internal HBA Card and HBA Super Capacitor (FRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvmq.html#scrolltoc

Table D-4 (Cont.) Replacement Parts for Oracle Server X7 Oracle Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7340897	Cable, SAS3 to HBA	DBO	CRU	See "Servicing the Internal HBA Cables" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvav.html#scrolltoc
7331123	Quad fan module	HS	CRU	See "Servicing Fan Modules (CRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gquhg.html#scrolltoc
7095106	CPU, P-8160, 2.1 GHz, 150W, 24 core, (SR3B0)	DBO	FRU	See "Servicing Processors (FRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqtmk.html#scrolltoc
7341145	Assembly, 8-slot disk backplane SAS3, 1U	DBO	FRU	See "Servicing the Disk Backplane (FRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvap.html#scrolltoc

Table D-4 (Cont.) Replacement Parts for Oracle Server X7 Oracle Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7317636	Assembly, motherboard, 1U	DBO	FRU	<ol style="list-style-type: none"> 1. See "Servicing the Motherboard (FRU)" in <i>Oracle Server X7-2 Service Manual</i> at http://docs.oracle.com/cd/E72435_01/html/E72445/gqvfa.html#scrolltoc 2. Press F2 after restarting the server. 3. Set the boot order.

D.4.2 Oracle Server X6-2 and X5-2 Parts

Figure D-3 (page D-13) shows the front view of the Oracle Server X6-2 or X5-2.

Figure D-3 Front View of Oracle Server X6-2 or X5-2



Figure D-4 (page D-13) shows the rear view of the Oracle Server X6-2 or X5-2.

Figure D-4 Rear View of Oracle Server X6-2 or X5-2



Table D-5 (page D-14) lists the replaceable parts for Oracle Server X6-2 compute servers. Unless otherwise noted, the repair procedures are in *Oracle Server X6-2 Service Manual*.

 **Note:**

All replacement parts listed in the following table are FRUs when covered under the Premier Support of Systems warranty.

Table D-5 Replacement Parts for Oracle Server X6-2 Compute Servers

Part Number	Description	Repair Category	Repair Procedure
350-1719	Assembly, slim rail, 1U - 2U, snap in	DBO	<ol style="list-style-type: none"> 1. If the server is running, then shut it down gracefully. See "Shutting Down Recovery Appliance (page 12-3)". 2. See "Reinstall the Server Into the Rack" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40004191003899.html
370-5621	3 volt lithium battery	DBO	See "Servicing the Battery (CRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40003f01423753.html
542-0287	Drive, 600 GB 10K, SAS2, SFF	DBO	See "Servicing Storage Drives (CRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40000091011460.html
7014226	Cable, interlock, 125 mm	DBO	See "Servicing the Power, FIM, Disk Backplane Data, and DVD Cables" at http://docs.oracle.com/cd/E62159_01/html/E62171/gnyov.html
7046442	Assembly, dual port, QDR CX3, host channel adapter	DBO	<ol style="list-style-type: none"> 1. If the server is running, then shut it down gracefully. See "Shutting Down Recovery Appliance (page 12-3)". 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40000f91037409.html 3. Check the InfiniBand configuration. See "Verifying InfiniBand Network Configuration" at http://docs.oracle.com/cd/E50790_01/doc/doc.121/e51951/components.htm#DBMMN21155 4. Use the following command to verify that the new hardware conforms to the hardware and firmware requirements: <code>/opt/oracle.SupportTools/CheckHWnFWprofile</code>
7051223	Assembly, Dual 10 GbE SFP+LP, X8 PCIe 2.0, RoHS:Y	DBO	See <i>Sun Dual 10GbE SFP+ PCIe 2.0 Low Profile Adapter User's Guide</i> at http://docs.oracle.com/cd/E19407-01/820-7898-11/index.html See "Servicing PCIe Cards (CRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40000f91037409.html
7067728	Heat sink, CPU 1U	DBO	See "Servicing Processors (FRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40001d31037480.html
7077680	Assembly, riser, 2S, 1U	DBO	See "Servicing PCIe Risers (CRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40000f91037394.html

Table D-5 (Cont.) Replacement Parts for Oracle Server X6-2 Compute Servers

Part Number	Description	Repair Category	Repair Procedure
7078072	DIMM, 32 GB, RDIMM, DDR4, 2400	DBO	See "Servicing the DIMMs (CRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40003f01425075.html
7079395	Power supply, AC, A256, F, 12V, 600W	HS	<ol style="list-style-type: none"> 1. Remove the rear left side of the CMA, and then swing the CMA back and out to the right. To access both power supplies, retract the telescopic cross member supporting the cables directly behind the server. 2. See "Servicing the Power Supplies (CRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z4000091014153.html 3. Replace the CMA.
7081071	Assembly, riser, single stack, 1U	DBO	See "Servicing PCIe Risers (CRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40000f91037394.html
7085209	HBA, 12 Gb SAS PCIe 8-port RAID internal with flash module	DBO	See "Servicing the Internal HBA Card (FRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/goeht.html
7086345	Assembly, super cap, 13.5 volt, 6.4F, SAS3 HBA	DBO	See "Servicing the Internal HBA Card (FRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/goeht.html
7086346	Cable, super cap, 620M, 6-pin, straight, SAS3 HBA	DBO	See "Servicing the Internal HBA SAS Cable assembly" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40003eb1420346.html
7089073	Fan module	HS	See "Servicing Fan Modules (CRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z4000091014194.html
7095106	CPU, E5-2699 v4, 2.2 GHz, 145W, 22 core, C-1 (SR1XD)	DBO	See "Servicing Processors (FRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40001d31037480.html
7097205	Assembly, 8DBP, 1U/2U	DBO	See "Servicing the Disk Backplane (FRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40011771417642.html
7098505	Assembly, motherboard, 1U	DBO	<ol style="list-style-type: none"> 1. See "Servicing the Motherboard (FRU)" at http://docs.oracle.com/cd/E62159_01/html/E62171/z40011771418436.html 2. Press F2 after restarting the server. 3. Set the boot order.

Table D-6 (page D-16) lists the replaceable parts for Oracle Server X5-2 compute servers. Unless otherwise noted, the repair procedures are in *Oracle Server X5-2 Service Manual*.



Note:

All replacement parts listed in the following table are FRUs when covered under the Premier Support of Systems warranty.

Table D-6 Replacement Parts for Oracle Server X5-2 Compute Servers

Part Number	Description	Repair Category	Repair Procedure
350-1719	Assembly, slim rail, 1U - 2U, snap in	DBO	<ol style="list-style-type: none"> 1. If the server is running, then shut it down gracefully. See "Shutting Down Recovery Appliance (page 12-3)". 2. See "Reinstall the Server Into the Rack" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40004191003899.html
370-5621	3 volt lithium battery	DBO	See "Servicing the Battery" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40003f01423753.html
542-0287	Drive, 600 GB 10K, SAS2, SFF	DBO	See "Servicing Storage Drives (CRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40000091011460.html
7014226	Cable, interlock, 125 mm	DBO	See "Servicing the Power, FIM, Disk Backplane Data, and DVD Cables" at http://docs.oracle.com/cd/E41059_01/html/E48320/gnyov.html
7046442	Assembly, dual port, QDR CX3, host channel adapter	DBO	<ol style="list-style-type: none"> 1. If the server is running, then shut it down gracefully. See "Shutting Down Recovery Appliance (page 12-3)". 2. See "Servicing PCIe Cards (CRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40000f91037409.html 3. Check the InfiniBand configuration. See "Verifying InfiniBand Network Configuration" at http://docs.oracle.com/cd/E50790_01/doc/doc.121/e51951/components.htm#DBMMN21155 4. Use the following command to verify that the new hardware conforms to the hardware and firmware requirements: <code>/opt/oracle.SupportTools/CheckHWnFWprofile</code>

Table D-6 (Cont.) Replacement Parts for Oracle Server X5-2 Compute Servers

Part Number	Description	Repair Category	Repair Procedure
7051223	Assembly, Dual 10 GbE SFP+LP, X8 PCIe 2.0, RoHS:Y	DBO	See <i>Sun Dual 10GbE SFP+ PCIe 2.0 Low Profile Adapter User's Guide</i> at http://docs.oracle.com/cd/E19407-01/820-7898-11/index.html See "Servicing PCIe Cards (CRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40000f91037409.html
7067728	Heat sink, CPU 1U	DBO	See "Servicing Processors (FRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40001d31037480.html
7077680	Assembly, riser, 2S, 1U	DBO	See "Servicing PCIe Risers (CRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40000f91037394.html
7078072	DIMM, 32 GB, LRDIMM, DDR4, 2133	DBO	See "Servicing the DIMMs (CRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40003f01425075.html
7079395	Power supply, AC, A256, F, 12V, 600W	HS	<ol style="list-style-type: none"> 1. Remove the rear left side of the CMA, and then swing the CMA back and out to the right. To access both power supplies, retract the telescopic cross member supporting the cables directly behind the server. 2. See "Servicing Power Supplies (CRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40000091014153.html 3. Replace the CMA.
7081071	Assembly, riser, single stack, 1U	DBO	See "Servicing PCIe Risers (CRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40000f91037394.html
7085209	HBA, 12 Gb SAS PCIe 8-port RAID internal with flash module	DBO	See "Servicing the Internal HBA Card (FRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/goeht.html
7086345	Assembly, super cap, 13.5 volt, 6.4F, SAS3 HBA	DBO	See "Servicing the Internal HBA Card (FRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/goeht.html
7086346	Cable, super cap, 620M, 6-pin, straight, SAS3 HBA	DBO	See "Servicing the Internal HBA SAS Cable assembly" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40003eb1420346.html
7089073	Fan module	HS	See "Servicing Fan Modules (CRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40000091014194.html

Table D-6 (Cont.) Replacement Parts for Oracle Server X5-2 Compute Servers

Part Number	Description	Repair Category	Repair Procedure
7095106	CPU, E5-2699 v3, 2.3 GHz, 145W, 18 core, C-1 (SR1XD)	DBO	See "Servicing Processors (FRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40001d31037480.html
7097205	Assembly, 8DBP, 1U/2U	DBO	See "Servicing the Disk Backplane (FRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40011771417642.html
7098505	Assembly, motherboard, 1U	DBO	<ol style="list-style-type: none"> 1. See "Servicing the Motherboard (FRU)" at http://docs.oracle.com/cd/E41059_01/html/E48320/z40011771418436.html 2. Press F2 after restarting the server. 3. Set the boot order.

D.4.3 Sun Server X4-2 Parts

Figure D-5 (page D-18) shows the front view of the Sun Server X4-2.

Figure D-5 Front View of Sun Server X4-2



Figure D-6 (page D-18) shows the rear view of the Sun Server X4-2.

Figure D-6 Rear View of Sun Server X4-2



Table D-7 (page D-19) lists the replaceable parts for Sun Server X4-2 compute servers. Unless otherwise noted, the repair procedures are in *Sun Server X4-2 Service Manual*.

 **Note:**

All replacement parts listed in the following table are FRUs when covered under the Premier Support of Systems warranty.

Table D-7 Replacement Parts for Sun Server X4-2 Compute Servers

Part Number	Description	Repair Category	Repair Procedure
7013526	Fan module	HS	See "Servicing Fan Modules" in <i>Sun Server X4-2 Service Manual</i> at http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091014194.html#scrolltoc
7018701	DIMM, 16 GB, DDR3, 2RX4, 1600, 1.35 V	DBO	See "Servicing the DIMMs" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40003f01425075.html#scrolltoc
7020774	Heat sink, CPU 1U, 60 mm FIN	DBO	See "Servicing Storage Drives (CRU)" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091011460.html#scrolltoc
7023303	Dual 16 Gigabit/Sec Fiber Channel / 10 GbS FCoE Host Adapter, QLogic		See "Install a PCIe Card in PCIe Slot 1 or 2" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40007611438967.html#scrolltoc
7027478	Disk Backplane, RoHS:Y	DBO	See "Servicing the Disk Backplane (FRU)" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40011771417642.html#scrolltoc
7039841	Front indicator module	DBO	See "Servicing PCIe Risers" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40011771417642.html#scrolltoc
7046442	Assembly, dual port, QDR CX3, host channel adapter	DBO	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down. See "Powering Down the Servers (page 12-4)". 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40000f91037409.html#scrolltoc 3. See "Connecting Recovery Appliance to Your Networks (page 10-3)". 4. Use the following command to verify that the new hardware meets the hardware and firmware requirements: <code>/opt/oracle.SupportTools/CheckHwnFWprofile</code>
7047503	Assembly, 6 GB, SAS RAID disk PCIe HBA, RoHS:Y	DBO	See "Servicing Storage Drives" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091011460.html#scrolltoc
7048707	Assembly, riser, double stack	DBO	See "Servicing PCIe Risers" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40011771417642.html#scrolltoc
7048710	1-slot PCI Express riser assembly	DBO	See "Servicing PCIe Risers" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40011771417642.html#scrolltoc

Table D-7 (Cont.) Replacement Parts for Sun Server X4-2 Compute Servers

Part Number	Description	Repair Category	Repair Procedure
7051223	Assembly, Dual 10 GbE SFP+LP, X8 PCIe 2.0, RoHS:Y	DBO	See <i>Sun Dual 10GbE SFP+ PCIe 2.0 Low Profile Adapter User's Guide</i> at http://docs.oracle.com/cd/E19407-01/820-7898-11/index.html
7057184	Assembly, battery carrier remote	HS	" Replacing a Disk Controller BBU on a Compute Server (page 12-44)"
7058153	Assembly, motherboard, X4-2	DBO	<ol style="list-style-type: none"> 1. See "Servicing the Motherboard" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40011771418436.html#scrolltoc 2. Press F2 after restarting the server. 3. Set the boot order.
7060020	Universal upgrade remote battery kit	EBO	" Replacing a Disk Controller BBU on a Compute Server (page 12-44)"
7060951	Power supply, AC, A256, F, 12V, 600W	HS	<ol style="list-style-type: none"> 1. Remove rear left side of the CMA, and then swing the CMA back and out to the right. To access both power supplies, retract the telescopic cross member supporting the cables directly behind the server. 2. See "Servicing the Power Supply" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091014153.html#scrolltoc 3. Replace the CMA.
7069232	CPU, E5-2697 v2, 12-core, 2.7 GHz	DBO	See "Servicing Processors" at http://docs.oracle.com/cd/E36975_01/html/E38045/z40001d31037480.html#scrolltoc



See Also:

Sun Server X4-2 Service Animations at

http://docs.oracle.com/cd/E36975_01/html/E49470/index.html

D.5 Parts for Storage Servers

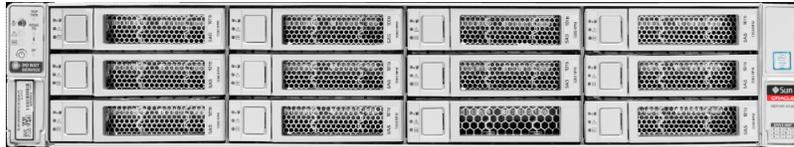
The following topics identify the replacement parts for the Recovery Appliance storage servers:

- [Oracle Server X6-2L and X5-2L Parts](#) (page D-24)
- [Sun Server X4-2L Parts](#) (page D-30)

D.5.1 Oracle Exadata Storage Server X7-2 High Capacity Server Parts

The following figure shows the front view of the Oracle Exadata Storage Server X7-2 High Capacity Server.

Figure D-7 Front View of Oracle Exadata Storage Server X7-2 High Capacity Server



The following figure shows the rear view of the Oracle Exadata Storage Server X7-2 High Capacity Server.

Figure D-8 Rear View of Oracle Exadata Storage Server X7-2 High Capacity Server



The following table lists the replaceable parts for Oracle Exadata Storage Server X7-2 High Capacity Servers.

 **Note:**

All replacement parts shown in the following table are considered FRU when covered under Premier Support of Systems warranty.

Table D-8 Replacement Parts for Oracle Exadata Storage Server X7-2 High Capacity Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
350-1719	Assembly, 1U-2U Snap-in, slim rail rack	EBO	CRU	<ol style="list-style-type: none"> 1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services (page 12-3)" for additional information. 2. See "Reinstall the Server Into the Rack" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqqa.html#scrolltoc
370-5621	3 volt lithium battery	EBO	FRU	See "Servicing the Battery" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtf.html#scrolltoc
7322171	Indicator assembly, LFIM	EBO	FRU	See "Servicing the Front LED Indicator Module" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtpg.html#scrolltoc
7341141	12-Disk backplane assembly	EBO	FRU	See "Servicing the Disk Backplane" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtsb.html#scrolltoc
7333459	Power Supply, A266, 1200W, AC Input	HS	FRU	<ol style="list-style-type: none"> 1. Remove rear left side of the CMA, and then swing it back and out to the right. The telescopic cross member supporting the cables directly behind the server can be retracted to access both power supplies. 2. See "Servicing Power Supplies" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtag.html#scrolltoc 3. Replace the CMA.
7322166	Fan assembly, 2U	HS	FRU	See "Servicing Fan Modules" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtak.html#scrolltoc

Table D-8 (Cont.) Replacement Parts for Oracle Exadata Storage Server X7-2 High Capacity Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7046442	Assembly, dual port, QDR CX3, host channel adapter	EBO	FRU	<ol style="list-style-type: none"> 1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services (page 12-3)" for additional information. 2. See "Servicing PCIe Cards" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtcq.html#scrolltoc 3. See "Configuring the InfiniBand Switches (page 10-3)" 4. Verify the new hardware meets the hardware and firmware specifications using the following command: <pre>/opt/oracle.SupportTools/ CheckHWnFWprofile</pre>
7322160	Heat sink, CPU	EBO	FRU	See "Servicing Processors (FRU)" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtmk.html#scrolltoc
7332759	Assembly, drive, 10 TB, 3.5 inch, SAS/ 7200rpm	HS	FRU	See "Servicing Storage Drives" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtcs.html#scrolltoc
7330697	DIMM, 16 GB, RDIMM, DDR4, 2666	EBO	FRU	See "Servicing the DIMMS" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtcn.html#scrolltoc
7332895	HBA, 12 Gb SAS PCIe 16-port RAID internal with flash module	EBO	FRU	See "Servicing an Internal HBA Card and HBA Super Capacitor" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtjv.html#scrolltoc
7086345	Assembly, super cap, 13.5 volt, 6.4F, SAS3 HBA	EBO	FRU	See "Servicing an Internal HBA Card and HBA Super Capacitor" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtjv.html#scrolltoc

Table D-8 (Cont.) Replacement Parts for Oracle Exadata Storage Server X7-2 High Capacity Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7315464	Cable, super cap, 620M, 6-pin, straight, SAS3 HBA	EBO	FRU	See "Servicing SAS Cables" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqthb.html#scrolltoc
7341587	M.2,150GB,SATA, 22x80mm,SSD,INTEL , DC S3520	EBO	FRU	See "Servicing the Internal M.2 Flash SSDs" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/grlly.html#scrolltoc
7335943	Assembly, 6.4 TB eMLC flash F640 NVMe card	EBO	FRU	See "Servicing PCIe Cards" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtqc.html#scrolltoc
7347511	CPU, XEON S-4114 , 2.2G, 85W, 10 CORE, (SR3GK)	EBO	FRU	See "Servicing Processors (FRU)" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtmk.html#scrolltoc
7335280	ASM, MB MECH, 2U	EBO	FRU	See "Servicing the Motherboard Assembly" in <i>Oracle Exadata Storage Server X7-2 EF and HC Service Manual</i> at http://docs.oracle.com/cd/E72463_01/html/E72479/gqtwc.html#scrolltoc

D.5.2 Oracle Server X6-2L and X5-2L Parts

Figure D-9 (page D-24) shows the front view of the Oracle Server X6-2L or X5-2L. The hard drives are numbered from left to right, starting in the lower left. The drives in the bottom row are numbers 0, 1, 2, and 3. The drives in the middle row are numbers 4, 5, 6, and 7. The drives in the top row are numbers 8, 9, 10, and 11.

Figure D-9 Front View of Oracle Server X6-2L or X5-2L

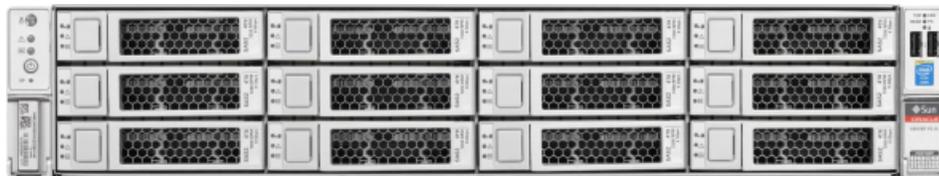


Figure D-10 (page D-25) shows the rear view of an Oracle Server X6-2L or X5-2L storage server.

Figure D-10 Rear View of Oracle Server X6-2L or X5-2L



Table D-9 (page D-25) lists the replaceable parts for the Oracle Server X6-2L. Unless otherwise noted, the repair procedures are in *Oracle Server X6-2L Service Manual*.

 **Note:**

All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table D-9 Replacement Parts for Oracle Server X6-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
350-1719	Assembly, 1U-2U Snap-in, slim rail rack	EBO	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down gracefully. See "Shutting Down Recovery Appliance (page 12-3)". 2. See "Reinstall the Server in the Rack" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40004191003899.html
370-5621	3 volt lithium battery	EBO	See "Servicing the Battery (CRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40000f91037456.html
7014226	Cable, interlock, 125 mm	EBO	See "Servicing SAS Cables (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/baecfjbc.html
7017231	Indicator assembly, right	EBO	See "Servicing the Front LED/USB Indicator Modules (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40001d31014236.html
7023515	Disk backplane assembly, 2 rear	EBO	See "Servicing the Front and Rear Storage Drive Backplanes (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40007ac1419009.html

Table D-9 (Cont.) Replacement Parts for Oracle Server X6-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
7044130	Power Supply, A258, 1000W, AC Input	HS	<ol style="list-style-type: none"> 1. Remove the rear left side of the CMA, and then swing the CMA back and out to the right. To access both power supplies, retract the telescopic cross member supporting the cables directly behind the server. 2. See "Servicing Power Supplies (CRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40000091014153.html 3. Replace the CMA.
7045042	Fan assembly, 2U, 80 x 80, Ferrite, center rotate	HS	See "Servicing Fan Modules (CRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40000091014194.html
7046442	Assembly, dual port, QDR CX3, host channel adapter	EBO	<ol style="list-style-type: none"> 1. If the server is running, then shut it down gracefully. See "Shutting Down Recovery Appliance (page 12-3)". 2. See "Servicing PCIe Cards (CRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40000f91037409.html 3. Use the following command to verify that the new hardware meets the hardware and firmware requirements: <code>/opt/oracle.SupportTools/CheckHWnFWprofile</code>
7064884	Heat sink, CPU	EBO	See "Servicing Processors (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40001d31037480.html
7066824	Assembly, drive, 8 TB, 3.5 inch, SAS/7200rpm	HS	See "Servicing Storage Drives and Rear Drives (CRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40000091011460.html
7069337	Assembly, backplane, 12X3.5, W/SAS expansion, 2U	EBO	See "Servicing the Front and Rear Storage Drive Backplanes (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40007ac1419009.html
7078070	DIMM, 16 GB, RDIMM, DDR4, 2400	EBO	See "Servicing the DIMMS (CRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/ceicjagi.html
7085209	HBA, 12 Gb SAS PCIe 8-port RAID internal with flash module	EBO	See "Servicing the Oracle Storage 12 Gb/s SAS PCIe RAID HBA Card (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/gocfy.html

Table D-9 (Cont.) Replacement Parts for Oracle Server X6-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
7086345	Assembly, super cap, 13.5 volt, 6.4F, SAS3 HBA	EBO	See "Servicing the Oracle Storage 12 Gb/s SAS PCIe RAID HBA Card (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/gocfy.html
7086346	Cable, super cap, 620M, 6-pin, straight, SAS3 HBA	EBO	See "Servicing SAS Cables (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/baecfjbc.html
7090170	Flash drive, MLC, OSA ID, USB, 8 GB	EBO	See "Servicing the Internal USB Flash Drives (CRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/ceicfbfd.html
7090698	Assembly, 3.2 TB eMLC flash F320 NVMe card	EBO	See <i>Oracle Flash Accelerator F320 PCIe Card User Guide</i> at http://docs.oracle.com/cd/E65386_01/html/E65387/goiip.html
7095102	CPU, E5-2630 V4, 2.2 G, 85 W, 10 CORE, R-2 (SR206)	EBO	See "Servicing Processors (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40001d31037480.html
7096185	Assembly, indicator, left	EBO	See "Servicing the Front LED/USB Indicator Modules (FRU)" at http://docs.oracle.com/cd/E62172_01/html/E62184/z40001d31014236.html
7098504	ASM, MOBO tray	EBO	

Table D-10 (page D-28) lists the replaceable parts for the Oracle Server X5-2L. Unless otherwise noted, the repair procedures are in *Oracle Server X5-2L Service Manual*.

 **Note:**

All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table D-10 Replacement Parts for Oracle Server X5-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
350-1719	Assembly, 1U-2U Snap-in, slim rail rack	EBO	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down gracefully. See "Shutting Down Recovery Appliance (page 12-3)". 2. See "Reinstall the Server in the Rack" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40004191003899.html
370-5621	3 volt lithium battery	EBO	See "Servicing the Battery (CRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40000f91037456.html
7014226	Cable, interlock, 125 mm	EBO	See "Servicing SAS Cables (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/baecfjbc.html
7017231	Indicator assembly, right	EBO	See "Servicing the Front LED/USB Indicator Modules (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40001d31014236.html
7023515	Disk backplane assembly, 2 rear	EBO	See "Servicing the Front and Rear Storage Drive Backplanes (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40007ac1419009.html
7044130	Power Supply, A258, 1000W, AC Input	HS	<ol style="list-style-type: none"> 1. Remove the rear left side of the CMA, and then swing the CMA back and out to the right. To access both power supplies, retract the telescopic cross member supporting the cables directly behind the server. 2. See "Servicing Power Supplies (CRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40000091014153.html 3. Replace the CMA.
7045042	Fan assembly, 2U, 80 x 80, Ferrite, center rotate	HS	See "Servicing Fan Modules (CRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40000091014194.html
7046442	Assembly, dual port, QDR CX3, host channel adapter	EBO	<ol style="list-style-type: none"> 1. If the server is running, then shut it down gracefully. See "Shutting Down Recovery Appliance (page 12-3)". 2. See "Servicing PCIe Cards (CRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40000f91037409.html 3. Use the following command to verify that the new hardware meets the hardware and firmware requirements: <code>/opt/oracle.SupportTools/CheckHwFWprofile</code>

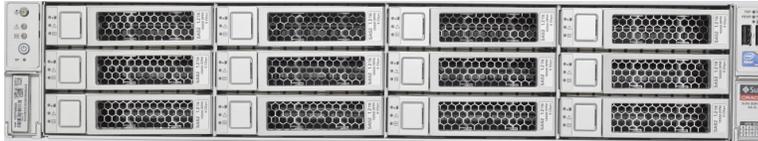
Table D-10 (Cont.) Replacement Parts for Oracle Server X5-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
7064884	Heat sink, CPU	EBO	See "Servicing Processors (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40001d31037480.html
7066824	Assembly, drive, 4 TB, 3.5 inch, SAS/7200rpm	HS	See "Servicing Storage Drives and Rear Drives (CRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40000091011460.html
7069337	Assembly, backplane, 12X3.5, W/SAS expansion, 2U	EBO	See "Servicing the Front and Rear Storage Drive Backplanes (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40007ac1419009.html
7078070	DIMM, 8 GB, RDIMM, DDR4, 2133	EBO	See "Servicing the DIMMS (CRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/ceicjagi.html
7085209	HBA, 12 Gb SAS PCIe 8-port RAID internal with flash module	EBO	See "Servicing the Oracle Storage 12 Gb/s SAS PCIe RAID HBA Card (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/gocfy.html
7086345	Assembly, super cap, 13.5 volt, 6.4F, SAS3 HBA	EBO	See "Servicing the Oracle Storage 12 Gb/s SAS PCIe RAID HBA Card (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/gocfy.html
7086346	Cable, super cap, 620M, 6-pin, straight, SAS3 HBA	EBO	See "Servicing SAS Cables (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/baecfjbc.html
7090170	Flash drive, MLC, OSA ID, USB, 8 GB	EBO	See "Servicing the Internal USB Flash Drives (CRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/ceicfbfd.html
7090698	Assembly, 1.6 TB eMLC flash F160 NVMe card, Aura 3.0 with LP BRKT	EBO	See <i>Oracle Flash Accelerator F160 PCIe Card User Guide</i> at http://docs.oracle.com/cd/E54943_01/html/E54947/index.html
7095102	CPU, E5-2630 V3, 2.4G, 85W, 8 CORE, R-2 (SR206)	EBO	See "Servicing Processors (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40001d31037480.html
7096185	Assembly, indicator, left	EBO	See "Servicing the Front LED/USB Indicator Modules (FRU)" at http://docs.oracle.com/cd/E41033_01/html/E48333/z40001d31014236.html
7098504	ASM, MOBO tray	EBO	

D.5.3 Sun Server X4-2L Parts

[Figure D-11](#) (page D-30) shows the front view of the Sun Server X4-2L. The hard drives are numbered from left to right, starting in the lower left. The drives in the bottom row are numbers 0, 1, 2, and 3. The drives in the middle row are numbers 4, 5, 6, and 7. The drives in the top row are numbers 8, 9, 10, and 11.

Figure D-11 Front View of Sun Server X4-2L



[Figure D-12](#) (page D-30) shows the rear view of a Sun Server X4-2L storage server.

Figure D-12 Rear View of Sun Server X4-2L



[Table D-11](#) (page D-30) lists the replaceable parts for the Sun Server X4-2L. Unless otherwise noted, the repair procedures are in *Sun Server X4-2L Service Manual*.

 **Note:**

All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table D-11 Replacement Parts for Sun Server X4-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
371-5002	USB drive, 4 GB	EBO	See "Servicing the Internal USB Flash Drives" in <i>Sun Server X4-2L Service Manual</i> at http://docs.oracle.com/cd/E36974_01/html/E38145/ceiicfbf.html#scrolltoc
7014226	Cable, interlock, 125 mm	EBO	See "Servicing Cables" at http://docs.oracle.com/cd/E36974_01/html/E38145/baecfjbc.html#scrolltoc

Table D-11 (Cont.) Replacement Parts for Sun Server X4-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
7014391	Expander assembly, 2U	EBO	See "Servicing the SAS Expander Module" at http://docs.oracle.com/cd/E36974_01/html/E38145/ceieidja.html#scrolltoc
7017231	Indicator assembly, right	EBO	See "Servicing the Front LED/USB Indicator Modules" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40001d31014236.html#scrolltoc
7018701	DIMM, 16 GB, DDR3, 1600, 2Rx4, 1.35V	EBO	See "Servicing the DIMMS" at http://docs.oracle.com/cd/E36974_01/html/E38145/ceicjagi.html#scrolltoc
7020486	DIMM, 8 GB, DDR3, 1600, 2Rx4, 1.35V	EBO	See "Servicing the DIMMS" at http://docs.oracle.com/cd/E36974_01/html/E38145/ceicjagi.html#scrolltoc
7020776	Heatsink, CPU, 2U, 60 MM FIN, RoHS:YU	EBO	See "Servicing Storage Drives and Rear Drives" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000091011460.html#scrolltoc
7023515	Disk backplane assembly, 2 rear	EBO	See "Servicing the Front and Rear Storage Drive Backplanes" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40007ac1419009.html#scrolltoc
7027479	Disk backplane assembly, 4-Slot 3.5" RoHS:Y	EBO	See "Servicing the Front and Rear Storage Drive Backplanes" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40007ac1419009.html#scrolltoc
7044130	Power Supply, A258, 1000 Watt AC Input	HS	<ol style="list-style-type: none"> 1. Remove rear left side of the CMA, and then swing the CMA back and out to the right. To access both power supplies, retract the telescopic cross member supporting the cables directly behind the server. 2. See "Servicing Power Supplies" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000091014153.html#scrolltoc 3. Replace the CMA.
7045042	Fan assembly, 2U, 80 x 80, Ferrite, center_rotate	HS	See "Servicing Fan Modules" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000091014194.html#scrolltoc

Table D-11 (Cont.) Replacement Parts for Sun Server X4-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
7046442	Assembly, dual port, QDR CX3, host channel adapter	EBO	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down. See "Powering Down the Servers (page 12-4)". 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000f91037409.html#scrolltoc 3. See "Connecting Recovery Appliance to Your Networks (page 10-3)". 4. Use the following command to verify that the new hardware meets the hardware and firmware requirements: <code>/opt/oracle.SupportTools/CheckHWnFWprofile</code>
7047503	Assembly, 6 GB, SAS RAID disk PCIe HBA, RoHS:Y	DBO	See "Servicing Storage Drives and Rear Drives" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000091011460.html#scrolltoc
7048864	Indicator assembly, left, 2U	EBO	See "Servicing the Front LED/USB Indicator Modules (FRU)" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40001d31014236.html#scrolltoc
7057184	Assembly, battery carrier remote	HS	" Replacing a Disk Controller BBU on a Storage Server (page 12-45)"
7058152	Assembly, motherboard, X4-2L	EBO	See "Servicing the Motherboard Assembly" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40001d31037512.html#scrolltoc
7060020	Universal upgrade remote battery kit	EBO	" Replacing a Disk Controller BBU on a Storage Server (page 12-45)"
7069200	Assembly, 800 GB eMLC flash PCIe card	EBO	See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000f91037409.html#scrolltoc
7069242	CPU, E5-2630 v2, 6-core, 2.6 Ghz	EBO	See "Servicing Processors" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40001d31037480.html#scrolltoc
7070656	Disk assembly, 1.2 TB, SAS2	HS	See "Servicing Storage Drives and Rear Drives" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000091011460.html#scrolltoc
7106353	4 TB 7200 rpm SAS HDD with bracket	DBO	See "Servicing Storage Drives and Rear Drives" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000091011460.html#scrolltoc

D.6 Parts for InfiniBand Switches

Table D-12 (page D-33) lists the replaceable parts for the InfiniBand switches used in Recovery Appliance. See "[Cables and Cable Accessories for Recovery Appliance](#) (page D-35)" for information about InfiniBand cables.

 **Note:**

All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table D-12 Replacement Parts for InfiniBand Switches

Part Number	Description	Repair Category	Repair
300-2143	760 watt power supply, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
300-2299	760 watt AC input power supply	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
350-1312	Fan module for Sun Datacenter InfiniBand Switch 36, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
350-1566	Fan module, 1 RU, Sanace, reverse	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
371-2210	CR2032 3 v battery, RoHS:Y	IR	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
541-3495	Sun Datacenter InfiniBand Switch 36, RoHS:YL	IR	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
7052970	Sun Datacenter InfiniBand Switch 36, RoHS:YL	IR	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
7061031	Fan module, 1 RU, Sanace, reverse	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html

D.7 Parts for Cisco Switch

Table D-13 (page D-34) lists replaceable parts for the Cisco switch used with Recovery Appliance.

 **Note:**

All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table D-13 Replacement Parts for Cisco Switch

Part Number	Description	Repair Category	Repair
7023685	Cisco Catalyst 4948 switch, RoHS:Y	IR	See chapter 3 in <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
7024423	Power supply for Cisco Catalyst 4948 switch, RoHS:Y	HS	See chapter 4 in <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
7024424	Cooling fan for Cisco Catalyst 4948 switch, RoHS:Y	HS	See chapter 4 in <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html

D.8 Parts for Power Distribution Units

Table D-14 (page D-35) lists the replacement parts for the power distribution units (PDUs). Given the number of cables and connections, Oracle recommends that you shut down the Recovery Appliance rack when replacing power distribution unit parts.

 **Note:**

All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table D-14 Replacement Parts for Power Distribution Units

Part Number	Description	Repair Category	Repair
180-2379	13A/250V 90 degree IEC 320 Sheet E to IEC 320 C13 cord, 1 meter, RoHS:Y	HS	Not applicable
180-2380	13A/250V 90 degree IEC 320 Sheet E to IEC 320 C13 cord, 2 meter, RoHS:Y	HS	Not applicable
371-4000	Three-phase PDU: 15kV,5-Pin, IEC309 32A, 5 PIN 230/400V, 3ph IP44, International, RoHS:Y	IR	See " Replacing a Power Distribution Unit (page 12-9) " of this guide.
371-4199	Single-phase PDU: 15kVA, with three 30A/250V 2-Pole/3-Wire NEMA L6-30P Plugs, North America, RoHS:Y	IR	See " Replacing a Power Distribution Unit (page 12-9) " of this guide.
371-4201	Single-phase PDU: 15kVA, with three Blue 32A/240V Splash-proof 2-Pole/3-Wire IEC 60309 Plugs, International, RoHS:Y	IR	See " Replacing a Power Distribution Unit (page 12-9) " of this guide.
371-5039	Single-phase PDU: 22k VA Hubbell CS8265C PlugNo X-Option	IR	See " Replacing a Power Distribution Unit (page 12-9) " of this guide.
371-5040	Single-phase PDU: 22k VA Blue 32A IEC 60309 Plug No X-Option	IR	See " Replacing a Power Distribution Unit (page 12-9) " of this guide.
7051908	Three-phase PDU: 15kVA, IEC309 60A, 4 PIN, 250VAC, 3ph IP67, North America, RoHS:Y	IR	See " Replacing a Power Distribution Unit (page 12-9) " of this guide.

D.9 Cables and Cable Accessories for Recovery Appliance

[Table D-15](#) (page D-36) lists the replacement cables and cable accessories for Recovery Appliance.

 **Note:**

All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table D-15 Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
350-1287	Slim rail, CMA, 1U-2U for a storage server	EBO	<ol style="list-style-type: none"> 1. If the server is currently running, then it should be shut down gracefully. See "Powering On and Off a Recovery Appliance Rack (page 12-2)" for additional information. 2. See "Installing the Cable Management Arm" in <i>Sun Fire X4170, X4270, and X4275 Servers Installation Manual</i> at http://docs.oracle.com/cd/E19477-01/820-5827-13/rack_install.html#50614309_99582
350-1546	Cable management arm (CMA) 1U, RoHS:Y for compute server	DBO	<ol style="list-style-type: none"> 1. If the server is currently running, then it should be shut down gracefully. See "Powering On and Off a Recovery Appliance Rack (page 12-2)" for additional information. 2. See "Installing the Cable Management Arm" in <i>Sun Fire X4170, X4270, and X4275 Servers Installation Manual</i> at http://docs.oracle.com/cd/E19477-01/820-5827-13/rack_install.html#50614309_99582
350-1519	Serial cable kit, RoHS:Y with the following: <ul style="list-style-type: none"> • USB to DB 9-pin M serial cable • DB 9-pin F to DB 9-pin F null modem cable 	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
350-1637	Label assembly, EV2, 2-way, upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4 (page F-1)
350-1638	Label assembly EV2, 3-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4 (page F-1)
350-1639	Label assembly EV2, 4-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4 (page F-1)
350-1640	Label assembly EV2, 5-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4 (page F-1)
350-1641	Label assembly EV2, 6-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4 (page F-1)
350-1642	Label assembly EV2, 7-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4 (page F-1)
350-1643	Label assembly EV2, 8-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4 (page F-1)

Table D-15 (Cont.) Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
530-4445	QSFP InfiniBand passive copper cable, 10 GbE, 3 m, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html Note: There is a loss of redundancy when replacing cables.
530-4446	QSFP InfiniBand passive copper cable, 10 GbE, 5 m, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html Note: There is a loss of redundancy when replacing cables.
530-4403	4x QSFP copper QDR InfiniBand cable, 2 m, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html Note: There is a loss of redundancy when replacing cables.
530-4404	4x QSFP copper QDR InfiniBand cable, 3 m, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
530-4415	4x QSFP copper QDR InfiniBand cable, 5 m, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
530-4432	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, blue, RoHS:Y	HS	See <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4433	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 10 foot, red, RoHS:Y	HS	See <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4434	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, red, RoHS:Y	HS	See <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4435	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 10 foot, black, RoHS:Y	HS	See <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html

Table D-15 (Cont.) Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
530-4436	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, black, RoHS:Y	HS	See <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4437	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, orange, RoHS:Y	HS	See <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4438	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 10 foot, black, RoHS:Y	HS	See <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
530-4444	QSFP to QSFP Passive Copper Cable, 1 m, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
530-4445	QSFP passive copper cable, 10 GE, 3 m	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html
530-4502-01	Ethernet cable, Cat 5/5E, 10 foot, blue, RoHS:Y	HS	See <i>Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> at http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html
7015400	Power cable, jumper, SR2, 2m C14RA, 10A, C13	HS	See "Replacing a Power Distribution Unit (page 12-9)" of this guide.
7042273	Slim rail, cable management arm, 1U-2U for a storage server	EBO	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down. See "Shutting Down Recovery Appliance (page 12-3)". 2. See "Installing the Second-Generation Cable Management Arm" in <i>Sun Server X4-2L Service Manual</i> at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000151427029.html#scrolltoc
7102869	Passive fiber InfiniBand cable, 10 m, RoHS:Y	HS	<ol style="list-style-type: none"> 1. See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html <p>Note: Purchase two QSFP SR Transceivers (Part Number X2124A-N) for each passive fiber cable.</p>

Table D-15 (Cont.) Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
7102870	Passive fiber InfiniBand cable, 20 m, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html Note: Purchase two QSFP SR Transceivers (Part Number X2124A-N) for each passive fiber cable.
7102871	Passive fiber InfiniBand cable, 50 m, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> at http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html Note: Purchase two QSFP SR Transceivers (Part Number X2124A-N) for each passive fiber cable.

E

In-Rack Cabling Tables for Recovery Appliance X7, X6, and X5

This appendix identifies the location of the various components in a Recovery Appliance X7, X6, and X5 full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

This appendix contains the following sections:

- [Cable Color Coding](#) (page E-1)
- [Recovery Appliance X7, X6, and X5 Rack Layout](#) (page E-1)
- [Oracle Integrated Lights Out Manager Cabling Connections](#) (page E-3)
- [Administrative Gigabit Ethernet Cabling Connections](#) (page E-4)
- [Single-Phase Power Distribution Unit Cabling Connections](#) (page E-5)
- [Three-Phase Power Distribution Unit Cabling Connections](#) (page E-6)
- [InfiniBand Switch-to-Switch Cabling Connections](#) (page E-7)
- [InfiniBand Switch-to-Server Cabling Connections](#) (page E-8)

E.1 Cable Color Coding

The cables used in Recovery Appliance are color coded as follows:

- **Black:** InfiniBand cables, AC power jumper cables, and Ethernet cables for storage servers
- **Blue:** Gigabit Ethernet cables for compute servers
- **Red:** Integrated Lights Out Manager (ILOM) cables

E.2 Recovery Appliance X7, X6, and X5 Rack Layout

[Figure E-1](#) (page E-2) shows the front and rear views of a Recovery Appliance X7, X6, or X5 full rack. The front is shown on the left, and the rear is shown on the right.

 **Note:**

The 1U empty space at the bottom of the rack is reserved for an optional InfiniBand spine switch, which is used in multitrack installations. Do not install any other unit in this space.

Figure E-1 Rack Layout of the Recovery Appliance X7, X6, or X5 Full Rack

Rear View				Front View			
U#	PDU A	Zero Data Loss Recovery Appliance	PDU B	U#	PDU B	Zero Data Loss Recovery Appliance	PDU A
42	INP2	Storage Server	INP0	42	INP0	Storage Server	INP2
41		Storage Server					
40		Storage Server					
39		Storage Server					
38		Storage Server					
37		Storage Server					
36		Storage Server					
35		Storage Server					
34		Storage Server					
33		Storage Server					
32		Storage Server					
31		Storage Server					
30	Storage Server	INP1	INP1				
29	Storage Server						
28	Storage Server						
27	Storage Server						
26	Storage Server						
25	Storage Server						
24	Storage Server						
23	Storage Server						
22	Infiniband Gateway Switch						
21	Ethernet Switch						
20	Infiniband Gateway Switch						
19	Storage Server						
18	Compute Server						
17	Compute Server						
16	Storage Server	INP0	INP0				
15	Storage Server						
14	Storage Server						
13	Storage Server						
12	Storage Server						
11	Storage Server						
10	Storage Server						
9	Storage Server						
8	Storage Server						
7	Storage Server						
6	Storage Server						
5	Storage Server						
4	Storage Server						
3	Storage Server						
2	Storage Server						
1	1U Solid Filler						

Figure E-2 (page E-3) shows the front and rear views of a Recovery Appliance X7, X6, or X5 minimum configuration.

Figure E-2 Rack Layout of the Recovery Appliance X7, X6, or X5 Minimum Configuration

Rear View				Front View			
U#	PDU A	Zero Data Loss Recovery Appliance	PDU B	U#	PDU B	Zero Data Loss Recovery Appliance	PDU A
42	INP2	1U Solid Filler	INP0	42	INP0	1U Solid Filler	INP2
41		1U Solid Filler					
40		1U Solid Filler					
39		1U Solid Filler					
38		1U Solid Filler					
37		1U Solid Filler					
36		1U Solid Filler					
35		1U Solid Filler					
34		1U Solid Filler					
33		1U Solid Filler					
32	1U Solid Filler	INP1	32	INP1	1U Solid Filler	INP1	
31	1U Solid Filler						
30	1U Solid Filler						
29	1U Solid Filler						
28	1U Solid Filler						
27	1U Solid Filler						
26	1U Solid Filler						
25	1U Solid Filler						
24	1U Solid Filler						
23	1U Solid Filler						
22	Infiniband Gateway Switch	INP0	22	INP2	1U Vented Filler	INP0	
21	Ethernet Switch						
20	Infiniband Gateway Switch						
19	1U Solid Filler						
18	1U Solid Filler						
17	Compute Server						
16	Compute Server						
15	1U Solid Filler						
14	1U Solid Filler						
13	1U Solid Filler		INP1		13		INP1
12	1U Solid Filler						
11	1U Solid Filler						
10	1U Solid Filler						
9	1U Solid Filler						
8	1U Solid Filler						
7	Storage Server						
6	Storage Server						
5	Storage Server						
4	Storage Server						
3	Storage Server	INP0	3	INP0	Storage Server	INP0	
2	Storage Server						
1	1U Solid Filler						

E.3 Oracle Integrated Lights Out Manager Cabling Connections

Table E-1 (page E-4) shows the cable connections from the servers to the Oracle ILOM switches in a Recovery Appliance X7, X6, or X5 full rack. The port numbers are the same in racks with fewer storage servers. The ILOM port on each server is labeled NET MGT and connects to the gigabit Ethernet switch located in rack unit 21.

ILOM cables are red.

Table E-1 ILOM Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit ¹	Component	To Ethernet Port
U41	Storage server	16
U39	Storage server	2
U37	Storage server	4
U35	Storage server	6
U33	Storage server	8
U31	Storage server	10
U29	Storage server	12
U27	Storage server	14
U25	Storage server	18
U23	Storage server	22
U18	Storage server	26
U17	Compute server	28
U16	Compute server	30
U14	Storage server	32
U12	Storage server	34
U10	Storage server	36
U8	Storage server	38
U6	Storage server	40
U4	Storage server	42
U2	Storage server	44

¹ Un is the unit location in the rack, where n is the number.

E.4 Administrative Gigabit Ethernet Cabling Connections

Table E-2 (page E-4) shows the cable connections from the servers to the gigabit Ethernet switch in a Recovery Appliance X7, X6, or X5 full rack. The port numbers are the same in racks with fewer storage servers. The Ethernet switch is located in rack unit 21. Ethernet cables are blue.

Table E-2 Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit ¹	Component	To Ethernet Port
U41	Storage server	15
U39	Storage server	1
U37	Storage server	3
U35	Storage server	5
U33	Storage server	7

Table E-2 (Cont.) Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit ¹	Component	To Ethernet Port
U31	Storage server	9
U29	Storage server	11
U27	Storage server	13
U25	Storage server	17
U23	Storage server	21
U22	InfiniBand leaf switch	45
U20	InfiniBand leaf switch	46
U18	Storage server	25
U17	Compute server	27
U16	Compute server	29
U14	Storage server	31
U12	Storage server	33
U10	Storage server	35
U8	Storage server	37
U6	Storage server	39
U4	Storage server	41
U2	Storage server	43

¹ Un is the unit location in the rack, where n is the number.

E.5 Single-Phase Power Distribution Unit Cabling Connections

Table E-3 (page E-5) shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X7, X6, or X5 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table E-3 Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U42	G5-5	G0-1	2 meters
U41	G4-5	G1-1	2 meters
U39	G5-6	G0-0	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-1	G0-5	2 meters

Table E-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U33	G5-0	G0-6	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-3	G1-3	2 meters
U27	G4-2	G1-4	2 meters
U25	G3-5	G2-1	2 meters
U23	G3-3	G2-3	2 meters
U22	G3-1	G2-5	2 meters
U21	G3-0	G2-6	2 meters
U20	G1-5	G4-1	2 meters
U18	G2-2	G3-4	2 meters
U17	G2-1	G3-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters
U1	G0-0	G5-6	2 meters

¹ Un is the unit location in the rack, where n is the number.

E.6 Three-Phase Power Distribution Unit Cabling Connections

Table E-4 (page E-6) shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X7, X6, or X5 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table E-4 Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U42	G5-4	G2-0	2 meters
U41	G5-5	G2-1	2 meters

Table E-4 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U39	G5-6	G2-0	2 meters
U37	G5-3	G2-3	2 meters
U35	G5-0	G2-6	2 meters
U33	G5-2	G2-4	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G4-1	G1-5	2 meters
U25	G3-5	G0-1	2 meters
U23	G3-3	G0-3	2 meters
U22	G3-2	G0-4	2 meters
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U18	G2-2	G5-4	2 meters
U17	G2-1	G5-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G1-5	G4-1	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters
U1	G0-0	G3-6	2 meters

¹ Un is the unit location in the rack, where n is the number.

E.7 InfiniBand Switch-to-Switch Cabling Connections

Table E-5 (page E-7) lists the ports and cable lengths for the InfiniBand connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) InfiniBand cables.

Table E-5 InfiniBand Switch-to-Switch Cabling

From InfiniBand Switch U20 Port ¹	To InfiniBand Switch U22 Port ¹	Cable Length
9B	9A	1 meter
10B	10A	1 meter

Table E-5 (Cont.) InfiniBand Switch-to-Switch Cabling

From InfiniBand Switch U20 Port ¹	To InfiniBand Switch U22 Port ¹	Cable Length
11B	11A	1 meter
8A	8A	1 meter
9A	9B	1 meter
10A	10B	1 meter
11A	11B	1 meter

¹ Un is the unit location in the rack, where n is the number.

E.8 InfiniBand Switch-to-Server Cabling Connections

Table E-6 (page E-8) lists the location, ports, and cables for the InfiniBand connections from the gateway leaf switches to the servers in a Recovery Appliance X6 or X5 full rack. The port numbers and cable lengths are the same in racks with fewer storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) InfiniBand cables.

Table E-6 InfiniBand Switch-to-Server Cabling for the Recovery Appliance X6 or X5 Full Racks

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U22	6B	U41	Storage server	PCIe 3, P1	3 meters
U22	0A	U39	Storage server	PCIe 3, P1	3 meters
U22	0B	U37	Storage server	PCIe 3, P1	3 meters
U22	1A	U35	Storage server	PCIe 3, P1	3 meters
U22	1B	U33	Storage server	PCIe 3, P1	3 meters
U22	2A	U31	Storage server	PCIe 3, P1	3 meters
U22	2B	U29	Storage server	PCIe 3, P1	3 meters
U22	3B	U27	Storage server	PCIe 3, P1	3 meters
U22	4A	U25	Storage server	PCIe 3, P1	2 meters
U22	5A	U23	Storage server	PCIe 3, P1	2 meters
U22	13A	U18	Storage server	PCIe 3, P2	2 meters
U22	13B	U17	Compute server	PCIe 2, P2	2 meters
U22	14B	U16	Compute server	PCIe 2, P2	2 meters
U22	14A	U14	Storage server	PCIe 3, P2	3 meters
U22	15A	U12	Storage server	PCIe 3, P2	3 meters
U22	15B	U10	Storage server	PCIe 3, P2	3 meters
U22	16A	U8	Storage server	PCIe 3, P2	3 meters
U22	16B	U6	Storage server	PCIe 3, P2	3 meters

Table E-6 (Cont.) InfiniBand Switch-to-Server Cabling for the Recovery Appliance X6 or X5 Full Racks

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U22	17A	U4	Storage server	PCIe 3, P2	3 meters
U22	17B	U2	Storage server	PCIe 3, P2	3 meters
U20	6B	U41	Storage server	PCIe 3, P2	3 meters
U20	0A	U39	Storage server	PCIe 3, P2	3 meters
U20	0B	U37	Storage server	PCIe 3, P2	3 meters
U20	1A	U35	Storage server	PCIe 3, P2	3 meters
U20	1B	U33	Storage server	PCIe 3, P2	3 meters
U20	2A	U31	Storage server	PCIe 3, P2	3 meters
U20	2B	U29	Storage server	PCIe 3, P2	3 meters
U20	3B	U27	Storage server	PCIe 3, P2	3 meters
U20	4A	U25	Storage server	PCIe 3, P2	2 meters
U20	5A	U23	Storage server	PCIe 3, P2	2 meters
U20	13A	U18	Storage server	PCIe 3, P1	2 meters
U20	13B	U17	Compute server	PCIe 2, P1	2 meters
U20	14B	U16	Compute server	PCIe 2, P1	2 meters
U20	14A	U14	Storage server	PCIe 3, P1	2 meters
U20	15A	U12	Storage server	PCIe 3, P1	3 meters
U20	15B	U10	Storage server	PCIe 3, P1	3 meters
U20	16A	U8	Storage server	PCIe 3, P1	3 meters
U20	16B	U6	Storage server	PCIe 3, P1	3 meters
U20	17A	U4	Storage server	PCIe 3, P1	3 meters
U20	17B	U2	Storage server	PCIe 3, P1	3 meters

¹ Un is the unit location in the rack, where n is the number.

F

In-Rack Cabling Tables for Recovery Appliance X4

This appendix identifies the location of the various components in a Recovery Appliance X4 full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

This appendix contains the following sections:

- [Recovery Appliance X4 Rack Layout](#) (page F-1)
- [Oracle Integrated Lights Out Manager Cabling Connections](#) (page F-3)
- [Administrative Gigabit Ethernet Cabling Connections](#) (page F-4)
- [Single-Phase Power Distribution Unit Cabling Connections](#) (page F-5)
- [Three-Phase Power Distribution Unit Cabling Connections](#) (page F-6)
- [InfiniBand Switch-to-Switch Cabling Connections](#) (page F-7)
- [InfiniBand Switch-to-Server Cabling Connections](#) (page F-7)

The cables used in Recovery Appliance are color-coded as follows:

- **Black:** InfiniBand cables, or AC power jumper cables
- **Blue:** Gigabit Ethernet cables
- **Red:** Integrated Lights Out Manager (ILOM) cables

F.1 Recovery Appliance X4 Rack Layout

[Figure F-1](#) (page F-2) shows the front and rear views of a Recovery Appliance X4 full rack. The front is shown on the left, and the rear is shown on the right.

Note:

The 1U empty space at the bottom of the rack is reserved for an optional InfiniBand spine switch, which is used in multirack installations. Do not install any other unit in this space.

Figure F-1 Rack Layout of the Recovery Appliance X4 Full Rack

Rear View				Front View			
U#	PDU A	Zero Data Loss Recovery Appliance	PDU B	U#	PDU B	Zero Data Loss Recovery Appliance	PDU A
42	INP2	1U Solid Filler	INP0	42	INP0	1U Solid Filler	INP2
41		1U Solid Filler					
40		Storage Server					
39		Storage Server					
38		Storage Server					
37		Storage Server					
36		Storage Server					
35		Storage Server					
34		Storage Server					
33		Storage Server					
32		Storage Server					
31		Storage Server					
30		Storage Server					
29	Storage Server	INP1	29	Storage Server	INP1		
28	Storage Server						
27	Storage Server						
26	1U Solid Filler						
25	1U Solid Filler						
24	1U Solid Filler						
23	1U Solid Filler						
22	Infiniband Gateway Switch						
21	Ethernet Switch						
20	Infiniband Gateway Switch						
19	1U Solid Filler						
18	1U Solid Filler						
17	Compute Server						
16	Compute Server						
15	Storage Server	INP0	15	Storage Server	INP0		
14	Storage Server						
13	Storage Server						
12	Storage Server						
11	Storage Server						
10	Storage Server						
9	Storage Server						
8	Storage Server						
7	Storage Server						
6	Storage Server						
5	Storage Server						
4	Storage Server						
3	Storage Server						
2	Storage Server						
1	1U Solid Filler						

Figure F-2 (page F-3) shows the front and rear views of a Recovery Appliance X4 minimum configuration.

Figure F-2 Rack Layout of the Recovery Appliance X4 Minimum Configuration

Rear View				Front View			
U#	PDU A	Zero Data Loss Recovery Appliance	PDU B	U#	PDU B	Zero Data Loss Recovery Appliance	PDU A
42	INP2	1U Solid Filler	INP0	42	INP0	1U Solid Filler	INP2
41		1U Solid Filler					
40		1U Solid Filler					
39		1U Solid Filler					
38		1U Solid Filler					
37		1U Solid Filler					
36		1U Solid Filler					
35		1U Solid Filler					
34		1U Solid Filler					
33		1U Solid Filler					
32	1U Solid Filler						
31	1U Solid Filler						
30	1U Solid Filler						
29	1U Solid Filler						
28	INP1	1U Solid Filler	INP1	28	INP1	1U Solid Filler	INP1
27		1U Solid Filler					
26		1U Solid Filler					
25		1U Solid Filler					
24		1U Solid Filler					
23		1U Solid Filler					
22		Infiniband Gateway Switch					
21		Ethernet Switch					
20		Infiniband Gateway Switch					
19		1U Solid Filler					
18	1U Solid Filler						
17	Compute Server						
16	Compute Server						
15	1U Solid Filler						
14	1U Solid Filler						
13	INP0	1U Solid Filler	INP2	13	INP2	1U Solid Filler	INP0
12		1U Solid Filler					
11		1U Solid Filler					
10		1U Solid Filler					
9		1U Solid Filler					
8		1U Solid Filler					
7		Storage Server					
6		Storage Server					
5		Storage Server					
4		Storage Server					
3	Storage Server						
2	Storage Server						
1	1U Solid Filler						

F.2 Oracle Integrated Lights Out Manager Cabling Connections

Table F-1 (page F-4) shows the cable connections from the servers to the Oracle ILOM switches in a Recovery Appliance X4 full rack. The port numbers are the same in racks with fewer storage servers. The ILOM port on each server is labeled NET MGT and connects to the gigabit Ethernet switch located in rack unit 21.

ILOM cables are red.

Table F-1 ILOM Cabling for the Recovery Appliance X4 Full Racks

From Rack Unit ¹	Component	To Ethernet Port
U39	Storage server	2
U37	Storage server	4
U35	Storage server	6
U33	Storage server	8
U31	Storage server	10
U29	Storage server	12
U27	Storage server	14
U17	Compute server	28
U16	Compute server	30
U14	Storage server	32
U12	Storage server	34
U10	Storage server	36
U8	Storage server	38
U6	Storage server	40
U4	Storage server	42
U2	Storage server	44

¹ Un is the unit location in the rack, where n is the number.

F.3 Administrative Gigabit Ethernet Cabling Connections

Table F-2 (page F-4) shows the cable connections from the servers to the gigabit Ethernet switch in a Recovery Appliance X4 full rack. The port numbers are the same in racks with fewer storage servers. The Ethernet switch is located in rack unit 21. Ethernet cables are blue.

Table F-2 Gigabit Ethernet Cabling for the Recovery Appliance X4 Full Racks

From Rack Unit ¹	Component	To Ethernet Port
U39	Storage server	1
U37	Storage server	3
U35	Storage server	5
U33	Storage server	7
U31	Storage server	9
U29	Storage server	11
U27	Storage server	13
U22	InfiniBand leaf switch	45
U20	InfiniBand leaf switch	46
U17	Compute server	27

Table F-2 (Cont.) Gigabit Ethernet Cabling for the Recovery Appliance X4 Full Racks

From Rack Unit ¹	Component	To Ethernet Port
U16	Compute server	29
U14	Storage server	31
U12	Storage server	33
U10	Storage server	35
U8	Storage server	37
U6	Storage server	39
U4	Storage server	41
U2	Storage server	43

¹ Un is the unit location in the rack, where n is the number.

F.4 Single-Phase Power Distribution Unit Cabling Connections

[Table F-3](#) (page F-5) shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X4 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-3 Single-Phase PDU Cabling for the Recovery Appliance X4 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U39	G5-6	G0-0	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-0	G0-6	2 meters
U33	G4-6	G1-0	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G3-6	G2-0	2 meters
U22	G3-1	G2-5	2 meters
U21	G3-0	G2-6	2 meters
U20	G2-4	G3-2	2 meters
U17	G2-1	G3-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters

Table F-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X4 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters

¹ Un is the unit location in the rack, where n is the number.

F.5 Three-Phase Power Distribution Unit Cabling Connections

Table F-4 (page F-6) shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X4 full rack. The destinations and the cable lengths are the same in racks with fewer storage servers. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table F-4 Three-Phase PDU Cabling for the Recovery Appliance X4 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U39	G5-6	G2-0	2 meters
U37	G5-3	G2-3	2 meters
U35	G5-0	G2-6	2 meters
U33	G4-6	G1-0	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G3-6	G0-0	2 meters
U22	G3-1	G0-5	2 meters
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U17	G2-1	G5-5	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G5-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters

¹ Un is the unit location in the rack, where n is the number.

F.6 InfiniBand Switch-to-Switch Cabling Connections

Table F-5 (page F-7) lists the ports and cable lengths for the InfiniBand connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) InfiniBand cables.

Table F-5 InfiniBand Switch-to-Switch Cabling

From InfiniBand Switch U20 Port ¹	To InfiniBand Switch U22 Port ¹	Cable Length
9B	9A	1 meter
10B	10A	1 meter
11B	11A	1 meter
8A	8A	1 meter
9A	9B	1 meter
10A	10B	1 meter
11A	11B	1 meter

¹ Un is the unit location in the rack, where n is the number.

F.7 InfiniBand Switch-to-Server Cabling Connections

Table F-6 (page F-7) lists the location, ports, and cables for the InfiniBand connections from the gateway leaf switches to the servers in a Recovery Appliance X4 full rack. The port numbers and cable lengths are the same in racks with fewer storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) InfiniBand cables.

Table F-6 InfiniBand Switch-to-Server Cabling for the Recovery Appliance X4 Full Racks

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit	To Rack Unit	Port	Cable Length
U22	0A	U39	Storage server	PCIe 3, P1	3 meters
U22	0B	U37	Storage server	PCIe 3, P1	3 meters
U22	1A	U35	Storage server	PCIe 3, P1	3 meters
U22	1B	U33	Storage server	PCIe 3, P1	3 meters
U22	2A	U31	Storage server	PCIe 3, P1	3 meters
U22	2B	U29	Storage server	PCIe 3, P1	3 meters
U22	3B	U27	Storage server	PCIe 3, P1	3 meters
U22	13B	U17	Compute server	PCIe 2, P2	2 meters
U22	14B	U16	Compute server	PCIe 2, P2	2 meters
U22	14A	U14	Storage server	PCIe 3, P2	3 meters

Table F-6 (Cont.) InfiniBand Switch-to-Server Cabling for the Recovery Appliance X4 Full Racks

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit	To Rack Unit	Port	Cable Length
U22	15A	U12	Storage server	PCIe 3, P2	3 meters
U22	15B	U10	Storage server	PCIe 3, P2	3 meters
U22	16A	U8	Storage server	PCIe 3, P2	3 meters
U22	16B	U6	Storage server	PCIe 3, P2	3 meters
U22	17A	U4	Storage server	PCIe 3, P2	3 meters
U22	17B	U2	Storage server	PCIe 3, P2	3 meters
U20	0A	U39	Storage server	PCIe 3, P2	3 meters
U20	0B	U37	Storage server	PCIe 3, P2	3 meters
U20	1A	U35	Storage server	PCIe 3, P2	3 meters
U20	1B	U33	Storage server	PCIe 3, P2	3 meters
U20	2A	U31	Storage server	PCIe 3, P2	3 meters
U20	2B	U29	Storage server	PCIe 3, P2	3 meters
U20	3B	U27	Storage server	PCIe 3, P2	3 meters
U20	13B	U17	Compute server	PCIe 2, P1	2 meters
U20	14B	U16	Compute server	PCIe 2, P1	2 meters
U20	14A	U14	Storage server	PCIe 3, P1	2 meters
U20	15A	U12	Storage server	PCIe 3, P1	3 meters
U20	15B	U10	Storage server	PCIe 3, P1	3 meters
U20	16A	U8	Storage server	PCIe 3, P1	3 meters
U20	16B	U6	Storage server	PCIe 3, P1	3 meters
U20	17A	U4	Storage server	PCIe 3, P1	3 meters
U20	17B	U2	Storage server	PCIe 3, P1	3 meters

¹ Un is the unit location in the rack, where n is the number.

G

Multirack Cabling Tables

InfiniBand cables connect multiple racks together, in addition to connecting the components within a rack. The cable lengths provided in these tables are for adjacent racks with the cables running underneath the racks. If the racks are apart or the cables run above the racks, then longer cables might be required. The cable lengths listed here are only guidelines. You should always verify the actual cable lengths required.

This appendix contains the following sections:

- [About Multirack Cabling](#) (page G-1)
- [Key to Cabling Table Abbreviations](#) (page G-7)
- [Two-Rack Cabling](#) (page G-7)
- [Three-Rack Cabling](#) (page G-8)
- [Four-Rack Cabling](#) (page G-10)
- [Five-Rack Cabling](#) (page G-12)
- [Six-Rack Cabling](#) (page G-15)
- [Seven-Rack Cabling](#) (page G-18)
- [Eight-Rack Cabling](#) (page G-22)
- [Nine-Rack Cabling](#) (page G-27)
- [10-Rack Cabling](#) (page G-27)
- [11-Rack Cabling](#) (page G-28)
- [12-Rack Cabling](#) (page G-29)
- [13-Rack Cabling](#) (page G-29)
- [14-Rack Cabling](#) (page G-30)
- [15-Rack Cabling](#) (page G-31)
- [16-Rack Cabling](#) (page G-31)
- [17-Rack Cabling](#) (page G-32)
- [18-Rack Cabling](#) (page G-33)

See Also:

[Extending Recovery Appliance with Multiple Racks](#) (page 8-1)

G.1 About Multirack Cabling

You can connect up to 18 full Recovery Appliance racks. The racks do not include spare cables or a third Sun Datacenter InfiniBand Switch 36, which functions as the

spine switch. To extend a rack, order the cables and a Sun Datacenter InfiniBand Switch 36.

This section contains the following topics:

- [Cables Provided With a Spine Switch](#) (page G-2)
- [Optical Cables Required for Four or More Racks](#) (page G-2)
- [Cabling Three Switches Within a Rack](#) (page G-3)
- [Cabling the Switches in Two Racks](#) (page G-4)
- [Cabling the Switches in Up to Eight Racks](#) (page G-5)
- [Cabling Switches in Up to 18 Racks](#) (page G-5)



See Also:

"[Cabling Recovery Appliance to Oracle Exadata Database Machine](#) (page 9-2)"

G.1.1 Cables Provided With a Spine Switch

The Sun Datacenter InfiniBand Switch 36 switches attach to standard Quad Small Form-factor Pluggable (QSFP) connectors at the ends of the InfiniBand cables.

[Table G-1](#) (page G-2) lists the spare cables provided with the Sun Datacenter InfiniBand Switch 36 spine switch.

Table G-1 InfiniBand Cables Provided as Spares

Cable	Spare	Interswitch	Total
InfiniBand 2 meter cable	0	7	7
InfiniBand 3 meter cable	6	2	8
InfiniBand 5 meter cable	10	0	10

G.1.2 Optical Cables Required for Four or More Racks

Connecting four or more racks requires additional 10-meter optical QSFP-to-QSFP QDR InfiniBand cables. Each optical cable requires two QSFP SR transceivers.

[Table G-2](#) (page G-2) lists the cables needed to connect the racks.

Table G-2 Cable Requirements for Multirack Installations

Number of Racks	3 Meter Cables	5 Meter Cables	10 Meter Cables
2	16	16	0
3	18	30	0
4	16	40	8
5	20	42	18

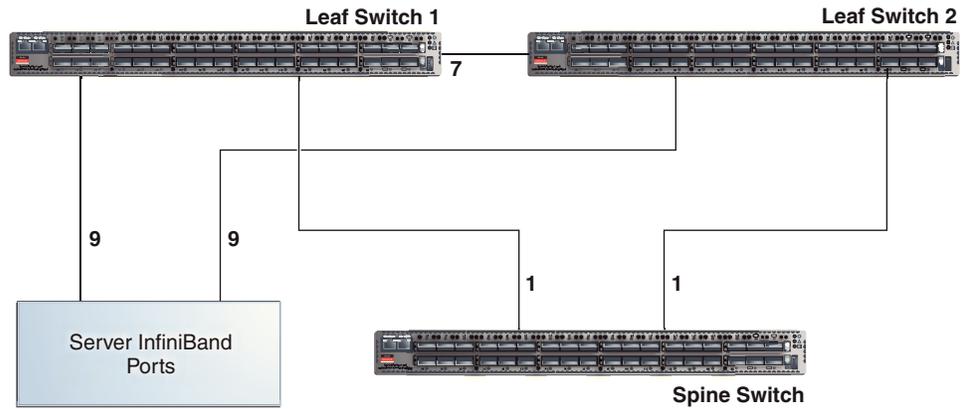
Table G-2 (Cont.) Cable Requirements for Multirack Installations

Number of Racks	3 Meter Cables	5 Meter Cables	10 Meter Cables
6	24	46	26
7	28	44	40
8	16	52	60

G.1.3 Cabling Three Switches Within a Rack

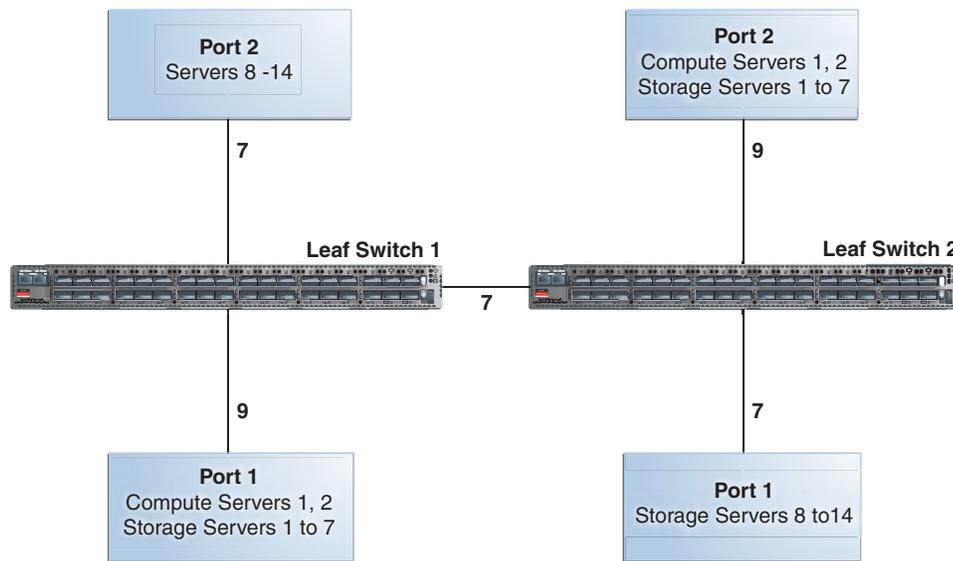
In a single rack, the two leaf switches are interconnected using seven connections. In addition, each leaf switch has one connection to the spine switch. The leaf switches connect to the spine switch as shown in [Figure G-1](#) (page G-3).

Figure G-1 In-Rack Connections Between Spine Switch and Leaf Switches



The servers connect to the leaf switches as shown in [Figure G-2](#) (page G-4).

Figure G-2 In-Rack Connections Between the Servers and the Leaf Switches



G.1.4 Cabling the Switches in Two Racks

When you connect up to eight racks, remove the seven interswitch connections between the leaf switches and the two connections between the leaf switches and the spine switch. From each leaf switch, distribute eight connections over the spine switches in all racks. In multirack environments, the leaf switches inside a rack are not directly connected.

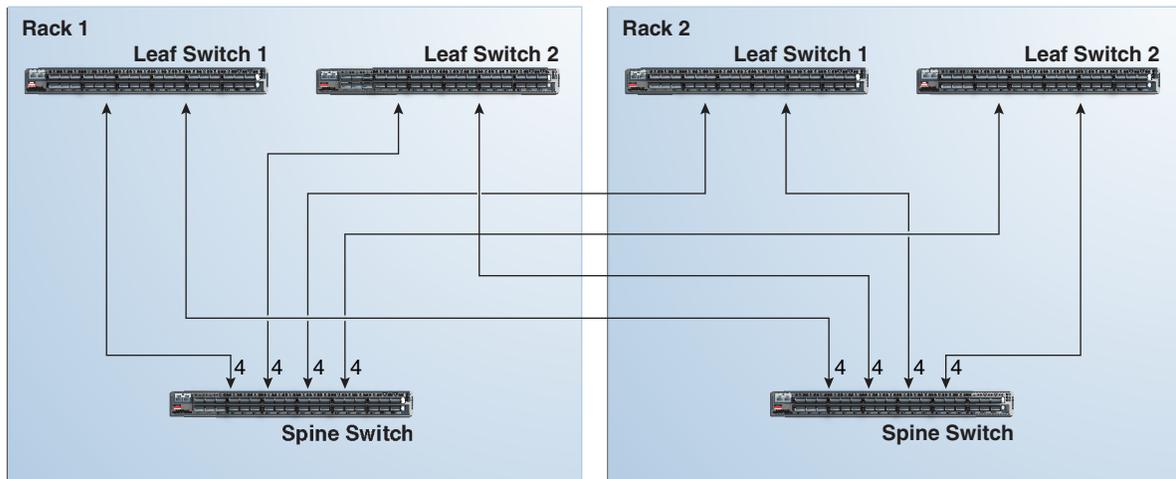
Figure G-3 (page G-5) shows that each leaf switch in rack 1 connects to the following switches:

- Four in-rack connections to the spine switch in rack 1
- Four external connections to the spine switch in rack 2

The spine switch in rack 1 connects to the following switches:

- Eight in-rack connections to both leaf switches in rack 1
- Eight external connections to both leaf switches in rack 2

Figure G-3 Connecting Switches Across Two Racks

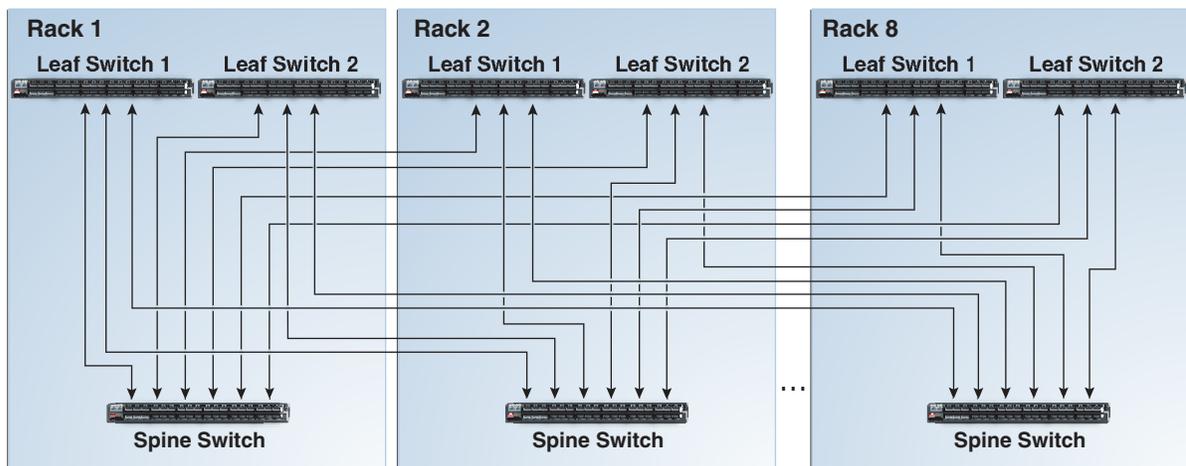


G.1.5 Cabling the Switches in Up to Eight Racks

As the number of racks increases from two to eight, the pattern continues, as shown in [Figure G-4](#) (page G-5). Because not all of the racks are shown, not all of the cables are shown either.

Each leaf switch has eight interswitch connections distributed across all spine switches. Each spine switch has 16 interswitch connections distributed across all leaf switches. The leaf switches are not directly connected to other leaf switches, and the spine switches are not directly connected to the other spine switches.

Figure G-4 Connecting Switches Across Eight Racks



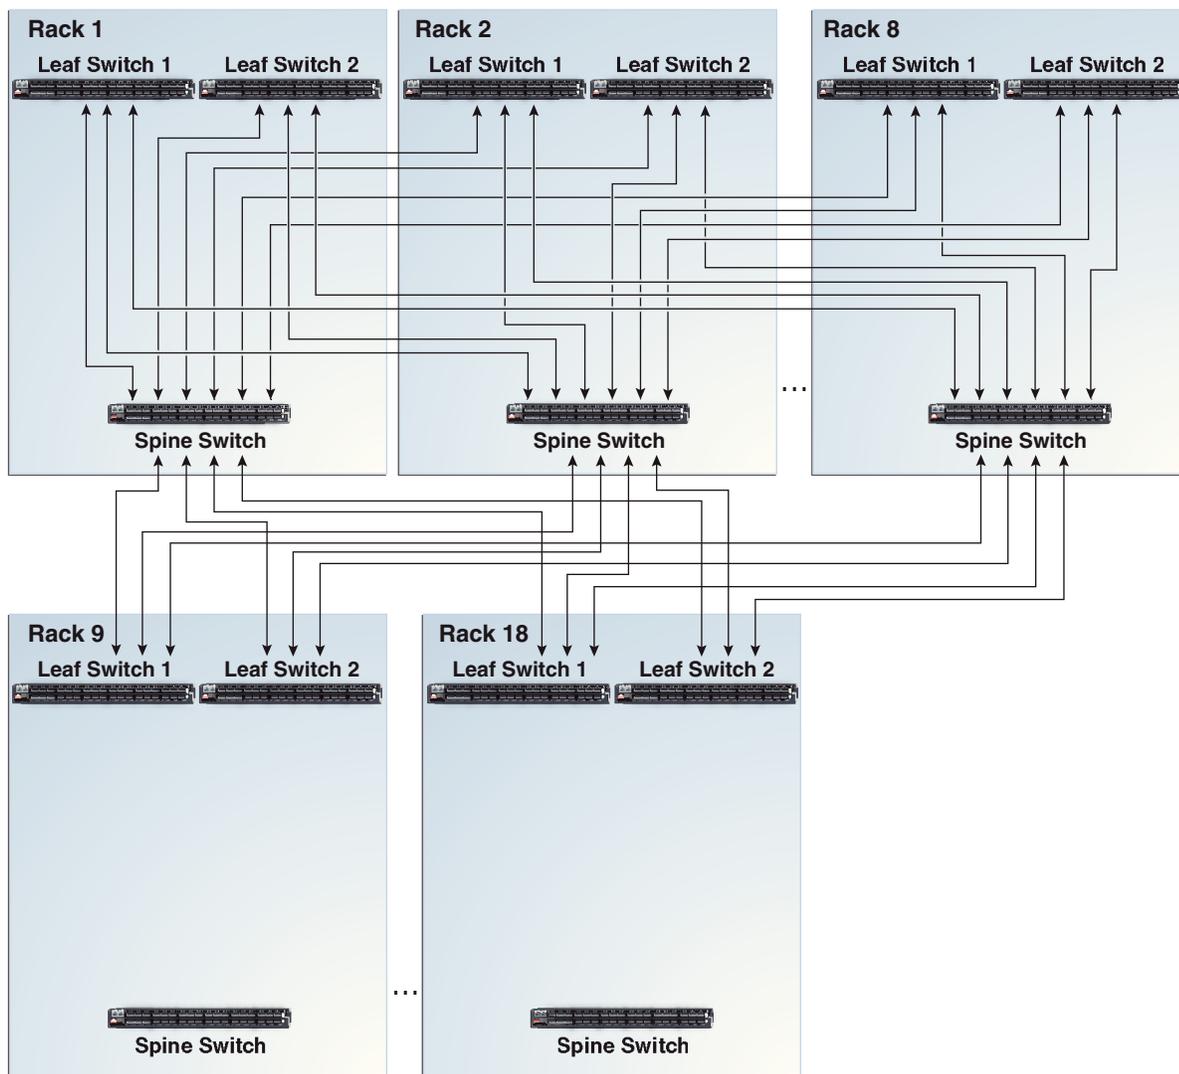
G.1.6 Cabling Switches in Up to 18 Racks

When you connect racks 9 through 18, remove the interswitch connections between the leaf switches, and disconnect the spine switches from the InfiniBand fabric.

Figure G-5 (page G-6) shows that the eight links from every leaf switch connect to the spine switches in racks 1 to 8.

The connections among racks 1 to 8 do not change you add the cables for racks 9 to 18.

Figure G-5 Connecting Switches Across 18 Racks



Note:

Cable lengths are specified for racks 1 through 8. Cable lengths for racks 9 through 18 vary depending on the layout of the racks.

G.2 Key to Cabling Table Abbreviations

The following abbreviations are used in the tables:

Abbreviation	Description
Rn	Rack <i>n</i> , where <i>n</i> is the number of the rack, such as R1.
IBn	Unit location in rack, where <i>n</i> is the number, such as IB3.
Pn	InfiniBand port <i>n</i> , where <i>n</i> is port number, such as P8A.

G.3 Two-Rack Cabling

Table G-3 (page G-7) shows the cable connections for the first spine switch (R1-U1) when two full racks are cabled together.



See Also:

["Cabling Two Racks Together \(page 8-4\)"](#)

Table G-3 Leaf Switch Connections for Rack 1 in a Two-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A R1-IB3-P8B to R1-U1-P4A R1-IB3-P9A to R1-U1-P5A R1-IB3-P9B to R1-U1-P6A	3 meters
R1 IB3 to Rack 2	R1-IB3-P10A to R2-U1-P7A R1-IB3-P10B to R2-U1-P8A R1-IB3-P11A to R2-U1-P9A R1-IB3-P11B to R2-U1-P10A	5 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B R1-IB2-P9A to R1-U1-P5B R1-IB2-P9B to R1-U1-P6B	3 meters
R1 IB2 to Rack 2	R1-IB2-P10A to R2-U1-P7B R1-IB2-P10B to R2-U1-P8B R1-IB2-P11A to R2-U1-P9B R1-IB2-P11B to R2-U1-P10B	5 meters

Table G-4 (page G-8) shows the cable connections for the second spine switch (R2-U1) when two full racks are cabled together.

Table G-4 Leaf Switch Connections for Rack 2 in a Two-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A R2-IB3-P9A to R2-U1-P5A R2-IB3-P9B to R2-U1-P6A	3 meters
R2 IB3 to Rack 1	R2-IB3-P10A to R1-U1-P7A R2-IB3-P10B to R1-U1-P8A R2-IB3-P11A to R1-U1-P9A R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B R2-IB2-P9A to R2-U1-P5B R2-IB2-P9B to R2-U1-P6B	3 meters
R2 IB2 to Rack 1	R2-IB2-P10A to R1-U1-P7B R2-IB2-P10B to R1-U1-P8B R2-IB2-P11A to R1-U1-P9B R2-IB2-P11B to R1-U1-P10B	5 meters

G.4 Three-Rack Cabling

Table G-5 (page G-8) shows the cable connections for the first spine switch (R1-U1) when three full racks are cabled together.

Table G-5 Leaf Switch Connections for Rack 1 in a Three-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A R1-IB3-P8B to R1-U1-P4A R1-IB3-P9A to R1-U1-P5A	3 meters
R1 IB3 to Rack 2	R1-IB3-P9B to R2-U1-P6A R1-IB3-P10A to R2-U1-P7A R1-IB3-P10B to R2-U1-P8A	5 meters
R1 IB3 to Rack 3	R1-IB3-P11A to R3-U1-P9A R1-IB3-P11B to R3-U1-P10A	5 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B R1-IB2-P9A to R1-U1-P5B	3 meters
R1 IB2 to Rack 2	R1-IB2-P9B to R2-U1-P6B R1-IB2-P10A to R2-U1-P7B R1-IB2-P10B to R2-U1-P8B	5 meters
R1 IB2 to Rack 3	R1-IB2-P11A to R3-U1-P9B R1-IB2-P11B to R3-U1-P10B	5 meters

Table G-6 (page G-9) shows the cable connections for the second spine switch (R2-U1) when three full racks are cabled together.

Table G-6 Leaf Switch Connections for Rack 2 in a Three-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A R2-IB3-P9A to R2-U1-P5A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11A to R1-U1-P9A R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9B to R3-U1-P6A R2-IB3-P10A to R3-U1-P7A R2-IB3-P10B to R3-U1-P8A	5 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B R2-IB2-P9A to R2-U1-P5B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11A to R1-U1-P9B R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9B to R3-U1-P6B R2-IB2-P10A to R3-U1-P7B R2-IB2-P10B to R3-U1-P8B	5 meters

Table G-7 (page G-9) shows the cable connections for the third spine switch (R3-U1) when three full racks are cabled together.

Table G-7 Leaf Switch Connections for Rack 3 in a Three-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A R3-IB3-P8B to R3-U1-P4A R3-IB3-P9A to R3-U1-P5A	3 meters
R3 IB3 to Rack 1	R3-IB3-P9B to R1-U1-P6A R3-IB3-P10A to R1-U1-P7A R3-IB3-P10B to R1-U1-P8A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11A to R2-U1-P9A R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B R3-IB2-P8B to R3-U1-P4B R3-IB2-P9A to R3-U1-P5B	3 meters
R3 IB2 to Rack 1	R3-IB2-P9B to R1-U1-P6B R3-IB2-P10A to R1-U1-P7B R3-IB2-P10B to R1-U1-P8B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11A to R2-U1-P9B R3-IB2-P11B to R2-U1-P10B	5 meters

G.5 Four-Rack Cabling

Table G-8 (page G-10) shows the cable connections for the first spine switch (R1-U1) when four full racks are cabled together.

Table G-8 Leaf Switch Connections for Rack 1 in a Four-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A R1-IB3-P8B to R1-U1-P4A	3 meters
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A R1-IB3-P9B to R2-U1-P6A	5 meters
R1 IB3 to Rack 3	R1-IB3-P10A to R3-U1-P7A R1-IB3-P10B to R3-U1-P8A	5 meters
R1 IB3 to Rack 4	R1-IB3-P11A to R4-U1-P9A R1-IB3-P11B to R4-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B	3 meters
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B R1-IB2-P9B to R2-U1-P6B	5 meters
R1 IB2 to Rack 3	R1-IB2-P10A to R3-U1-P7B R1-IB2-P10B to R3-U1-P8B	5 meters
R1 IB2 to Rack 4	R1-IB2-P11A to R4-U1-P9B R1-IB2-P11B to R4-U1-P10B	10 meters

Table G-9 (page G-10) shows the cable connections for the second spine switch (R2-U1) when four full racks are cabled together.

Table G-9 Leaf Switch Connections for Rack 2 in a Four-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11A to R1-U1-P9A R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-U1-P5A R2-IB3-P9B to R3-U1-P6A	5 meters
R2 IB3 to Rack 4	R2-IB3-P10A to R4-U1-P7A R2-IB3-P10B to R4-U1-P8A	5 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11A to R1-U1-P9B R2-IB2-P11B to R1-U1-P10B	5 meters

Table G-9 (Cont.) Leaf Switch Connections for Rack 2 in a Four-Rack System

Leaf Switch	Connection	Cable Length
R2 IB2 to Rack 3	R2-IB2-P9A to R3-U1-P5B R2-IB2-P9B to R3-U1-P6B	5 meters
R2 IB2 to Rack 4	R2-IB2-P10A to R4-U1-P7B R2-IB2-P10B to R4-U1-P8B	5 meters

Table G-10 (page G-11) shows the cable connections for the third spine switch (R3-U1) when four full racks are cabled together.

Table G-10 Leaf Switch Connections for Rack 3 in a Four-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A R3-IB3-P8B to R3-U1-P4A	3 meters
R3 IB3 to Rack 1	R3-IB3-P10A to R1-U1-P7A R3-IB3-P10B to R1-U1-P8A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11A to R2-U1-P9A R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-U1-P5A R3-IB3-P9B to R4-U1-P6A	5 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B R3-IB2-P8B to R3-U1-P4B	3 meters
R3 IB2 to Rack 1	R3-IB2-P10A to R1-U1-P7B R3-IB2-P10B to R1-U1-P8B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11A to R2-U1-P9B R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-U1-P5B R3-IB2-P9B to R4-U1-P6B	5 meters

Table G-11 (page G-11) shows the cable connections for the fourth spine switch (R4-U1) when four full racks are cabled together.

Table G-11 Leaf Switch Connections for Rack 4 in a Four-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A R4-IB3-P8B to R4-U1-P4A	3 meters
R4 IB3 to Rack 1	R4-IB3-P9A to R1-U1-P5A R4-IB3-P9B to R1-U1-P6A	10 meters
R4 IB3 to Rack 2	R4-IB3-P10A to R2-U1-P7A R4-IB3-P10B to R2-U1-P8A	5 meters

Table G-11 (Cont.) Leaf Switch Connections for Rack 4 in a Four-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 to Rack 3	R4-IB3-P11A to R3-U1-P9A R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B R4-IB2-P8B to R4-U1-P4B	3 meters
R4 IB2 to Rack 1	R4-IB2-P9A to R1-U1-P5B R4-IB2-P9B to R1-U1-P6B	10 meters
R4 IB2 to Rack 2	R4-IB2-P10A to R2-U1-P7B R4-IB2-P10B to R2-U1-P8B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11A to R3-U1-P9B R4-IB2-P11B to R3-U1-P10B	5 meters

G.6 Five-Rack Cabling

[Table G-12](#) (page G-12) shows the cable connections for the first spine switch (R1-U1) when five full racks are cabled together.

Table G-12 Leaf Switch Connections for Rack 1 in a Five-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A R1-IB3-P8B to R1-U1-P4A	3 meters
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A R1-IB3-P9B to R2-U1-P6A	5 meters
R1 IB3 to Rack 3	R1-IB3-P10A to R3-U1-P7A R1-IB3-P10B to R3-U1-P8A	5 meters
R1 IB3 to Rack 4	R1-IB3-P11A to R4-U1-P9A	10 meters
R1 IB3 to Rack 5	R1-IB3-P11B to R5-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B	3 meters
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B R1-IB2-P9B to R2-U1-P6B	5 meters
R1 IB2 to Rack 3	R1-IB2-P10A to R3-U1-P7B R1-IB2-P10B to R3-U1-P8B	5 meters
R1 IB2 to Rack 4	R1-IB2-P11A to R4-U1-P9B	10 meters
R1 IB2 to Rack 5	R1-IB2-P11B to R5-U1-P10B	10 meters

[Table G-13](#) (page G-13) shows the cable connections for the second spine switch (R2-U1) when five full racks are cabled together.

Table G-13 Leaf Switch Connections for Rack 2 in a Five-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-U1-P5A R2-IB3-P9B to R3-U1-P6A	5 meters
R2 IB3 to Rack 4	R2-IB3-P10A to R4-U1-P7A R2-IB3-P10B to R4-U1-P8A	5 meters
R2 IB3 to Rack 5	R2-IB3-P11A to R5-U1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-U1-P5B R2-IB2-P9B to R3-U1-P6B	5 meters
R2 IB2 to Rack 4	R2-IB2-P10A to R4-U1-P7B R2-IB2-P10B to R4-U1-P8B	5 meters
R2 IB2 to Rack 5	R2-IB2-P11A to R5-U1-P9B	10 meters

[Table G-14](#) (page G-13) shows the cable connections for the third spine switch (R3-U1) when five full racks are cabled together.

Table G-14 Leaf Switch Connections for Rack 3 in a Five-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A R3-IB3-P8B to R3-U1-P4A	3 meters
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-U1-P5A R3-IB3-P9B to R4-U1-P6A	5 meters
R3 IB3 to Rack 5	R3-IB3-P10A to R5-U1-P7A R3-IB3-P10B to R5-U1-P8A	5 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B R3-IB2-P8B to R3-U1-P4B	3 meters
R3 IB2 to Rack 1	R3-IB2-P11A to R1-U1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-U1-P5B R3-IB2-P9B to R4-U1-P6B	5 meters
R3 IB2 to Rack 5	R3-IB2-P10A to R5-U1-P7B R3-IB2-P10B to R5-U1-P8B	5 meters

Table G-15 (page G-14) shows the cable connections for the fourth spine switch (R4-U1) when five full racks are cabled together.

Table G-15 Leaf Switch Connections for Rack 4 in a Five-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A R4-IB3-P8B to R4-U1-P4A	3 meters
R4 IB3 to Rack 1	R4-IB3-P10A to R1-U1-P7A R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-U1-P5A R4-IB3-P9B to R5-U1-P6A	5 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B R4-IB2-P8B to R4-U1-P4B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10A to R1-U1-P7B R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-U1-P5B R4-IB2-P9B to R5-U1-P6B	5 meters

Table G-16 (page G-14) shows the cable connections for the fifth spine switch (R5-U1) when five full racks are cabled together.

Table G-16 Leaf Switch Connections for Rack 5 in a Five-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A R5-IB3-P8B to R5-U1-P4A	3 meters
R5 IB3 to Rack 1	R5-IB3-P9A to R1-U1-P5A R5-IB3-P9B to R1-U1-P6A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10A to R2-U1-P7A R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B R5-IB2-P8B to R5-U1-P4B	3 meters
R5 IB2 to Rack 1	R5-IB2-P9A to R1-U1-P5B R5-IB2-P9B to R1-U1-P6B	10 meters
R5 IB2 to Rack 2	R5-IB2-P10A to R2-U1-P7B R5-IB3-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters

Table G-16 (Cont.) Leaf Switch Connections for Rack 5 in a Five-Rack System

Leaf Switch	Connection	Cable Length
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters

G.7 Six-Rack Cabling

Table G-17 (page G-15) shows the cable connections for the first spine switch (R1-U1) when six full racks are cabled together.

Table G-17 Leaf Switch Connections for Rack 1 in a Six-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A R1-IB3-P8B to R1-U1-P4A	3 meters
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A R1-IB3-P9B to R2-U1-P6A	5 meters
R1 IB3 to Rack 3	R1-IB3-P10A to R3-U1-P7A	5 meters
R1 IB3 to Rack 4	R1-IB3-P10B to R4-U1-P8A	10 meters
R1 IB3 to Rack 5	R1-IB3-P11A to R5-U1-P9A	10 meters
R1 IB3 to Rack 6	R1-IB3-P11B to R6-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B	3 meters
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B R1-IB2-P9B to R2-U1-P6B	5 meters
R1 IB2 to Rack 3	R1-IB2-P10A to R3-U1-P7B	5 meters
R1 IB2 to Rack 4	R1-IB2-P10B to R4-U1-P8B	10 meters
R1 IB2 to Rack 5	R1-IB2-P11A to R5-U1-P9B	10 meters
R1 IB2 to Rack 6	R1-IB2-P11B to R6-U1-P10B	10 meters

Table G-18 (page G-15) shows the cable connections for the second spine switch (R2-U1) when six full racks are cabled together.

Table G-18 Leaf Switch Connections for Rack 2 in a Six-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-U1-P5A R2-IB3-P9B to R3-U1-P6A	5 meters
R2 IB3 to Rack 4	R2-IB3-P10A to R4-U1-P7A	5 meters
R2 IB3 to Rack 5	R2-IB3-P10B to R5-U1-P8A	10 meters

Table G-18 (Cont.) Leaf Switch Connections for Rack 2 in a Six-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 to Rack 6	R2-IB3-P11A to R6-U1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-U1-P5B R2-IB2-P9B to R3-U1-P6B	5 meters
R2 IB2 to Rack 4	R2-IB2-P10A to R4-U1-P7B	5 meters
R2 IB2 to Rack 5	R2-IB2-P10B to R5-U1-P8B	10 meters
R2 IB2 to Rack 6	R2-IB2-P11A to R6-U1-P9B	10 meters

[Table G-19](#) (page G-16) shows the cable connections for the third spine switch (R3-U1) when six full racks are cabled together.

Table G-19 Leaf Switch Connections for Rack 3 in a Six-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A R3-IB3-P8B to R3-U1-P4A	3 meters
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-U1-P5A R3-IB3-P9B to R4-U1-P6A	5 meters
R3 IB3 to Rack 5	R3-IB3-P10A to R5-U1-P7A	5 meters
R3 IB3 to Rack 6	R3-IB3-P10B to R6-U1-P8A	10 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B R3-IB2-P8B to R3-U1-P4B	3 meters
R3 IB2 to Rack 1	R3-IB2-P11A to R1-U1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-U1-P5B R3-IB2-P9B to R4-U1-P6B	5 meters
R3 IB2 to Rack 5	R3-IB2-P10A to R5-U1-P7B	5 meters
R3 IB2 to Rack 6	R3-IB2-P10B to R6-U1-P8B	10 meters

[Table G-20](#) (page G-17) shows the cable connections for the fourth spine switch (R4-U1) when six full racks are cabled together.

Table G-20 Leaf Switch Connections for Rack 4 in a Six-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A R4-IB3-P8B to R4-U1-P4A	3 meters
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-U1-P5A R4-IB3-P9B to R5-U1-P6A	5 meters
R4 IB3 to Rack 6	R4-IB3-P10A to R6-U1-P7A	5 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B R4-IB2-P8B to R4-U1-P4B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-U1-P5B R4-IB2-P9B to R5-U1-P6B	5 meters
R4 IB2 to Rack 6	R4-IB2-P10A to R6-U1-P7B	5 meters

[Table G-21](#) (page G-17) shows the cable connections for the fifth spine switch (R5-U1) when six full racks are cabled together.

Table G-21 Leaf Switch Connections for Rack 5 in a Six-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A R5-IB3-P8B to R5-U1-P4A	3 meters
R5 IB3 to Rack 1	R5-IB3-P10A to R1-U1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P9A to R6-U1-P5A R5-IB3-P9B to R6-U1-P6A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B R5-IB2-P8B to R5-U1-P4B	3 meters
R5 IB2 to Rack 1	R5-IB2-P10A to R1-U1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB3-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P9A to R6-U1-P5B R5-IB2-P9B to R6-U1-P6B	5 meters

Table G-22 (page G-18) shows the cable connections for the sixth spine switch (R6-U1) when six full racks are cabled together.

Table G-22 Leaf Switch Connections for Rack 6 in a Six-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-U1-P3A R6-IB3-P8B to R6-U1-P4A	3 meters
R6 IB3 to Rack 1	R6-IB3-P9A to R1-U1-P5A R6-IB3-P9B to R1-U1-P6A	10 meters
R6 IB3 to Rack 2	R6-IB3-P10A to R2-U1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-U1-P8A	10 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-U1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-U1-P10A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-U1-P3B R6-IB2-P8B to R6-U1-P4B	3 meters
R6 IB2 to Rack 1	R6-IB2-P9A to R1-U1-P5B R6-IB2-P9B to R1-U1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB2-P10A to R2-U1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-U1-P8B	10 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters

G.8 Seven-Rack Cabling

Table G-23 (page G-18) shows the cable connections for the first spine switch (R1-U1) when seven racks are cabled together.

Table G-23 Leaf Switch Connections for Rack 1 in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A R1-IB3-P8B to R1-U1-P4A	3 meters
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A	5 meters
R1 IB3 to Rack 3	R1-IB3-P9B to R3-U1-P6A	5 meters
R1 IB3 to Rack 4	R1-IB3-P10A to R4-U1-P7A	10 meters
R1 IB3 to Rack 5	R1-IB3-P10B to R5-U1-P8A	10 meters
R1 IB3 to Rack 6	R1-IB3-P11A to R6-U1-P9A	10 meters
R1 IB3 to Rack 7	R1-IB3-P11B to R7-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B	3 meters
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B	5 meters
R1 IB2 to Rack 3	R1-IB2-P9B to R3-U1-P6B	5 meters

Table G-23 (Cont.) Leaf Switch Connections for Rack 1 in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R1 IB2 to Rack 4	R1-IB2-P10A to R4-U1-P7B	10 meters
R1 IB2 to Rack 5	R1-IB2-P10B to R5-U1-P8B	10 meters
R1 IB2 to Rack 6	R1-IB2-P11A to R6-U1-P9B	10 meters
R1 IB2 to Rack 7	R1-IB2-P11B to R7-U1-P10B	10 meters

Table G-24 (page G-19) shows the cable connections for the second spine switch (R2-U1) when seven full racks are cabled together.

Table G-24 Leaf Switch Connections for Rack 2 in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-U1-P5A	5 meters
R2 IB3 to Rack 4	R2-IB3-P9B to R4-U1-P6A	5 meters
R2 IB3 to Rack 5	R2-IB3-P10A to R5-U1-P7A	10 meters
R2 IB3 to Rack 6	R2-IB3-P10B to R6-U1-P8A	10 meters
R2 IB3 to Rack 7	R2-IB3-P11A to R7-U1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-U1-P5B	5 meters
R2 IB2 to Rack 4	R2-IB2-P9B to R4-U1-P6B	5 meters
R2 IB2 to Rack 5	R2-IB2-P10A to R5-U1-P7B	10 meters
R2 IB2 to Rack 6	R2-IB2-P10B to R6-U1-P8B	10 meters
R2 IB2 to Rack 7	R2-IB2-P11A to R7-U1-P9B	10 meters

Table G-25 (page G-19) shows the cable connections for the third spine switch (R3-U1) when seven full racks are cabled together.

Table G-25 Leaf Switch Connections for Rack 3 in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A R3-IB3-P8B to R3-U1-P4A	3 meters
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-U1-P5A	5 meters

Table G-25 (Cont.) Leaf Switch Connections for Rack 3 in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 to Rack 5	R3-IB3-P9B to R5-U1-P6A	5 meters
R3 IB3 to Rack 6	R3-IB3-P10A to R6-U1-P7A	10 meters
R3 IB3 to Rack 7	R3-IB3-P10B to R7-U1-P8A	10 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B R3-IB2-P8B to R3-U1-P4B	3 meters
R3 IB2 to Rack 1	R3-IB2-P11A to R1-U1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P9A to R4-U1-P5B	5 meters
R3 IB2 to Rack 5	R3-IB2-P9B to R5-U1-P6B	5 meters
R3 IB2 to Rack 6	R3-IB2-P10A to R6-U1-P7B	10 meters
R3 IB2 to Rack 7	R3-IB2-P10B to R7-U1-P8B	10 meters

Table G-26 (page G-20) shows the cable connections for the fourth spine switch (R4-U1) when seven full racks are cabled together.

Table G-26 Leaf Switch Connections for Rack 4 in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A R4-IB3-P8B to R4-U1-P4A	3 meters
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-U1-P5A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9B to R6-U1-P6A	5 meters
R4 IB3 to Rack 7	R4-IB3-P10A to R7-U1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B R4-IB2-P8B to R4-U1-P4B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-U1-P5B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9B to R6-U1-P6B	5 meters
R4 IB2 to Rack 7	R4-IB2-P10A to R7-U1-P7B	10 meters

Table G-27 (page G-21) shows the cable connections for the fifth spine switch (R5-U1) when seven full racks are cabled together.

Table G-27 Leaf Switch Connections for Rack 5 in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A R5-IB3-P8B to R5-U1-P4A	3 meters
R5 IB3 to Rack 1	R5-IB3-P10A to R1-U1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P9A to R6-U1-P5A	5 meters
R5 IB3 to Rack 7	R5-IB3-P9B to R7-U1-P6A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B R5-IB2-P8B to R5-U1-P4B	3 meters
R5 IB2 to Rack 1	R5-IB2-P10A to R1-U1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB3-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P9A to R6-U1-P5B	5 meters
R5 IB2 to Rack 7	R5-IB2-P9B to R7-U1-P6B	5 meters

[Table G-28](#) (page G-21) shows the cable connections for the sixth spine switch (R6-U1) when seven full racks are cabled together.

Table G-28 Leaf Switch Connections for Rack 6 in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-U1-P3A R6-IB3-P8B to R6-U1-P4A	3 meters
R6 IB3 to Rack 1	R6-IB3-P9B to R1-U1-P6A	10 meters
R6 IB3 to Rack 2	R6-IB3-P10A to R2-U1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-U1-P8A	10 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-U1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-U1-P10A	5 meters
R6 IB3 to Rack 7	R6-IB3-P9A to R7-U1-P5A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-U1-P3B R6-IB2-P8B to R6-U1-P4B	3 meters
R6 IB2 to Rack 1	R6-IB2-P9B to R1-U1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB3-P10A to R2-U1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-U1-P8B	10 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters
R6 IB2 to Rack 7	R6-IB3-P9A to R7-U1-P5B	5 meters

Table G-29 (page G-22) shows the cable connections for the seventh spine switch (R7-U1) when seven full racks are cabled together.

Table G-29 Leaf Switch Connections for Rack 7 in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R7 IB3 within Rack 7	R7-IB3-P8A to R7-U1-P3A R7-IB3-P8B to R7-U1-P4A	3 meters
R7 IB3 to Rack 1	R7-IB3-P9A to R1-U1-P5A	10 meters
R7 IB3 to Rack 2	R7-IB3-P9B to R2-U1-P6A	10 meters
R7 IB3 to Rack 3	R7-IB3-P10A to R3-U1-P7A	10 meters
R7 IB3 to Rack 4	R7-IB3-P10B to R4-U1-P8A	10 meters
R7 IB3 to Rack 5	R7-IB3-P11A to R5-U1-P9A	5 meters
R7 IB3 to Rack 6	R7-IB3-P11B to R6-U1-P10A	5 meters
R7 IB2 within Rack 7	R7-IB2-P8A to R7-U1-P3B R7-IB2-P8B to R7-U1-P4B	3 meters
R7 IB2 to Rack 1	R7-IB2-P9A to R1-U1-P5B	10 meters
R7 IB2 to Rack 2	R7-IB2-P9B to R2-U1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-U1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-U1-P8B	10 meters
R7 IB2 to Rack 5	R7-IB2-P11A to R5-U1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB2-P11B to R6-U1-P10B	5 meters

G.9 Eight-Rack Cabling

Table G-30 (page G-22) shows the cable connections for the first spine switch (R1-U1) when eight full racks are cabled together.

Table G-30 Leaf Switch Connections for Rack 1 in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A	3 meters
R1 IB3 to Rack 2	R1-IB3-P8B to R2-U1-P4A	5 meters
R1 IB3 to Rack 3	R1-IB3-P9A to R3-U1-P5A	5 meters
R1 IB3 to Rack 4	R1-IB3-P9B to R4-U1-P6A	10 meters
R1 IB3 to Rack 5	R1-IB3-P10A to R5-U1-P7A	10 meters
R1 IB3 to Rack 6	R1-IB3-P10B to R6-U1-P8A	10 meters
R1 IB3 to Rack 7	R1-IB3-P11A to R7-U1-P9A	10 meters
R1 IB3 to Rack 8	R1-IB3-P11B to R8-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
R1 IB2 to Rack 2	R1-IB2-P8B to R2-U1-P4B	5 meters
R1 IB2 to Rack 3	R1-IB2-P9A to R3-U1-P5B	5 meters
R1 IB2 to Rack 4	R1-IB2-P9B to R4-U1-P6B	10 meters

Table G-30 (Cont.) Leaf Switch Connections for Rack 1 in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R1 IB2 to Rack 5	R1-IB2-P10A to R5-U1-P7B	10 meters
R1 IB2 to Rack 6	R1-IB2-P10B to R6-U1-P8B	10 meters
R1 IB2 to Rack 7	R1-IB2-P11A to R7-U1-P9B	10 meters
R1 IB2 to Rack 8	R1-IB2-P11B to R8-U1-P10B	10 meters

[Table G-31](#) (page G-23) shows the cable connections for the second spine switch (R2-U1) when eight full racks are cabled together.

Table G-31 Leaf Switch Connections for Rack 2 in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P8B to R3-U1-P4A	5 meters
R2 IB3 to Rack 4	R2-IB3-P9A to R4-U1-P5A	5 meters
R2 IB3 to Rack 5	R2-IB3-P9B to R5-U1-P6A	10 meters
R2 IB3 to Rack 6	R2-IB3-P10A to R6-U1-P7A	10 meters
R2 IB3 to Rack 7	R2-IB3-P10B to R7-U1-P8A	10 meters
R2 IB3 to Rack 8	R2-IB3-P11A to R8-U1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P8B to R3-U1-P4B	5 meters
R2 IB2 to Rack 4	R2-IB2-P9A to R4-U1-P5B	5 meters
R2 IB2 to Rack 5	R2-IB2-P9B to R5-U1-P6B	10 meters
R2 IB2 to Rack 6	R2-IB2-P10A to R6-U1-P7B	10 meters
R2 IB2 to Rack 7	R2-IB2-P10B to R7-U1-P8B	10 meters
R2 IB2 to Rack 8	R2-IB2-P11A to R8-U1-P9B	10 meters

[Table G-32](#) (page G-23) shows the cable connections for the third spine switch (R3-U1) when eight full racks are cabled together.

Table G-32 Leaf Switch Connections for Rack 3 in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A	3 meters
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P8B to R4-U1-P4A	5 meters
R3 IB3 to Rack 5	R3-IB3-P9A to R5-U1-P5A	5 meters

Table G-32 (Cont.) Leaf Switch Connections for Rack 3 in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 to Rack 6	R3-IB3-P9B to R6-U1-P6A	10 meters
R3 IB3 to Rack 7	R3-IB3-P10A to R7-U1-P7A	10 meters
R3 IB3 to Rack 8	R3-IB3-P10B to R8-U1-P8A	10 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B	3 meters
R3 IB2 to Rack 1	R3-IB2-P11A to R1-U1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P8B to R4-U1-P4B	5 meters
R3 IB2 to Rack 5	R3-IB2-P9A to R5-U1-P5B	5 meters
R3 IB2 to Rack 6	R3-IB2-P9B to R6-U1-P6B	10 meters
R3 IB2 to Rack 7	R3-IB2-P10A to R7-U1-P7B	10 meters
R3 IB2 to Rack 8	R3-IB2-P10B to R8-U1-P8B	10 meters

[Table G-33](#) (page G-24) shows the cable connections for the fourth spine switch (R4-U1) when eight full racks are cabled together.

Table G-33 Leaf Switch Connections for Rack 4 in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A	3 meters
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P8B to R5-U1-P4A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9A to R6-U1-P5A	5 meters
R4 IB3 to Rack 7	R4-IB3-P9B to R7-U1-P6A	10 meters
R4 IB3 to Rack 8	R4-IB3-P10A to R8-U1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P8B to R5-U1-P4B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9A to R6-U1-P5B	5 meters
R4 IB2 to Rack 7	R4-IB2-P9B to R7-U1-P6B	10 meters
R4 IB2 to Rack 8	R4-IB2-P10A to R8-U1-P7B	10 meters

[Table G-34](#) (page G-25) shows the cable connections for the fifth spine switch (R5-U1) when eight full racks are cabled together.

Table G-34 Leaf Switch Connections for Rack 5 in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A	3 meters
R5 IB3 to Rack 1	R5-IB3-P10A to R1-U1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P8B to R6-U1-P4A	5 meters
R5 IB3 to Rack 7	R5-IB3-P9A to R7-U1-P5A	5 meters
R5 IB3 to Rack 8	R5-IB3-P9B to R8-U1-P6A	10 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B	3 meters
R5 IB2 to Rack 1	R5-IB2-P10A to R1-U1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB3-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P8B to R6-U1-P4B	5 meters
R5 IB2 to Rack 7	R5-IB2-P9A to R7-U1-P5B	5 meters
R5 IB2 to Rack 8	R5-IB2-P9B to R8-U1-P6B	10 meters

[Table G-35](#) (page G-25) shows the cable connections for the sixth spine switch (R6-U1) when eight full racks are cabled together.

Table G-35 Leaf Switch Connections for Rack 6 in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-U1-P3A	3 meters
R6 IB3 to Rack 1	R6-IB3-P9B to R1-U1-P6A	10 meters
R6 IB3 to Rack 2	R6-IB3-P10A to R2-U1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-U1-P8A	10 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-U1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-U1-P10A	5 meters
R6 IB3 to Rack 7	R6-IB3-P8B to R7-U1-P4A	5 meters
R6 IB3 to Rack 8	R6-IB3-P9A to R8-U1-P5A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-U1-P3B	3 meters
R6 IB2 to Rack 1	R6-IB2-P9B to R1-U1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB3-P10A to R2-U1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-U1-P8B	10 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters
R6 IB2 to Rack 7	R6-IB3-P8B to R7-U1-P4B	5 meters

Table G-35 (Cont.) Leaf Switch Connections for Rack 6 in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R6 IB2 to Rack 8	R6-IB2-P9A to R8-U1-P5B	5 meters

Table G-36 (page G-26) shows the cable connections for the seventh spine switch (R7-U1) when eight full racks are cabled together.

Table G-36 Leaf Switch Connections for Rack 7 in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R7 IB3 within Rack 7	R7-IB3-P8A to R7-U1-P3A	3 meters
R7 IB3 to Rack 1	R7-IB3-P9A to R1-U1-P5A	10 meters
R7 IB3 to Rack 2	R7-IB3-P9B to R2-U1-P6A	10 meters
R7 IB3 to Rack 3	R7-IB3-P10A to R3-U1-P7A	10 meters
R7 IB3 to Rack 4	R7-IB3-P10B to R4-U1-P8A	10 meters
R7 IB3 to Rack 5	R7-IB3-P11A to R5-U1-P9A	5 meters
R7 IB3 to Rack 6	R7-IB3-P11B to R6-U1-P10A	5 meters
R7 IB3 to Rack 8	R7-IB3-P8B to R8-U1-P4A	5 meters
R7 IB2 within Rack 7	R7-IB2-P8A to R7-U1-P3B	3 meters
R7 IB2 to Rack 1	R7-IB2-P9A to R1-U1-P5B	10 meters
R7 IB2 to Rack 2	R7-IB3-P9B to R2-U1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-U1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-U1-P8B	10 meters
R7 IB2 to Rack 5	R7-IB2-P11A to R5-U1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB3-P11B to R6-U1-P10B	5 meters
R7 IB2 to Rack 8	R7-IB3-P8B to R8-U1-P4B	5 meters

Table G-37 (page G-26) shows the cable connections for the eighth spine switch (R8-U1) when eight full racks are cabled together.

Table G-37 Leaf Switch Connections for Rack 8 in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R8 IB3 within Rack 8	R8-IB3-P8A to R8-U1-P3A	3 meters
R8 IB3 to Rack 1	R8-IB3-P8B to R1-U1-P4A	10 meters
R8 IB3 to Rack 2	R8-IB3-P9A to R2-U1-P5A	10 meters
R8 IB3 to Rack 3	R8-IB3-P9B to R3-U1-P6A	10 meters
R8 IB3 to Rack 4	R8-IB3-P10A to R4-U1-P7A	10 meters
R8 IB3 to Rack 5	R8-IB3-P10B to R5-U1-P8A	10 meters
R8 IB3 to Rack 6	R8-IB3-P11A to R6-U1-P9A	5 meters
R8 IB3 to Rack 7	R8-IB3-P11B to R7-U1-P10A	5 meters

Table G-37 (Cont.) Leaf Switch Connections for Rack 8 in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R8 IB2 within Rack 8	R8-IB2-P8A to R8-U1-P3B	3 meters
R8 IB2 to Rack 1	R8-IB2-P8B to R1-U1-P4B	10 meters
R8 IB2 to Rack 2	R8-IB3-P9A to R2-U1-P5B	10 meters
R8 IB2 to Rack 3	R8-IB2-P9B to R3-U1-P6B	10 meters
R8 IB2 to Rack 4	R8-IB2-P10A to R4-U1-P7B	10 meters
R8 IB2 to Rack 5	R8-IB2-P10B to R5-U1-P8B	10 meters
R8 IB2 to Rack 6	R8-IB3-P11A to R6-U1-P9B	5 meters
R8 IB2 to Rack 7	R8-IB3-P11B to R7-U1-P10B	5 meters

G.10 Nine-Rack Cabling

[Table G-38](#) (page G-27) shows the cable connections for the switches when nine full racks are cabled together. The cabling among racks 1 through 8 does not change.

Table G-38 Leaf Switch Connections for Rack 9

Leaf Switch	Connection
R9 IB3 to Rack 1	R9-IB3-P8B to R1-U1-P11A
R9 IB3 to Rack 2	R9-IB3-P9A to R2-U1-P11A
R9 IB3 to Rack 3	R9-IB3-P9B to R3-U1-P11A
R9 IB3 to Rack 4	R9-IB3-P10A to R4-U1-P11A
R9 IB3 to Rack 5	R9-IB3-P10B to R5-U1-P11A
R9 IB3 to Rack 6	R9-IB3-P11A to R6-U1-P11A
R9 IB3 to Rack 7	R9-IB3-P11B to R7-U1-P11A
R9 IB3 to Rack 8	R9-IB3-P9A to R8-U1-P11A
R9 IB2 to Rack 1	R9-IB2-P8B to R1-U1-P11B
R9 IB2 to Rack 2	R9-IB2-P9A to R2-U1-P11B
R9 IB2 to Rack 3	R9-IB2-P9B to R3-U1-P11B
R9 IB2 to Rack 4	R9-IB2-P10A to R4-U1-P11B
R9 IB2 to Rack 5	R9-IB2-P10B to R5-U1-P11B
R9 IB2 to Rack 6	R9-IB2-P11A to R6-U1-P11B
R9 IB2 to Rack 7	R9-IB2-P11B to R7-U1-P11B
R9 IB2 to Rack 8	R9-IB2-P8A to R8-U1-P11B

G.11 10-Rack Cabling

[Table G-39](#) (page G-28) shows the cable connections for the switches when 10 full racks are cabled together. The cabling among racks 1 through 8 does not change.

Table G-39 Leaf Switch Connections for Rack 10

Leaf Switch	Connection
R10 IB3 to Rack 1	R10-IB3-P8B to R1-U1-P12A
R10 IB3 to Rack 2	R10-IB3-P9A to R2-U1-P12A
R10 IB3 to Rack 3	R10-IB3-P9B to R3-U1-P12A
R10 IB3 to Rack 4	R10-IB3-P10A to R4-U1-P12A
R10 IB3 to Rack 5	R10-IB3-P10B to R5-U1-P12A
R10 IB3 to Rack 6	R10-IB3-P12A to R6-U1-P12A
R10 IB3 to Rack 7	R10-IB3-P12B to R7-U1-P12A
R10 IB3 to Rack 8	R10-IB3-P8A to R8-U1-P12A
R10 IB2 to Rack 1	R10-IB2-P8B to R1-U1-P12B
R10 IB2 to Rack 2	R10-IB2-P9A to R2-U1-P12B
R10 IB2 to Rack 3	R10-IB2-P9B to R3-U1-P12B
R10 IB2 to Rack 4	R10-IB2-P10A to R4-U1-P12B
R10 IB2 to Rack 5	R10-IB2-P10B to R5-U1-P12B
R10 IB2 to Rack 6	R10-IB2-P12A to R6-U1-P12B
R10 IB2 to Rack 7	R10-IB2-P12B to R7-U1-P12B
R10 IB2 to Rack 8	R10-IB2-P8A to R8-U1-P12B

G.12 11-Rack Cabling

[Table G-40](#) (page G-28) shows the cable connections for the switches when 11 full racks are cabled together. The cabling among racks 1 through 8 does not change.

Table G-40 Leaf Switch Connections for Rack 11

Leaf Switch	Connection
R11 IB3 to Rack 1	R11-IB3-P8B to R1-U1-P13A
R11 IB3 to Rack 2	R11-IB3-P9A to R2-U1-P13A
R11 IB3 to Rack 3	R11-IB3-P9B to R3-U1-P13A
R11 IB3 to Rack 4	R11-IB3-P10A to R4-U1-P13A
R11 IB3 to Rack 5	R11-IB3-P10B to R5-U1-P13A
R11 IB3 to Rack 6	R11-IB3-P13A to R6-U1-P13A
R11 IB3 to Rack 7	R11-IB3-P13B to R7-U1-P13A
R11 IB3 to Rack 8	R11-IB3-P8A to R8-U1-P13A
R11 IB2 to Rack 1	R11-IB2-P8B to R1-U1-P13B
R11 IB2 to Rack 2	R11-IB2-P9A to R2-U1-P13B
R11 IB2 to Rack 3	R11-IB2-P9B to R3-U1-P13B
R11 IB2 to Rack 4	R11-IB2-P10A to R4-U1-P13B
R11 IB2 to Rack 5	R11-IB2-P10B to R5-U1-P13B

Table G-40 (Cont.) Leaf Switch Connections for Rack 11

Leaf Switch	Connection
R11 IB2 to Rack 6	R11-IB2-P13A to R6-U1-P13B
R11 IB2 to Rack 7	R11-IB2-P13B to R7-U1-P13B
R11 IB2 to Rack 8	R11-IB2-P8A to R8-U1-P13B

G.13 12-Rack Cabling

[Table G-41](#) (page G-29) shows the cable connections for the switches when cabling twelve full racks together. The cabling between racks 1 through 8 does not change.

Table G-41 Leaf Switch Connections for Rack 12

Leaf Switch	Connection
R12 IB3 to Rack 1	R12-IB3-P8B to R1-U1-P14A
R12 IB3 to Rack 2	R12-IB3-P9A to R2-U1-P14A
R12 IB3 to Rack 3	R12-IB3-P9B to R3-U1-P14A
R12 IB3 to Rack 4	R12-IB3-P10A to R4-U1-P14A
R12 IB3 to Rack 5	R12-IB3-P10B to R5-U1-P14A
R12 IB3 to Rack 6	R12-IB3-P14A to R6-U1-P14A
R12 IB3 to Rack 7	R12-IB3-P14B to R7-U1-P14A
R12 IB3 to Rack 8	R12-IB3-P8A to R8-U1-P14A
R12 IB2 to Rack 1	R12-IB2-P8B to R1-U1-P14B
R12 IB2 to Rack 2	R12-IB2-P9A to R2-U1-P14B
R12 IB2 to Rack 3	R12-IB2-P9B to R3-U1-P14B
R12 IB2 to Rack 4	R12-IB2-P10A to R4-U1-P14B
R12 IB2 to Rack 5	R12-IB2-P10B to R5-U1-P14B
R12 IB2 to Rack 6	R12-IB2-P14A to R6-U1-P14B
R12 IB2 to Rack 7	R12-IB2-P14B to R7-U1-P14B
R12 IB2 to Rack 8	R12-IB2-P8A to R8-U1-P14B

G.14 13-Rack Cabling

[Table G-42](#) (page G-29) shows the cable connections for the switches when 13 full racks are cabled together. The cabling among racks 1 through 8 does not change.

Table G-42 Leaf Switch Connections for Rack 13

Leaf Switch	Connection
R13 IB3 to Rack 1	R13-IB3-P8B to R1-U1-P15A
R13 IB3 to Rack 2	R13-IB3-P9A to R2-U1-P15A

Table G-42 (Cont.) Leaf Switch Connections for Rack 13

Leaf Switch	Connection
R13 IB3 to Rack 3	R13-IB3-P9B to R3-U1-P15A
R13 IB3 to Rack 4	R13-IB3-P10A to R4-U1-P15A
R13 IB3 to Rack 5	R13-IB3-P10B to R5-U1-P15A
R13 IB3 to Rack 6	R13-IB3-P15A to R6-U1-P15A
R13 IB3 to Rack 7	R13-IB3-P15B to R7-U1-P15A
R13 IB3 to Rack 8	R13-IB3-P8A to R8-U1-P15A
R13 IB2 to Rack 1	R13-IB2-P8B to R1-U1-P15B
R13 IB2 to Rack 2	R13-IB2-P9A to R2-U1-P15B
R13 IB2 to Rack 3	R13-IB2-P9B to R3-U1-P15B
R13 IB2 to Rack 4	R13-IB2-P10A to R4-U1-P15B
R13 IB2 to Rack 5	R13-IB2-P10B to R5-U1-P15B
R13 IB2 to Rack 6	R13-IB2-P15A to R6-U1-P15B
R13 IB2 to Rack 7	R13-IB2-P15B to R7-U1-P15B
R13 IB2 to Rack 8	R13-IB2-P8A to R8-U1-P15B

G.15 14-Rack Cabling

[Table G-43](#) (page G-30) shows the cable connections for the switches when cabling 14 racks together. The cabling among racks 1 through 8 does not change.

Table G-43 Leaf Switch Connections for Rack 14

Leaf Switch	Connection
R14 IB3 to Rack 1	R14-IB3-P8B to R1-U1-P16A
R14 IB3 to Rack 2	R14-IB3-P9A to R2-U1-P16A
R14 IB3 to Rack 3	R14-IB3-P9B to R3-U1-P16A
R14 IB3 to Rack 4	R14-IB3-P10A to R4-U1-P16A
R14 IB3 to Rack 5	R14-IB3-P10B to R5-U1-P16A
R14 IB3 to Rack 6	R14-IB3-P16A to R6-U1-P16A
R14 IB3 to Rack 7	R14-IB3-P16B to R7-U1-P16A
R14 IB3 to Rack 8	R14-IB3-P8A to R8-U1-P16A
R14 IB2 to Rack 1	R14-IB2-P8B to R1-U1-P16B
R14 IB2 to Rack 2	R14-IB2-P9A to R2-U1-P16B
R14 IB2 to Rack 3	R14-IB2-P9B to R3-U1-P16B
R14 IB2 to Rack 4	R14-IB2-P10A to R4-U1-P16B
R14 IB2 to Rack 5	R14-IB2-P10B to R5-U1-P16B
R14 IB2 to Rack 6	R14-IB2-P16A to R6-U1-P16B
R14 IB2 to Rack 7	R14-IB2-P16B to R7-U1-P16B

Table G-43 (Cont.) Leaf Switch Connections for Rack 14

Leaf Switch	Connection
R14 IB2 to Rack 8	R14-IB2-P8A to R8-U1-P16B

G.16 15-Rack Cabling

[Table G-44](#) (page G-31) shows the cable connections for the switches when 15 racks are cabled together. The cabling among racks 1 through 8 does not change.

Table G-44 Leaf Switch Connections for Rack 15

Leaf Switch	Connection
R15 IB3 to Rack 1	R15-IB3-P8B to R1-U1-P17A
R15 IB3 to Rack 2	R15-IB3-P9A to R2-U1-P17A
R15 IB3 to Rack 3	R15-IB3-P9B to R3-U1-P17A
R15 IB3 to Rack 4	R15-IB3-P10A to R4-U1-P17A
R15 IB3 to Rack 5	R15-IB3-P10B to R5-U1-P17A
R15 IB3 to Rack 6	R15-IB3-P17A to R6-U1-P17A
R15 IB3 to Rack 7	R15-IB3-P17B to R7-U1-P17A
R15 IB3 to Rack 8	R15-IB3-P8A to R8-U1-P17A
R15 IB2 to Rack 1	R15-IB2-P8B to R1-U1-P17B
R15 IB2 to Rack 2	R15-IB2-P9A to R2-U1-P17B
R15 IB2 to Rack 3	R15-IB2-P9B to R3-U1-P17B
R15 IB2 to Rack 4	R15-IB2-P10A to R4-U1-P17B
R15 IB2 to Rack 5	R15-IB2-P10B to R5-U1-P17B
R15 IB2 to Rack 6	R15-IB2-P17A to R6-U1-P17B
R15 IB2 to Rack 7	R15-IB2-P17B to R7-U1-P17B
R15 IB2 to Rack 8	R15-IB2-P8A to R8-U1-P17B

G.17 16-Rack Cabling

[Table G-45](#) (page G-31) shows the cable connections for the switches when 16 full racks are cabled together. The cabling among racks 1 through 8 does not change.

Table G-45 Leaf Switch Connections for Rack 16

Leaf Switch	Connection
R16 IB3 to Rack 1	R16-IB3-P8B to R1-U1-P0A
R16 IB3 to Rack 2	R16-IB3-P9A to R2-U1-P0A
R16 IB3 to Rack 3	R16-IB3-P9B to R3-U1-P0A
R16 IB3 to Rack 4	R16-IB3-P10A to R4-U1-P0A

Table G-45 (Cont.) Leaf Switch Connections for Rack 16

Leaf Switch	Connection
R16 IB3 to Rack 5	R16-IB3-P10B to R5-U1-P0A
R16 IB3 to Rack 6	R16-IB3-P0A to R6-U1-P0A
R16 IB3 to Rack 7	R16-IB3-P0B to R7-U1-P0A
R16 IB3 to Rack 8	R16-IB3-P8A to R8-U1-P0A
R16 IB2 to Rack 1	R16-IB2-P8B to R1-U1-P0B
R16 IB2 to Rack 2	R16-IB2-P9A to R2-U1-P0B
R16 IB2 to Rack 3	R16-IB2-P9B to R3-U1-P0B
R16 IB2 to Rack 4	R16-IB2-P10A to R4-U1-P0B
R16 IB2 to Rack 5	R16-IB2-P10B to R5-U1-P0B
R16 IB2 to Rack 6	R16-IB2-P0A to R6-U1-P0B
R16 IB2 to Rack 7	R16-IB2-P0B to R7-U1-P0B
R16 IB2 to Rack 8	R16-IB2-P8A to R8-U1-P0B

G.18 17-Rack Cabling

[Table G-46](#) (page G-32) shows the cable connections for the switches when 17 racks are cabled together. The cabling among racks 1 through 8 does not change.

Table G-46 Leaf Switch Connections for Rack 17

Leaf Switch	Connection
R17 IB3 to Rack 1	R17-IB3-P8B to R1-U1-P1A
R17 IB3 to Rack 2	R17-IB3-P9A to R2-U1-P1A
R17 IB3 to Rack 3	R17-IB3-P9B to R3-U1-P1A
R17 IB3 to Rack 4	R17-IB3-P10A to R4-U1-P1A
R17 IB3 to Rack 5	R17-IB3-P10B to R5-U1-P1A
R17 IB3 to Rack 6	R17-IB3-P1A to R6-U1-P1A
R17 IB3 to Rack 7	R17-IB3-P1B to R7-U1-P1A
R17 IB3 to Rack 8	R17-IB3-P8A to R8-U1-P1A
R17 IB2 to Rack 1	R17-IB2-P8B to R1-U1-P1B
R17 IB2 to Rack 2	R17-IB2-P9A to R2-U1-P1B
R17 IB2 to Rack 3	R17-IB2-P9B to R3-U1-P1B
R17 IB2 to Rack 4	R17-IB2-P10A to R4-U1-P1B
R17 IB2 to Rack 5	R17-IB2-P10B to R5-U1-P1B
R17 IB2 to Rack 6	R17-IB2-P1A to R6-U1-P1B
R17 IB2 to Rack 7	R17-IB2-P1B to R7-U1-P1B
R17 IB2 to Rack 8	R17-IB2-P8A to R8-U1-P1B

G.19 18-Rack Cabling

Table G-47 (page G-33) shows the cable connections for the switches when 18 racks are cabled together. The cabling among racks 1 through 8 does not change.

Table G-47 Leaf Switch Connections for Rack 18

Leaf Switch	Connection
R18 IB3 to Rack 1	R18-IB3-P8B to R1-U1-P2A
R18 IB3 to Rack 2	R18-IB3-P9A to R2-U1-P2A
R18 IB3 to Rack 3	R18-IB3-P9B to R3-U1-P2A
R18 IB3 to Rack 4	R18-IB3-P10A to R4-U1-P2A
R18 IB3 to Rack 5	R18-IB3-P10B to R5-U1-P2A
R18 IB3 to Rack 6	R18-IB3-P2A to R6-U1-P2A
R18 IB3 to Rack 7	R18-IB3-P2B to R7-U1-P2A
R18 IB3 to Rack 8	R18-IB3-P8A to R8-U1-P2A
R18 IB2 to Rack 1	R18-IB2-P8B to R1-U1-P2B
R18 IB2 to Rack 2	R18-IB2-P9A to R2-U1-P2B
R18 IB2 to Rack 3	R18-IB2-P9B to R3-U1-P2B
R18 IB2 to Rack 4	R18-IB2-P10A to R4-U1-P2B
R18 IB2 to Rack 5	R18-IB2-P10B to R5-U1-P2B
R18 IB2 to Rack 6	R18-IB2-P2A to R6-U1-P2B
R18 IB2 to Rack 7	R18-IB2-P2B to R7-U1-P2B
R18 IB2 to Rack 8	R18-IB2-P8A to R8-U1-P2B

Index

A

access route requirements, [2-6](#)
acoustic levels, [2-1](#)
ACSLs servers, [5-26](#)
add
 racli Command, [A-4](#)
add acfs
 racli Command, [A-4](#)
add cloud_storage
 racli Command, [A-4](#)
add credential
 racli Command, [A-5](#)
add credential_wallet
 racli Command, [A-6](#)
add keystore
 racli Command, [A-7](#)
add nfs
 racli Command, [A-7](#)
add storage_cells
 racli Command, [A-8](#)
add storage_disks
 racli Command, [A-8](#)
add vpc_user
 racli Command, [A-9](#)
adding
 Sun Datacenter InfiniBand Switch 36 switch,
 [8-2](#)
administration network configuration, [5-8](#)
Africa power requirements, [2-10](#)
airflow, [2-13](#), [2-15](#)
alert messages
 battery learn cycle, [12-30](#)
 HALRT-00002 to HALRT-00500, [C-2](#)
 HALRT-00513 to HALRT-00600, [C-5](#)
 HALRT-00616 to HALRT-00700, [C-6](#)
 HALRT-00717 to HALRT-00800, [C-6](#)
 HALRT-00912 to HALRT-01000, [C-7](#)
 HALRT-01015 to HALRT-01100, [C-8](#)
 HALRT-01209 to HALRT-01300, [C-28](#)
 HALRT-01301 to HALRT-01400, [C-32](#)
 HALRT-02001 to HALRT-02008, [C-34](#)
 HALRT-02101 to HALRT-02200, [C-36](#)
 HALRT-10001 to HALRT-10500, [C-39](#)
 HALRT-10503 to HALRT-11000, [C-44](#)

alert messages (*continued*)
 HALRT-11001 to HALRT-11100, [C-45](#)
 hardware, [C-1](#)
alter
 racli Command, [A-10](#)
alter cloud_storage
 racli Command, [A-10](#)
alter credential_wallet
 racli Command, [A-11](#)
alter keystore
 racli Command, [A-11](#)
alter passphrase
 racli Command, [A-12](#)
alter rasys_user
 racli Command, [A-12](#)
alter vpc_user
 racli Command, [A-13](#)
altitude requirements, [2-13](#)
ambient conditions, [2-13](#)
ambient relative humidity range, [2-14](#)
ambient temperature ranges, [2-13](#)
Asia Pacific power requirements, [2-10](#)
ASR,
 about, [4-1](#)
 configuration, [4-2](#)
 troubleshooting, [4-5](#)
ASR Manager
 configuration, [5-25](#)
 installation, [4-3](#)
ASR network connections, [4-2](#)
ASR trap destinations, [4-4](#)
Auto Service Request, [4-1](#)
 configuration, [5-24](#)
 See also ASR

B

batteries
 learn cycle, [12-30](#)
 life expectancy, [12-7](#)
 monitoring, [12-8](#)
 replacing, [12-8](#)
 replacing disk controller BBUs, [12-43](#)
bend radius requirements, [7-5](#)
BONDIB0, [3-4](#)

buttons, Deployment Assistant, [5-2](#)

C

cable management arm, [7-5](#)

cable management arms
replacement part, [D-35](#)

cables

color coding, [E-1](#), [F-1](#)
replacement, [D-35](#)

cabling

bend radius, [7-5](#)
bundling InfiniBand, [7-5](#)
estimating cable length, [7-1](#)
InfiniBand precautions, [7-4](#)
multi-rack, [8-4](#), [G-1](#)
new storage servers, [7-15](#)

cabling requirements

multirack, [G-2](#)

cell alert configuration, [5-19](#)

CELLBOOT USB flash drive, [12-45](#)

channel adapter, [1-2](#), [1-3](#)

checking

cell configuration, [7-2](#)
cell disk configuration, [7-2](#)
firmware version, [7-3](#)
flash logs, [7-2](#)
HugePages memory configuration, [7-2](#)
image history, [7-2](#)
InfiniBand switch information, [7-2](#)
IORM plans, [7-2](#)
IP addresses, [7-2](#)

checkip script, [5-3](#), [5-4](#), [11-2](#)

checklist postinstallation, [6-12](#)

circuit breakers, [2-12](#)

circuit breakers for PDU, [6-14](#), [6-15](#)

Cisco switch

configuring default gateway, [10-12](#)
configuring host name, [10-12](#)
default port speed, [10-11](#)
model, [1-2](#), [1-3](#)
replacement parts, [D-34](#)
setting password, [10-12](#), [13-4](#)

cloning Oracle Database homes

Linux-based servers, [12-22](#)

cloning Oracle Grid infrastructure

Linux-based servers, [12-20](#)

Cloud Control agent, [5-25](#)

Cloud Control discovery, [11-5](#)

cluster configuration, [5-15](#)

cluster verification utility, [12-15](#)

CMA guidelines, [7-5](#)

compute servers

components, [1-3](#)
model, [1-2](#), [1-3](#)

compute servers (*continued*)

replacement parts, [D-13](#), [D-18](#)

configuration, [5-1](#)

configuration files

descriptions, [5-4](#)
generating, [5-2](#)
revising, [5-3](#)

Configuration Manager configuration, [5-23](#)

configuring

cell alerts, [5-19](#)
InfiniBand interface, [10-5](#)
PDUs, [10-17](#)

configuring ASR, [4-2](#)

connecting power cords, [6-12](#)

cooling, [2-14](#)

cord connections, power, [6-12](#)

CSI, [7-8](#)

Customer Support Identifier, [7-8](#)

D

database configuration, [5-16](#)

Oracle Database

users and groups, [5-16](#)

default IP addresses, [3-10](#)

default passwords, [B-1](#)

default serial port speed

Cisco switch, [10-11](#)

DELTA disk group, [5-16](#)

Deployment Assistant, [5-1](#)

dimensions, [2-6](#)

disable

racli Command, [A-13](#)

disable cloud_storage

racli Command, [A-13](#)

disable keystore

racli Command, [A-14](#)

disable nfs

racli Command, [A-14](#)

disk controller battery BBUs, [12-43](#)

disk controllers, [12-13](#)

disk drives

server status indicators, [D-3](#)

disk failures

replacing physical disks, [12-31](#)

disk groups, [5-16](#)

DNS, [3-10](#)

domain name system, [3-10](#)

E

earth ground cable, [6-10](#)

electrostatic discharge, [2-14](#)

email accounts, My Oracle Support (ASR), [4-1](#)

enable

- enable (*continued*)
 - racli Command, [A-14](#)
 - enable cloud_storage
 - racli Command, [A-14](#)
 - enable keystore
 - racli Command, [A-15](#)
 - enable nfs
 - racli Command, [A-15](#)
 - Enterprise Manager agent, [5-25](#)
 - environmental requirements, [2-1](#)
 - error messages
 - contacting Oracle Support Services, [C-2](#)
 - Ethernet cabling connections, [E-4](#), [F-4](#)
 - Ethernet card configuration, [5-7](#)
 - Ethernet switch
 - model, [1-2](#), [1-3](#)
 - replacement part, [D-34](#)
 - exachk utility, [7-2](#)
 - Exadata Deployment Assistant, [5-1](#)
 - starting, [5-3](#)
 - Exadata software installation, [11-3](#)
 - export
 - racli Command, [A-16](#)
 - export appliance
 - racli Command, [A-16](#)
- ## F
-
- fan modules
 - server status indicators, [D-3](#)
 - switch status indicators, [D-6](#)
 - FDOM, [12-37](#)
 - fiber channel adapters, [3-7](#)
 - installation, [6-11](#)
 - fibre channel network, [3-3](#)
 - firewalls
 - port assignments, [3-11](#)
 - firmware levels, [12-32](#), [12-33](#)
 - flash disks, [12-37](#)
 - flexible configuration, [5-5](#)
 - floor load requirements, [2-7](#)
- ## G
-
- gateway configuration, [5-8](#), [5-11](#)
 - Grid Control agent, [5-25](#)
 - ground cable, [6-10](#)
 - grounding to earth, [2-13](#)
 - groups
 - database, [5-16](#)
- ## H
-
- HALRT-00002 to HALRT-00500 alert messages, [C-2](#)
 - HALRT-00513 to HALRT-00600 alert messages, [C-5](#)
 - HALRT-00616 to HALRT-00700 alert messages, [C-6](#)
 - HALRT-00717 to HALRT-00800 alert messages, [C-6](#)
 - HALRT-00912 to HALRT-01000 alert messages, [C-7](#)
 - HALRT-01015 to HALRT-01100 alert messages, [C-8](#)
 - HALRT-01209 to HALRT-01300 alert messages, [C-28](#)
 - HALRT-01301 to HALRT-01400 alert messages, [C-32](#)
 - HALRT-02001 to HALRT-02008 alert messages, [C-34](#)
 - HALRT-02101 to HALRT-02200 alert messages, [C-36](#)
 - HALRT-10001 to HALRT-10500 alert messages, [C-39](#)
 - HALRT-10503 to HALRT-11000 alert messages, [C-44](#)
 - HALRT-11001 to HALRT-11100 alert messages, [C-45](#)
 - hardware alert messages, [C-1](#)
 - hardware components, [1-2](#)
 - hardware configuration, [5-5](#)
 - hardware failures, automatic reporting, [4-1](#)
 - hardware installation, [6-1](#)
 - hardware modifications, [1-9](#)
 - help
 - racli Command, [A-16](#)
 - high-voltage PDUs, Europe power requirements, [2-10](#)
 - humidity range, [2-14](#)
 - humidity requirements, [2-13](#)
- ## I
-
- identifying
 - predictive failure flash disk, [12-39](#)
 - ILOM cabling connections, [E-3](#), [F-3](#)
 - ILOM configuration, [5-8](#)
 - InfiniBand
 - private network, [3-4](#)
 - switch time zone, [13-9](#)
 - using for ingest, [9-1](#)
 - InfiniBand cables
 - spares, [G-2](#)
 - InfiniBand cabling, [E-7](#), [F-7](#)
 - InfiniBand network configuration, [5-12](#)
 - InfiniBand switch cabling, [E-8](#), [F-7](#)
 - InfiniBand switches
 - LEDs, [D-4](#)
 - model, [1-2](#), [1-3](#)

ingest network, [3-4](#)
 ingest network configuration, [5-10](#), [5-17](#)
 ingest over InfiniBand, [9-1](#)
 inspecting the rack, [6-12](#)
 install
 racli Command, [A-16](#)
 install appliance
 racli Command, [A-16](#)
 install okv_endpoint
 racli Command, [A-17](#)
 install utility
 step descriptions, [11-7](#)
 syntax, [11-6](#)
 installation
 safety guidelines, [7-4](#)
 installation directory, [11-2](#)
 InstallationTemplate file, [5-4](#)
 installing software, [11-3](#)
 installing Sun Datacenter InfiniBand Switch 36
 switch, [8-2](#)
 IP address configuration, [5-8](#), [5-10](#), [5-12](#), [5-13](#),
 [5-17](#)
 IP addresses
 factory default IP addresses, [3-10](#)
 Recovery Appliance totals, [2-2-2-5](#)
 ipconf utility, [12-47](#)
 IPMItool
 resetting the Service Processor, [12-12](#)

J

Japan power requirements, [2-8](#)

L

language support, [5-16](#)
 leaf switches, [10-8](#)
 LED color codes, [D-2](#), [D-4](#)
 LED status, [D-2](#)
 LED verification, [6-15](#)
 leveling feet, [6-7](#), [6-8](#)
 license restrictions, [1-9](#)
 list
 racli Command, [A-18](#)
 list cloud_storage
 racli Command, [A-18](#)
 list nfs
 racli Command, [A-18](#)
 list okv_endpoint
 racli Command, [A-19](#)
 listing PCI slot numbers, [12-37](#)
 locating
 PDU circuit breakers, [6-14](#), [6-15](#)
 low-voltage PDUs, [2-8](#)

M

maintenance access requirements, [2-7](#)
 management network, [3-3](#)
 media managers
 installation of third-party, [1-9](#)
 Middle East power requirements, [2-10](#)
 model numbers, [1-2](#)
 moving drives, [12-34](#)
 multitrack cabling
 cable requirements, [G-2](#)
 My Oracle Support email account, [4-1](#)

N

name prefix, [5-5](#)
 naming format mask, [5-9](#)
 naming format masks, [5-11-5-13](#)
 NET0, [3-3](#)
 network cables, [2-15](#)
 network components, [3-1](#)
 network configuration, [5-6](#)
 administration, [5-8](#)
 network connections
 ASR, [4-2](#)
 network diagram, [3-5](#)
 network preparation, [3-1](#)
 Network Time Protocol, [13-8](#)
 networks
 fibre channel SAN, [3-3](#)
 InfiniBand, [3-4](#)
 ingest, [3-4](#)
 management, [3-3](#)
 replication, [3-3](#)
 nonoperating requirements, [2-13](#)
 North America power requirements, [2-8](#)
 NTP servers, [13-8](#)

O

OneCommand download, [5-2](#)
 OneCommand platforms, [5-2](#)
 operating requirements, [2-13](#)
 Oracle Configuration Manager configuration,
 [5-23](#)
 Oracle Enterprise Manager Cloud Control
 discovery, [11-5](#)
 Oracle Enterprise Manager installation, [1-9](#)
 Oracle Management Server, [5-26](#), [11-2](#)
 Oracle Maximum Availability Architecture, [12-27](#)
 Oracle Secure Backup, [1-9](#)
 Oracle Server X5-2
 replacement parts, [D-13](#)
 Oracle Server X5-2L
 replacement parts, [D-24](#)

Oracle Server X6-2
 replacement parts, [D-13](#)
 Oracle Server X6-2L
 replacement parts, [D-24](#)
 Oracle Support Services
 contacting, [C-2](#)

P

passwords

changing

Cisco Ethernet password, [13-4](#)
 database server passwords, [13-1](#)
 ILOM password, [13-3](#)
 InfiniBand switch passwords, [13-3](#)
 PDU passwords, [13-2](#)
 storage server passwords, [13-2](#)

patch

racli Command, [A-19](#)

patch appliance

racli Command, [A-19](#)

PCI slot numbers, [12-37](#)

PCIe cards, [12-37](#)

PDU cabling, [E-5](#), [E-6](#), [F-5](#), [F-6](#)

PDUs

circuit breakers, [6-14](#), [6-15](#)
 configuration, [10-17](#)
 connecting hardware, [2-9–2-11](#)
 PDU-A location, [12-9](#)
 PDU-B location, [12-9](#)
 replacement parts, [D-34](#)
 single phase, [2-8](#)
 three-phase, [2-9](#), [2-11](#)

perforated tiles, [2-15](#)

physical disks

removing in error, [12-36](#)
 replacing, [12-31](#)
 status, [12-27](#)

platinum.csv, [5-4](#)

port assignments, [3-11](#)

postinstallation checklist, [6-12](#)

power cord connections, [6-12](#)

power cord routing, [6-13](#)

power off

emergency, [12-2](#)
 switch, [12-2](#)

power on

sequence for servers, [12-5](#)
 using ILOM, [12-6](#)

power requirements, [2-8](#), [2-12](#)

power supplies

server status indicators, [D-4](#)
 switch status indicators, [D-6](#)

powering the system, [6-12](#)

preconf_rack_n.csv, [5-4](#)

problem reporting, [4-1](#)

proxy configuration

for ASR, [5-25](#)
 for Oracle Configuration Manager, [5-24](#)

Q

QLogic card, [1-3](#)

QLogic fiber channel adapters, [6-11](#)

R

rack component prefix, [5-5](#)

rack components, [1-2](#)

rack configuration, [5-1](#)

rack dimensions, [2-6](#)

rack installation, [6-1](#)

rack layout

X4 full rack, [F-1](#)
 X4 minimum configuration, [F-2](#)
 X7, X6, or X5 full rack, [E-1](#)
 X7, X6, or X5 minimum configuration, [E-2](#)

rack weight, [2-6](#)

racli Commands

add, [A-4](#)
 add acfs, [A-4](#)
 add cloud_storage, [A-4](#)
 add credential, [A-5](#)
 add credential_wallet, [A-6](#)
 add keystore, [A-7](#)
 add nfs, [A-7](#)
 add storage_cells, [A-8](#)
 add storage_disks, [A-8](#)
 add vpc_user, [A-9](#)
 alter, [A-10](#)
 alter cloud_storage, [A-10](#)
 alter credential_wallet, [A-11](#)
 alter keystore, [A-11](#)
 alter passphrase, [A-12](#)
 alter rasys_user, [A-12](#)
 alter vpc_user, [A-13](#)
 disable, [A-13](#)
 disable cloud_storage, [A-13](#)
 disable keystore, [A-14](#)
 disable nfs, [A-14](#)
 enable, [A-14](#)
 enable cloud_storage, [A-14](#)
 enable keystore, [A-15](#)
 enable nfs, [A-15](#)
 export, [A-16](#)
 export appliance, [A-16](#)
 help, [A-16](#)
 install, [A-16](#)
 install appliance, [A-16](#)
 install okv_endpoint, [A-17](#)

racli Commands (*continued*)

- list, [A-18](#)
- list cloud_storage, [A-18](#)
- list nfs, [A-18](#)
- list okv_endpoint, [A-19](#)
- patch, [A-19](#)
- patch appliance, [A-19](#)
- remove, [A-20](#)
- remove acfs, [A-20](#)
- remove cloud_storage, [A-20](#)
- remove keystore, [A-21](#)
- remove nfs, [A-21](#)
- remove vpc_user, [A-21](#)
- start, [A-22](#)
- start nfs, [A-22](#)
- status, [A-22](#)
- status appliance, [A-22](#)
- status cloud_storage, [A-22](#)
- status disk_group, [A-23](#)
- status keystore, [A-23](#)
- status nfs, [A-23](#)
- status okv_endpoint, [A-23](#)
- stop, [A-24](#)
- stop nfs, [A-24](#)
- uninstall, [A-24](#)
- uninstall appliance, [A-24](#)
- uninstall okv_endpoint, [A-25](#)
- update, [A-26](#)
- update storage_location, [A-26](#)
- update support_tools, [A-26](#)
- upgrade, [A-26](#)
- upgrade appliance, [A-26](#)
- version, [A-27](#)

Recovery Appliance software installation, [11-3](#)

recreating

- CELLBOOT USB flash drive, [12-49](#)

reenabling a rejected disk, [12-36](#)

relative humidity, [2-14](#)

relocating voting disks

- Linux-based servers, [12-15](#)

remove

- racli Command, [A-20](#)

remove acfs

- racli Command, [A-20](#)

remove cloud_storage

- racli Command, [A-20](#)

remove keystore

- racli Command, [A-21](#)

remove nfs

- racli Command, [A-21](#)

remove vpc_user

- racli Command, [A-21](#)

removing a disk in error, [12-36](#)

repair categories, [D-1](#)

replacement parts, [D-1](#)

replacing

- disk controller BBUs, [12-43](#)
- physical disks, [12-31](#)

replication network, [3-3](#)

- SCAN configuration, [5-18](#)
- VIP configuration, [5-18](#)

replication network configuration, [5-13](#), [5-17](#)

rescue flash drive, [12-45](#)

resetting

- ILOM, [12-11](#)

S

safety guidelines, [6-1](#)

SAN network, [3-3](#)

SCAN addresses, [3-4](#), [3-10](#)

SCAN configuration, [5-18](#)

- replication, [5-18](#)

server name prefix, [5-5](#)

servers

- LED color codes, [D-2](#), [D-4](#)

setting

- battery learn cycle, [12-30](#)
- Cisco switch time zone, [10-14](#)
- InfiniBand
 - switch date, [13-9](#)
 - switch time zone, [13-9](#)

shipping brackets, [6-2](#)

shipping kit, [6-2](#), [6-3](#)

shipping weights, [2-6](#)

single-phase PDUs, [2-8](#), [2-10](#)

SMTP alerts, [4-1](#)

SNMP alerts, [4-1](#)

software installation, [11-3](#)

software modifications, [1-10](#)

South America power requirements, [2-8](#)

spares kit, [1-7](#)

spine switches

- rack location, [10-8](#)

stabilizing rack, [6-7](#)

stabilizing ZDLRA Rack, [6-8](#)

start

- racli Command, [A-22](#)

start nfs

- racli Command, [A-22](#)

status

- component LEDs, [D-2](#)
- physical disks, [12-27](#)
- racli Command, [A-22](#)

status appliance

- racli Command, [A-22](#)

status cloud_storage

- racli Command, [A-22](#)

status disk_group

- racli Command, [A-23](#)

status keystore
 racli Command, [A-23](#)

status nfs
 racli Command, [A-23](#)

status okv_endpoint
 racli Command, [A-23](#)

stop
 racli Command, [A-24](#)

stop nfs
 racli Command, [A-24](#)

storage servers
 adding to a rack, [7-1](#)
 components, [1-5](#)
 increasing storage location size, [7-18](#)
 model, [1-2](#), [1-3](#)
 replacement parts, [D-24](#), [D-30](#)

Subnet Manager Master, [10-8](#)

Sun Datacenter InfiniBand Switch 36
 LEDs, [D-4](#)
 replacement parts, [D-33](#)

Sun Server X4-2
 replacement parts, [D-18](#)

Sun Server X4-2L
 replacement parts, [D-30](#)

support accounts, [4-2](#)

Support Hub, [5-24](#)

support policies, [5-24](#)

switch name prefix, [5-5](#)

switch-to-server cabling, [E-8](#), [F-7](#)

switch-to-switch cabling, [E-7](#), [F-7](#)

T

Taiwan power requirements, [2-8](#)

tape backup
 network connections, [3-7](#)
 Oracle red stack, [1-8](#)
 third-party restrictions, [1-8](#)

tape library
 configuration for Oracle red stack, [5-26](#)

tape network, [3-3](#)

temperature ranges, [2-13](#)

temperature requirements, [2-13](#)

temporary_config.xml, [5-4](#)

three-phase PDUs, [2-9](#), [2-11](#)

tile flooring, [2-15](#)

time zones, [10-14](#)

trap destinations (ASR), [4-4](#)

U

underfloor ventilation, [2-14](#)

uninstall
 racli Command, [A-24](#)

uninstall appliance
 racli Command, [A-24](#)

uninstall okv_endpoint
 racli Command, [A-25](#)

update
 racli Command, [A-26](#)

update storage_location
 racli Command, [A-26](#)

update support_tools
 racli Command, [A-26](#)

upgrade
 racli Command, [A-26](#)

upgrade appliance
 racli Command, [A-26](#)

USB devices, [1-10](#)

USB flash drive, [12-45](#)

users
 database, [5-16](#)

utilities
 cluster verification, [12-15](#)
 exachk, [7-2](#)
 ipconf, [12-47](#)
 verify-topology, [10-24](#)

V

ventilation, [2-14](#)

verify-topology utility, [10-24](#)

verifying
 database server configuration, [12-13](#)
 firmware, [12-32](#), [12-33](#), [D-7](#), [D-14](#), [D-16](#),
 [D-19](#), [D-21](#), [D-26](#), [D-28](#), [D-32](#)
 hardware, [D-7](#), [D-14](#), [D-16](#), [D-19](#), [D-21](#),
 [D-26](#), [D-28](#), [D-32](#)
 InfiniBand topology, [10-24](#)
 virtual drive configuration, [12-13](#)

version
 racli Command, [A-27](#)

VIP addresses, [3-4](#), [3-10](#)

VIP network configuration, [5-18](#)

virtual local area networks, [3-3](#)

VLAN tagging, [3-9](#)

VLANs, [3-3](#), [10-11](#)

W

warranty restrictions, [1-9](#)

weight, [2-6](#)

WorkDir subdirectory, [11-2](#)