

Zero Data Loss Recovery Appliance

Owner's Guide



Release 23.1
F72079-03
June 2024

ORACLE®

Zero Data Loss Recovery Appliance Owner's Guide, Release 23.1

F72079-03

Copyright © 2014, 2024, Oracle and/or its affiliates.

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

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

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

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

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

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

Contents

Recovery Appliance Preface

Audience	xxix
Documentation Accessibility	xxix
Conventions	xxix

Changes in Zero Data Loss Recovery Appliance Owner's Guide Release 23.1

1 Introduction to the Recovery Appliance

About the Recovery Appliance	1-1
Ransomware and Cybersecurity	1-2
Recovery Appliance Rack Components	1-4
Compute Server Components	1-8
Storage Server Components	1-14
About Tape Backup Infrastructure	1-17
Oracle Recommended Stack	1-18
Third-Party Tape Systems	1-18
Restrictions on Modifying Recovery Appliance	1-19

Part I Preinstallation

2 Site Requirements for Recovery Appliance and Oracle Exadata Storage Expansion Rack

Reviewing Safety Guidelines	2-1
Ensuring That the Site is Ready	2-2
General Environmental Requirements	2-2
General Environmental Requirements for ZDLRA Rack X6 and Later	2-3
Space Requirements	2-3
Space Requirements for Racks up to Recovery Appliance X6	2-3
Receiving, Unpacking, and Access Route Requirements	2-3

Rack Weights for Recovery Appliance	2-4
Maintenance Access Requirements	2-5
Flooring Requirements	2-5
Electrical Power Requirements	2-6
PDU Power Requirements	2-6
Low-Voltage 15 kVA Single Phase PDU for North America, South America, Japan and Taiwan	2-6
Low-Voltage 22 kVA Single Phase PDU for North America, South America, Japan and Taiwan	2-7
Low-Voltage 15 kVA Three Phase PDU for North America, South America, Japan and Taiwan	2-8
Low-Voltage 24 kVA Three Phase PDU for North America, South America, Japan and Taiwan	2-8
High-Voltage 15 kVA Single Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	2-9
High-Voltage 22 kVA Single Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	2-10
High-Voltage 15 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	2-10
High-Voltage 24 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	2-11
Facility Power Requirements	2-11
Circuit Breaker Requirements	2-12
Electrical Grounding Guidelines	2-12
Temperature and Humidity Requirements	2-12
Ventilation and Cooling Requirements	2-14
Network Connection and IP Address Requirements for ZDLRA Rack	2-15
Network Connection Requirements for ZDLRA Rack	2-15
DNS Configuration for ZDLRA Rack	2-16
IP Address Requirements for ZDLRA Rack	2-16
Dedicated Backup Network between Oracle Exadata and Oracle Zero Data Loss Recovery Appliance	2-17

3 Preparing Your Networks for Recovery Appliance

Overview of Network Requirements	3-1
About the Network Components and Interfaces	3-4
Example of Network Connections for Recovery Appliance	3-8
Connecting Recovery Appliance Rack Components to the Networks	3-10
Connecting Recovery Appliance to a Tape Library	3-10
Oracle Recommended Stack	3-10
Third-Party Tape Systems	3-11
Using Network VLAN Tagging with Recovery Appliance	3-12
Registering Recovery Appliance in the Domain Name System	3-13

Factory IP Address Settings	3-13
Port Assignments When Using a Firewall	3-14
Proxy Servers	3-17

4 Setting Up Auto Service Request

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

5 Using Oracle Exadata Deployment Assistant

Overview of Oracle Exadata Deployment Assistant	5-1
Considerations and Requirement for Using OEDA	5-2
Getting Started with the OEDA Browser-based User Interface	5-3
Using the Browser-based Version of Oracle Exadata Deployment Assistant	5-4

Part II Installation and Configuration

6 Unpacking and Installing ZDLRA Rack

7 Expanding a Recovery Appliance Rack with Additional Storage Servers

Preparing to Extend Recovery Appliance	7-2
Estimating RDMA Network Fabric Cable Path Lengths	7-2
Obtaining the Current Configuration Information	7-2
Generating the Network Configuration Files	7-3
Moving the Audit and Diagnostic Files	7-4
Synchronizing Release and Patch Levels	7-4
Precautions and Guidelines	7-4
General Safety Precautions	7-4
InfiniBand Cable Precautions	7-5
InfiniBand Cable Bundling Guidelines	7-5
Cable Management Arm Guidelines	7-5
Adding Storage Servers to a Rack	7-6
Checking the Recovery Appliance Rack	7-6
Unpacking and Labeling the New Storage Servers	7-7
Recovery Appliance RA23, RA21, X8M, and X8 through X5 Population Rules	7-8

Removing the Doors	7-11
Adding New Servers	7-13
Cabling the Storage Servers	7-18
Closing the Rack	7-20
Making New Storage Available	7-21

8 Configuring a Recovery Appliance Rack

Supporting Auto Service Request	8-1
Prerequisites for Using Auto Service Request	8-1
Checking an Existing ASR Manager Installation	8-1
Installing ASR Manager	8-2
Installing the Tape Hardware	8-2
Verifying the Network Configuration Prior to Configuring the Rack	8-2
Configuring the RDMA Network Fabric Switch	8-3
Configuring the InfiniBand Switches	8-3
Configuring an InfiniBand Switch	8-3
Setting the Time on an InfiniBand Switch	8-5
Setting the Serial Number on a Spine Switch	8-6
Checking the Health of an InfiniBand Switch	8-6
Setting a Spine Switch as the Subnet Manager Master	8-8
Configuring the Cisco Nexus 9336C-FX2 Switch	8-9
Configuring the Cisco Ethernet Switch	8-14
Configuring the Cisco Catalyst 4948 Ethernet Switch	8-14
Configuring the Cisco Nexus 93108-1G or 9348 Ethernet Switch	8-21
Performing the Initial Switch Configuration for the Cisco Nexus 93108-1G or 9348 Ethernet Switch	8-21
Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch	8-25
Disabling Spanning Tree on the Ethernet Switch	8-25
Configuring the Power Distribution Units	8-26
Assigning Network Addresses to the PDUs	8-26
Configuring the PDU System Time Settings	8-27
Upgrading the PDU Firmware	8-27
Checking the Health of the Compute Servers	8-28
Checking the Health of the Storage Servers	8-31
Verifying the RoCE Network Fabric Configuration	8-33
Verifying the InfiniBand Network Fabric Network	8-33
Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack	8-35

9 Installing the Recovery Appliance Software

About the Installation Software	9-1

Software Installation Prerequisites	9-2
About Adding a Rack to an Existing Rack	9-2
About Adding New Storage Servers to an Existing Rack	9-2
Reclaiming Disk Space on the Compute Servers	9-2
Installing the Software on Recovery Appliance	9-3
Install Utility Syntax	9-6
Install Utility Parameters	9-6
Install Utility Syntax Examples	9-6
About the Install Utility Steps	9-8
Setting Up Enterprise Manager to Administer Recovery Appliance	9-10
Deploying the Enterprise Manager Agents	9-10
Discovering the Cluster	9-11
Discovering the Cluster Database Targets	9-12
Discovering the Recovery Appliance	9-12
Monitoring TLS-enabled Targets	9-14
Monitoring TLS-enabled Targets with EM Agents	9-15
Discovering TLS-Enabled Recovery Appliance	9-16

Part III Security and Maintenance of Recovery Appliance

10 User Security on Recovery Appliance

Default User Accounts for Oracle Zero Data Loss Recovery Appliance	10-1
User Roles for the Recovery Appliance	10-3
RACLI Non-Root User	10-4
Securing the Operations of the Recovery Appliance	10-5
Kerberos Authentication and the Recovery Appliance	10-6
LDAP Authentication and the Recovery Appliance	10-7
Quorum	10-9
Default Password Requirements	10-10
Default Security Settings Implemented by OEDA	10-11

11 Keeping the Recovery Appliance Secure

Securing the Hardware	11-1
Getting the Rack Serial Number	11-1
Getting the Serial Numbers for Rack Components	11-2
Getting the Rack Serial Number for a Cisco 9336C or 9348 Switch	11-4
Getting the Rack Serial Number for a Sun Datacenter InfiniBand Switch 36	11-5
Getting the Serial Number for a Cisco 4948 Ethernet Switch	11-5
Securing the Software	11-6
Maintaining a Secure Environment	11-6

12 Recovering the Recovery Appliance

13 Maintaining the Recovery Appliance Hardware

Cautions and Warnings	13-1
Determining the Server Model	13-2
Powering On and Off a Recovery Appliance Rack	13-2
Emergency Power-Off Procedure	13-2
Powering Off in an Emergency	13-2
About the Emergency Power-Off Switch	13-2
Shutting Down Recovery Appliance	13-3
Stopping Recovery Appliance Services	13-3
Powering Down the Servers	13-5
Powering the Network Switches	13-6
Starting Up Recovery Appliance	13-6
Starting Up Recovery Appliance Components	13-6
Startup Sequence	13-6
Powering On Servers Remotely	13-6
Starting the Recovery Appliance Software	13-7
Replacing the Disk Controller Batteries	13-8
Monitoring Batteries in the Compute Servers	13-9
Replacing Batteries in Disk Controllers	13-9
Replacing a Power Distribution Unit	13-10
PDU Replacement Guidelines	13-10
Replacing a PDU	13-10
Replacing a Tape Drive	13-12
Save Configuration before Updating Firmware	13-13
Resetting a Non-Responsive Oracle ILOM	13-15
Resetting Oracle ILOM Using SSH	13-15
Resetting Oracle ILOM Using the Remote Console	13-16
Resetting Oracle ILOM Using IPMITool	13-16
Resetting Oracle ILOM By Removing Power	13-16
Maintaining the Compute Servers	13-16
Verifying the RAID Status of a Compute Server	13-17
Reimaging a Compute Server	13-18
Contacting Oracle Support Services	13-18
Downloading the Latest Release of the Cluster Verification Utility	13-18
Removing the Failed Compute Server from the Cluster	13-18

Preparing the USB Flash Drive for Imaging	13-20
Copying the Image to the New Compute Server	13-20
Configuring the Replacement Compute Server	13-21
Preparing the Replacement Compute Server for the Cluster	13-22
Applying Patch Bundles to a Replacement Compute Server	13-24
Cloning the Oracle Grid Infrastructure	13-24
Clone Oracle Database Homes to the Replacement Compute Server	13-26
Maintaining the Storage Servers	13-27
Shutting Down a Storage Server	13-27
Enabling Network Connectivity Using the Diagnostics ISO	13-29
Maintaining the Physical Disks of Storage Servers	13-30
About System Disks and Data Disks	13-30
Monitoring the Status of Physical Disks	13-30
What Happens When Disk Errors Occur?	13-31
About Detecting Underperforming Disks	13-32
About Rebalancing the Data	13-33
Monitoring Hard Disk Controller Write-Through Cache Mode	13-33
Replacing a Failed Physical Disk	13-34
Replacing a Faulty Physical Disk	13-35
Removing an Underperforming Physical Disk	13-37
Moving All Drives from One Storage Server to Another	13-37
Removing and Replacing the Same Physical Disk	13-39
Reenabling a Rejected Physical Disk	13-39
Maintaining the Flash Disks of Storage Servers	13-40
About the Flash Disks	13-40
Faulty Status Indicators	13-41
Identifying Flash Disks in Poor Health	13-42
Identifying Underperforming Flash Disks	13-42
When Is It Safe to Replace a Faulty Flash Disk?	13-43
Replacing a Failed Flash Disk	13-43
Replacing a Faulty Flash Disk	13-44
Removing an Underperforming Flash Disk	13-45
About Write-Back Flash Cache	13-46
Replacing a Disk Controller Battery Backup Unit	13-46
Replacing a Disk Controller BBU on a Compute Server	13-47
Replacing a Disk Controller BBU on a Storage Server	13-47
Using the Storage Server Rescue Procedure	13-48
First Steps Before Rescuing the Storage Server	13-48
If the Server Has Normal Redundancy	13-48
If the Server Has High Redundancy	13-49
About the Rescue Procedure	13-49
Rescuing a Server Using the CELLBOOT USB Flash Drive	13-50

Reconfiguring the Rescued Storage Server	13-51
Recreating a Damaged CELLBOOT USB Flash Drive	13-51

14 Reconfiguring the Hardware

Changing Component Passwords	14-1
Changing the Compute Server Passwords	14-1
Changing the User Account Password	14-1
Changing the GRUB Account Password	14-2
Changing the Storage Server Passwords	14-2
Creating New Users and Changing Passwords on a PDU	14-2
Changing the Oracle ILOM Password	14-3
Changing an InfiniBand Switch Password	14-3
Changing the Cisco Ethernet Switch Password	14-4
Changing the Recovery Appliance Database User Password	14-4
Changing OSB Tape Backup Passwords	14-5
Changing the DNS Servers	14-5
Changing the DNS Server Address on an InfiniBand Switch	14-5
Changing the DNS Server Address on the Cisco Ethernet Switch	14-6
Changing the DNS Server Address on the Compute Server	14-7
Changing the DNS Server on Storage Servers	14-7
Changing the chrony/NTP Servers	14-8
Setting the NTP Server Address on the Compute Servers	14-8
Setting the CHRONY Server Address on the InfiniBand Switches	14-8
Set the NTP Server Address on the Cisco Ethernet Switch	14-9
Setting the CHRONY Server on the Storage Servers	14-10
Restart Recovery Appliance	14-10
Changing the Time Zone Settings	14-10
Changing the Time Zone Settings on the Storage Servers	14-11
Change Time Zone Settings on the Compute Servers	14-12
Change the Time Zone Settings on the InfiniBand Switches	14-13
Change the Time Zone Settings on the Cisco Switch	14-13

Part IV Recovery Appliance Command Line Interface Reference

15 RACLI Overview

racli add	15-7
racli add admin_user	15-7
racli add certificate	15-7
racli add cloud_key	15-8
racli add cloud_location	15-9

racli add cloud_user	15-10
racli add credential	15-11
racli add credential_wallet	15-12
racli add db_user	15-12
racli add host	15-13
racli add keystore	15-13
racli add nfs	15-14
racli add rasys_wallet	15-14
racli add remote_syslog	15-15
racli add storage_cells	15-15
racli add storage_disks	15-16
racli add tag	15-16
racli add vpc_user	15-17
racli alter	15-18
racli alter admin_user	15-18
racli alter adr_parameter	15-19
racli alter cloud_key	15-19
racli alter cloud_location	15-20
racli alter db_user	15-20
racli alter keystore	15-21
racli alter network	15-22
racli alter parameter	15-23
racli alter passphrase	15-23
racli alter rasys_user	15-24
racli alter vpc_user	15-25
racli approve	15-25
racli approve request	15-25
racli configure	15-26
racli configure export	15-26
racli create	15-26
racli create certificate	15-26
racli deny	15-27
racli deny request	15-27
racli disable	15-27
racli disable aide	15-28
racli disable archived_log_grouping	15-28
racli disable check	15-28
racli disable cloud_location	15-29
racli disable keystore	15-29
racli disable nfs	15-29
racli disable rasys_user	15-30
racli disable root_access	15-30

racli disable rpc	15-30
racli disable ssh	15-31
racli disable sys_remote_access	15-31
racli enable	15-31
racli enable acfs	15-32
racli enable aide	15-32
racli enable archived_log_grouping	15-32
racli enable check	15-33
racli enable cloud_location	15-33
racli enable keystore	15-33
racli enable nfs	15-34
racli enable rasys_user	15-34
racli enable rpc	15-34
racli enable ssh	15-35
racli enable sys_remote_access	15-35
racli export	15-35
racli export appliance	15-35
racli install	15-36
racli install appliance	15-36
racli install okv_endpoint	15-37
racli list	15-38
racli list admin_user	15-38
racli list certificate	15-38
racli list check	15-38
racli list cloud_key	15-39
racli list cloud_location	15-39
racli list cloud_user	15-40
racli list compute_nodes	15-40
racli list db_user	15-41
racli list env	15-41
racli list export	15-41
racli list grid_home	15-42
racli list host	15-42
racli list ip_addr	15-42
racli list listener	15-43
racli list nfs	15-43
racli list okv_endpoint	15-43
racli list oracle_home	15-44
racli list parameter	15-44
racli list remote_syslog	15-45
racli list san	15-45
racli list scan	15-45

racli list scan_listener	15-45
racli list storage_cells	15-46
racli list tag	15-46
racli remove	15-46
racli remove acfs	15-47
racli remove admin_user	15-47
racli remove certificate	15-47
racli remove cloud_key	15-48
racli remove cloud_location	15-48
racli remove cloud_user	15-48
racli remove db_user	15-49
racli remove grid_home	15-49
racli remove host	15-49
racli remove keystore	15-50
racli remove nfs	15-50
racli remove oracle_home	15-50
racli remove osb_pieces	15-51
racli remove remote_syslog	15-51
racli remove tag	15-51
racli remove vpc_user	15-51
racli run	15-52
racli run check	15-52
racli run diagnostics	15-53
racli run preinstall	15-54
racli start	15-55
racli start appliance	15-55
racli start nfs	15-55
racli status	15-56
racli status aide	15-56
racli status appliance	15-56
racli status archived_log_grouping	15-56
racli status cloud_location	15-57
racli status compliance	15-57
racli status disk_group	15-57
racli status keystore	15-58
racli status listener	15-58
racli status nfs	15-58
racli status okv_endpoint	15-59
racli status request	15-59
racli status scan	15-59
racli status scan_listener	15-59
racli stop	15-60

racli stop appliance	15-60
racli stop nfs	15-60
racli uninstall	15-60
racli uninstall appliance	15-61
racli uninstall okv_endpoint	15-62
racli update	15-62
racli update aide	15-63
racli update appliance	15-63
racli update parameter	15-64
racli update ssh_keys	15-64
racli update ssh_known_hosts	15-64
racli update storage_location	15-64
racli update support_tools	15-65
racli version	15-65

Part V Cabling Information for Oracle Zero Data Loss Recovery Appliance

16 In-Rack Cabling Tables for Recovery Appliance RA23

Cable Color Coding	16-1
Recovery Appliance RA23 Rack Layout	16-1
Oracle Integrated Lights Out Manager Cabling Connections RA23	16-2
Administrative Ethernet Port Cabling Tables RA23	16-3
Single-Phase Power Distribution Unit Cabling Connections RA23	16-5
Three-Phase Power Distribution Unit Cabling Connections RA23	16-6
RDMA Network Fabric Switch-to-Switch Cabling Connections RA23	16-7
RDMA Network Fabric Switch-to-Server Cabling Connections RA23	16-8

17 In-Rack Cabling Tables for Recovery Appliance RA21

Cable Color Coding	17-1
Recovery Appliance RA21 Rack Layout	17-1
Oracle Integrated Lights Out Manager Cabling Connections RA21	17-6
Administrative Ethernet Port Cabling Tables RA21	17-6
Single-Phase Power Distribution Unit Cabling Connections RA21	17-8
Three-Phase Power Distribution Unit Cabling Connections RA21	17-9
RDMA Network Fabric Switch-to-Switch Cabling Connections RA21	17-10
RDMA Network Fabric Switch-to-Server Cabling Connections RA21	17-11

18 In-Rack Cabling Tables for Recovery Appliance X8M

Cable Color Coding	18-1
Recovery Appliance RA21 and X8M Rack Layout	18-1
Oracle Integrated Lights Out Manager Cabling Connections X8M	18-3
Administrative Ethernet Port Cabling Tables X8M	18-4
Single-Phase Power Distribution Unit Cabling Connections X8M	18-5
Three-Phase Power Distribution Unit Cabling Connections X8M	18-6
RDMA Network Fabric Switch-to-Switch Cabling Connections X8M	18-7
RDMA Network Fabric Switch-to-Server Cabling Connections X8M	18-8

19 In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5

Cable Color Coding	19-1
Recovery Appliance X8, X7, X6, and X5 Rack Layout	19-1
Oracle Integrated Lights Out Manager Cabling Connections	19-3
Administrative Gigabit Ethernet Cabling Connections	19-4
Single-Phase Power Distribution Unit Cabling Connections X8	19-5
Single-Phase Power Distribution Unit Cabling Connections X7 and Earlier	19-6
Three-Phase Power Distribution Unit Cabling Connections X8	19-7
Three-Phase Power Distribution Unit Cabling Connections	19-8
RDMA Network Fabric Switch-to-Switch Cabling Connections	19-9
RDMA Network Fabric Switch-to-Server Cabling Connections	19-10

20 In-Rack Cabling Tables for Recovery Appliance X4

Recovery Appliance X4 Rack Layout	20-1
Oracle Integrated Lights Out Manager Cabling Connections	20-3
Administrative Gigabit Ethernet Cabling Connections	20-4
Single-Phase Power Distribution Unit Cabling Connections	20-5
Three-Phase Power Distribution Unit Cabling Connections	20-6
InfiniBand Switch-to-Switch Cabling Connections	20-7
InfiniBand Switch-to-Server Cabling Connections	20-7

Part VI Extending Oracle Zero Data Loss Recovery Appliance

21 Extending a Rack by Adding Another Rack

Overview of Adding Another Rack to an Existing System	21-1
Cabling Two Racks Together	21-2
Cabling Two RoCE Network Fabric Racks Together with No Down Time	21-2

Extending an RA21 or Later Model Rack with No Down Time by Adding Another RA21 or Later Model Rack	21-2
Extending an X8M Rack with No Down Time by Adding an RA21 or Later Model Rack	21-41
Extending an X8M Rack with No Down Time by Adding Another X8M Rack	21-83
Cabling Two RoCE Network Fabric Racks Together with Down Time Allowed	21-121
Cabling Two InfiniBand Network Fabric Racks Together	21-127
Cabling Several Racks Together	21-128
Cabling Several RoCE Network Fabric Racks Together	21-128
Cabling Several InfiniBand Network Fabric Racks Together	21-137

22 Multi-Rack Cabling Tables for Recovery Appliance RA21 and Later Models

Understanding Multi-Rack Cabling for RA21 and Later Model Racks	22-1
Preparing for Multi-Rack Cabling with RA21 and Later Model Racks	22-3
Two-Rack Cabling for RA21 and Later Model Racks	22-5
Three-Rack Cabling for RA21 and Later Model Racks	22-9
Four-Rack Cabling for RA21 and Later Model Racks	22-15
Five-Rack Cabling for RA21 and Later Model Racks	22-22
Six-Rack Cabling for RA21 and Later Model Racks	22-30
Seven-Rack Cabling for RA21 and Later Model Racks	22-39
Eight-Rack Cabling for RA21 and Later Model Racks	22-51
Nine-Rack Cabling for RA21 and Later Model Racks	22-63
Ten-Rack Cabling for RA21 and Later Model Racks	22-77
Eleven-Rack Cabling for RA21 and Later Model Racks	22-92
Twelve-Rack Cabling for RA21 and Later Model Racks	22-109
Thirteen-Rack Cabling for RA21 and Later Model Racks	22-127
Fourteen-Rack Cabling for RA21 and Later Model Racks	22-147

23 Multi-Rack Cabling Tables for Recovery Appliance X8M Models

Understanding Multi-Rack Cabling for X8M Racks	23-1
Preparing for Multi-Rack Cabling with X8M Racks	23-4
Cabling Multiple Recovery Appliance X8M Racks	23-6
Two-Rack Cabling for X8M Racks	23-7
Three-Rack Cabling for X8M Racks	23-9
Four-Rack Cabling for X8M Racks	23-13
Five-Rack Cabling for X8M Racks	23-16
Six-Rack Cabling for X8M Racks	23-21
Seven-Rack Cabling for X8M Racks	23-27
Eight-Rack Cabling for X8M Racks	23-34

24 Multi-Rack Cabling Tables for ZDLRA Rack Models with InfiniBand Network Fabric (X2 to X8)

Understanding Multi-Rack Cabling for Racks with InfiniBand Network Fabric	24-1
Preparing for Multi-Rack Cabling with InfiniBand Network Fabric	24-4
Cabling Oracle Exadata Quarter Racks and Oracle Exadata Eighth Racks with InfiniBand Network Fabric	24-6
Two-Rack Cabling with InfiniBand Network Fabric	24-8
Three-Rack Cabling with InfiniBand Network Fabric	24-10
Four-Rack Cabling with InfiniBand Network Fabric	24-12
Five-Rack Cabling with InfiniBand Network Fabric	24-15
Six-Rack Cabling with InfiniBand Network Fabric	24-18
Seven-Rack Cabling with InfiniBand Network Fabric	24-21
Eight-Rack Cabling with InfiniBand Network Fabric	24-25

Part VII Appendixes

25 Error Messages

Contacting Oracle Support Services	25-1
Generic Alert Messages: HALRT-00001 to HALRT-00500	25-2
Temperature Alert Messages: HALRT-00513 to HALRT-00600	25-5
Fan Alert Messages: HALRT-00616 to HALRT-00700	25-6
Power Supply/Converters Alert Messages: HALRT-00717 to HALRT-00800	25-6
Network Alert Messages: HALRT-00912 to HALRT-01000	25-7
Hard Disk and Flash Disk Alert Messages: HALRT-01015 to HALRT-01200	25-7
Accelerator/Controller Alert Messages: HALRT-01209 to HALRT-01300	25-27
Voltage Alert Messages: HALRT-01301 to HALRT-01400	25-30
Physical Disk Alert Messages: HALRT-01401 to HALRT-01500	25-30
Auto Service Request Alert Messages for Storage Servers: HALRT-02001 to HALRT-02013	25-33
HDD Controller Battery Alert Messages: HALRT-02101 to HALRT-02200	25-34
Generic Actions for Alert Messages: HALRT-10001 to HALRT-10500	25-37
Generic Actions for Alert Messages: HALRT-10503 to HALRT-11000	25-42
Generic Actions for Alert Messages: HALRT-11001 to HALRT-11100	25-43

26 Replacement Units

About Warranty Coverage	26-1
Understanding Repair Categories	26-1
LED Status Descriptions	26-2
Oracle Server LEDs	26-2
Sun Datacenter InfiniBand Switch 36 LEDs	26-4

Parts for Compute Servers	26-6
Database Server Components of Recovery Appliance X9M	26-6
Database Server Components of Recovery Appliance X8M	26-7
Oracle Server X8 and X8M Database Server Parts	26-8
Parts for Storage Servers	26-11
Oracle Exadata Storage Server X9M, X8M, and X8 High Capacity and Extended (XT) Server Parts	26-11
Parts for Cisco Switch	26-15
Parts for InfiniBand Switches	26-15
Parts for Power Distribution Units	26-16
Cables and Cable Accessories for Recovery Appliance	26-17

Index

List of Figures

2-1	Typical Data Center Configuration for Perforated Floor Tiles	2-15
2-2	ZDLRA RA23 100GE to Exadata X10	2-19
3-1	Oracle Zero Data Loss Recovery Appliance RA23 Networking	3-5
3-2	Oracle Zero Data Loss Recovery Appliance RA21 Backplane External Network Connectivity	3-5
3-3	Oracle Zero Data Loss Recovery Appliance X8M Backplane External Network Connectivity	3-6
3-4	Oracle Zero Data Loss Recovery Appliance X8 and X7 Backplane External Network Connectivity	3-7
3-5	Network Diagram for Recovery Appliance	3-9
3-6	Network Connections to the Recovery Appliance Rack Components	3-10
3-7	Recovery Appliance Connection to an Oracle Tape System	3-11
3-8	Recovery Appliance Connection to a Third-Party Tape System	3-12
4-1	Auto Service Request Network Connections	4-3
7-1	Server Expansion Kit Components	7-8
7-2	RA23 Rack Layout and Population Rules	7-10
7-3	Removing the Rack Doors	7-12
7-4	Removing the Filler Panels	7-13
7-5	Location of the Trough	7-14
7-6	Aligning the Side-Rail Assembly With the Rack	7-15
7-7	Inserting the Server with Mounting Brackets into the Slide Rails	7-17
16-1	RA23 Rack Population	16-2
17-1	RA21 Rack Population	17-2
17-2	Rack Layout of the Recovery Appliance RA21 Full Rack	17-3
17-3	Rack Layout of the Recovery Appliance RA21 Minimum Configuration	17-5
18-1	Rack Layout of the Recovery Appliance RA21 and X8M Full Rack	18-2
18-2	Rack Layout of the Recovery Appliance RA21 and X8M Minimum Configuration	18-3
19-1	Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Full Rack	19-2
19-2	Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Minimum Configuration	19-3
20-1	Rack Layout of the Recovery Appliance X4 Full Rack	20-2
20-2	Rack Layout of the Recovery Appliance X4 Minimum Configuration	20-3
23-1	Connections Between Spine Switches and Leaf Switches for up to 8 Racks	23-4
24-1	Connections Between Spine Switch and Leaf Switches in a Single Rack	24-2
24-2	Connections Between Database Servers and Storage Servers and Leaf Switches	24-2
24-3	Connections Between Spine Switches and Leaf Switches Across Two Racks	24-3
24-4	Connections Between Spine Switches and Leaf Switches for up to 8 Racks	24-3
24-5	Leaf Switch Connections Between 2 Quarter Racks	24-7
24-6	Leaf and Spine Switch Connections Between a Quarter Rack and a Half or Full Rack	24-8

24-7	Leaf and Spine Switch Connections for a Quarter Rack Connected to One or More Half or Full Racks	24-8
26-1	Front View of Oracle Server X8 and X8M Database Servers	26-8
26-2	Rear View of Oracle Server X8 and X8M Database Servers	26-8
26-3	Front View of Oracle Exadata Storage Server X9M, X8M and X8 High Capacity and XT servers	26-11
26-4	Rear View of Oracle Exadata Storage Server X9M, X8M and X8 High Capacity and XT servers	26-12

List of Tables

1-1	Components of a Oracle Zero Data Loss Recovery Appliance RA23 Rack	1-5
1-2	Components of a Oracle Zero Data Loss Recovery Appliance RA21 Rack	1-5
1-3	Components of a Oracle Zero Data Loss Recovery Appliance X8M Rack	1-6
1-4	Components of a Recovery Appliance X8 Rack	1-6
1-5	Components of a Recovery Appliance X7 Rack	1-7
1-6	Components of a Recovery Appliance X6 Rack	1-7
1-7	Components of a Recovery Appliance X5 Rack	1-7
1-8	Components of a Recovery Appliance X4 Rack	1-8
1-9	RA21 Compute Server Components	1-9
1-10	X9M Compute Server Components	1-10
1-11	Oracle Zero Data Loss Recovery Appliance X8M Compute Server Components	1-11
1-12	Oracle X8 Compute Server Components	1-12
1-13	Oracle X7-2 Compute Server Components	1-12
1-14	Oracle Server X6-2 Components	1-13
1-15	Oracle Server X5-2 Components	1-13
1-16	Sun Server X4-2 Components	1-14
1-17	Oracle Zero Data Loss Recovery Appliance X8 and X8M Server Storage Components	1-15
1-18	Oracle Server X7–2L Storage Components	1-16
1-19	Oracle X6-2L Storage Server Components	1-16
1-20	Oracle Server X5-2L Components	1-16
1-21	Sun Server X4-2L Components	1-17
1-22	Oracle Tape System Components	1-18
2-1	Access Route Requirements	2-4
2-2	Shipping Weight	2-4
2-3	Maintenance Access Requirements for ZDLRA Rack (All Models)	2-5
2-4	Floor Load Requirements for ZDLRA Rack	2-5
2-5	Specifications for each Low-Voltage 15 kVA Single Phase PDU for North America, South America, Japan and Taiwan	2-6
2-6	Specifications for each Low-Voltage 22 kVA Single Phase PDU for North America, South America, Japan and Taiwan	2-7
2-7	Specifications for each Low-Voltage 15 kVA Three Phase PDU for North America, South America, Japan and Taiwan	2-8
2-8	Specifications for each Low-Voltage 24 kVA Three Phase PDU for North America, South America, Japan and Taiwan	2-8
2-9	Specifications for each High-Voltage 15 kVA Single Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	2-9

2-10	Specifications for each High-Voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	2-10
2-11	Specifications for each High-Voltage 15 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	2-10
2-12	Specifications for each High-Voltage 24 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan	2-11
2-13	Temperature, Humidity and Altitude Requirements	2-13
2-14	Ethernet Cabling for the Recovery Appliance RA21 Racks	2-17
3-1	Factory IP Addresses for Recovery Appliance	3-13
3-2	Open Ports for the Firewall	3-15
7-1	Recovery Appliance RA23 Rack Population Rules	7-9
7-2	Recovery Appliance RA21, X8M, and X8 through X5 Rack Population Rules	7-11
9-1	Options of the Install Utility	9-6
10-1	Default Users and Passwords for Oracle Zero Data Loss Recovery Appliance	10-1
13-1	Battery Life Expectancy	13-9
13-2	Alerts Indicating Poor Disk Performance	13-32
16-1	ILOM Cabling for the Recovery Appliance RA23 Full Racks	16-3
16-2	Ethernet Cabling for the Recovery Appliance RA23 Racks	16-3
16-3	Single-Phase PDU Cabling for the Recovery Appliance RA23 Full Racks	16-5
16-4	Three-Phase PDU Cabling for the Recovery Appliance RA23 Full Racks	16-6
16-5	RDMA Network Fabric Switch-to-Switch Cabling RA23	16-7
16-6	RDMA Network Fabric Switch to Compute Server Cabling for the Recovery Appliance RA23 Full Racks	16-8
16-7	RDMA Network Fabric Switch to HC Server Cabling for the Recovery Appliance RA23 Full Racks	16-8
17-1	ILOM Cabling for the Recovery Appliance RA21 Full Racks	17-6
17-2	Ethernet Cabling for the Recovery Appliance RA21 Racks	17-7
17-3	Single-Phase PDU Cabling for the Recovery Appliance RA21 Full Racks	17-8
17-4	Three-Phase PDU Cabling for the Recovery Appliance RA21 Full Racks	17-9
17-5	RDMA Network Fabric Switch-to-Switch Cabling RA21	17-10
17-6	RDMA Network Fabric Switch to Compute Server Cabling for the Recovery Appliance RA21 Full Racks	17-11
17-7	RDMA Network Fabric Switch to HC Server Cabling for the Recovery Appliance RA21 Full Racks	17-11
18-1	ILOM Cabling for the Recovery Appliance X8M Full Racks	18-4
18-2	Ethernet Cabling for the Recovery Appliance X8M Full Racks	18-5
18-3	Single-Phase PDU Cabling for the Recovery Appliance X8M Full Racks	18-5
18-4	Three-Phase PDU Cabling for the Recovery Appliance X8M Full Racks	18-6
18-5	RDMA Network Fabric Switch-to-Switch Cabling X8M	18-7

18-6	RDMA Network Fabric Switch-to-Server Cabling for the Recovery Appliance X8M Full Racks	18-8
19-1	ILOM Cabling for the Recovery Appliance X7, X6, or X5 Full Racks	19-4
19-2	Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks	19-4
19-3	Single-Phase PDU Cabling for the Recovery Appliance X8 Full Racks	19-5
19-4	Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks	19-6
19-5	Three-Phase PDU Cabling for the Recovery Appliance X8 Full Racks	19-7
19-6	Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks	19-8
19-7	RDMA Network Fabric Switch-to-Switch Cabling	19-9
19-8	RDMA Network Fabric Switch-to-Server Cabling for the Recovery Appliance X8, X7, X6 or X5 Full Racks	19-10
20-1	ILOM Cabling for the Recovery Appliance X4 Full Racks	20-4
20-2	Gigabit Ethernet Cabling for the Recovery Appliance X4 Full Racks	20-4
20-3	Single-Phase PDU Cabling for the Recovery Appliance X4 Full Racks	20-5
20-4	Three-Phase PDU Cabling for the Recovery Appliance X4 Full Racks	20-6
20-5	InfiniBand Switch-to-Switch Cabling	20-7
20-6	InfiniBand Switch-to-Server Cabling for the Recovery Appliance X4 Full Racks	20-7
21-1	Leaf Switch Connections for the X8M Rack (R1)	21-82
21-2	Leaf Switch Connections for the RA21 or Later Model Rack (R2)	21-82
22-1	Leaf Switch Connections for the First Rack in a Two-Rack System	22-6
22-2	Leaf Switch Connections for the Second Rack in a Two-Rack System	22-7
22-3	Two-Rack System Connections Sorted By The Spine Switch Port Location	22-8
22-4	Leaf Switch Connections for the First Rack in a Three-Rack System	22-10
22-5	Leaf Switch Connections for the Second Rack in a Three-Rack System	22-11
22-6	Leaf Switch Connections for the Third Rack in a Three-Rack System	22-12
22-7	Three-Rack System Connections Sorted By The Spine Switch Port Location	22-13
22-8	Leaf Switch Connections for the First Rack in a Four-Rack System	22-16
22-9	Leaf Switch Connections for the Second Rack in a Four-Rack System	22-16
22-10	Leaf Switch Connections for the Third Rack in a Four-Rack System	22-17
22-11	Leaf Switch Connections for the Fourth Rack in a Four-Rack System	22-18
22-12	Four-Rack System Connections Sorted By The Spine Switch Port Location	22-19
22-13	Leaf Switch Connections for the First Rack in a Five-Rack System	22-22
22-14	Leaf Switch Connections for the Second Rack in a Five-Rack System	22-23
22-15	Leaf Switch Connections for the Third Rack in a Five-Rack System	22-24
22-16	Leaf Switch Connections for the Fourth Rack in a Five-Rack System	22-25
22-17	Leaf Switch Connections for the Fifth Rack in a Five-Rack System	22-26
22-18	Five-Rack System Connections Sorted By The Spine Switch Port Location	22-26
22-19	Leaf Switch Connections for the First Rack in a Six-Rack System	22-31

22-20	Leaf Switch Connections for the Second Rack in a Six-Rack System	22-31
22-21	Leaf Switch Connections for the Third Rack in a Six-Rack System	22-32
22-22	Leaf Switch Connections for the Fourth Rack in a Six-Rack System	22-33
22-23	Leaf Switch Connections for the Fifth Rack in a Six-Rack System	22-34
22-24	Leaf Switch Connections for the Sixth Rack in a Six-Rack System	22-35
22-25	Six-Rack System Connections Sorted By The Spine Switch Port Location	22-35
22-26	Leaf Switch Connections for the First Rack in a Seven-Rack System	22-40
22-27	Leaf Switch Connections for the Second Rack in a Seven-Rack System	22-41
22-28	Leaf Switch Connections for the Third Rack in a Seven-Rack System	22-42
22-29	Leaf Switch Connections for the Fourth Rack in a Seven-Rack System	22-43
22-30	Leaf Switch Connections for the Fifth Rack in a Seven-Rack System	22-43
22-31	Leaf Switch Connections for the Sixth Rack in a Seven-Rack System	22-44
22-32	Leaf Switch Connections for the Seventh Rack in a Seven-Rack System	22-45
22-33	Seven-Rack System Connections Sorted By The Spine Switch Port Location	22-46
22-34	Leaf Switch Connections for the First Rack in a Eight-Rack System	22-52
22-35	Leaf Switch Connections for the Second Rack in a Eight-Rack System	22-52
22-36	Leaf Switch Connections for the Third Rack in a Eight-Rack System	22-53
22-37	Leaf Switch Connections for the Fourth Rack in a Eight-Rack System	22-54
22-38	Leaf Switch Connections for the Fifth Rack in a Eight-Rack System	22-55
22-39	Leaf Switch Connections for the Sixth Rack in a Eight-Rack System	22-55
22-40	Leaf Switch Connections for the Seventh Rack in a Eight-Rack System	22-56
22-41	Leaf Switch Connections for the Eighth Rack in a Eight-Rack System	22-57
22-42	Eight-Rack System Connections Sorted By The Spine Switch Port Location	22-58
22-43	Leaf Switch Connections for the First Rack in a 9 Rack System	22-64
22-44	Leaf Switch Connections for the Second Rack in a 9 Rack System	22-65
22-45	Leaf Switch Connections for the Third Rack in a 9 Rack System	22-66
22-46	Leaf Switch Connections for the Fourth Rack in a 9 Rack System	22-66
22-47	Leaf Switch Connections for the Fifth Rack in a 9 Rack System	22-67
22-48	Leaf Switch Connections for the Sixth Rack in a 9 Rack System	22-68
22-49	Leaf Switch Connections for the Seventh Rack in a 9 Rack System	22-69
22-50	Leaf Switch Connections for the Eighth Rack in a 9 Rack System	22-69
22-51	Leaf Switch Connections for the Ninth Rack in a 9 Rack System	22-70
22-52	Nine-Rack System Connections Sorted By The Spine Switch Port Location	22-71
22-53	Leaf Switch Connections for the First Rack in a 10 Rack System	22-78
22-54	Leaf Switch Connections for the Second Rack in a 10 Rack System	22-79
22-55	Leaf Switch Connections for the Third Rack in a 10 Rack System	22-80
22-56	Leaf Switch Connections for the Fourth Rack in a 10 Rack System	22-80

22-57	Leaf Switch Connections for the Fifth Rack in a 10 Rack System	22-81
22-58	Leaf Switch Connections for the Sixth Rack in a 10 Rack System	22-82
22-59	Leaf Switch Connections for the Seventh Rack in a 10 Rack System	22-83
22-60	Leaf Switch Connections for the Eighth Rack in a 10 Rack System	22-83
22-61	Leaf Switch Connections for the Ninth Rack in a 10 Rack System	22-84
22-62	Leaf Switch Connections for the Tenth Rack in a 10 Rack System	22-85
22-63	Ten-Rack System Connections Sorted By The Spine Switch Port Location	22-86
22-64	Leaf Switch Connections for the First Rack in a 11 Rack System	22-93
22-65	Leaf Switch Connections for the Second Rack in a 11 Rack System	22-94
22-66	Leaf Switch Connections for the Third Rack in a 11 Rack System	22-95
22-67	Leaf Switch Connections for the Fourth Rack in a 11 Rack System	22-95
22-68	Leaf Switch Connections for the Fifth Rack in a 11 Rack System	22-96
22-69	Leaf Switch Connections for the Sixth Rack in a 11 Rack System	22-97
22-70	Leaf Switch Connections for the Seventh Rack in a 11 Rack System	22-98
22-71	Leaf Switch Connections for the Eighth Rack in a 11 Rack System	22-98
22-72	Leaf Switch Connections for the Ninth Rack in a 11 Rack System	22-99
22-73	Leaf Switch Connections for the Tenth Rack in a 11 Rack System	22-100
22-74	Leaf Switch Connections for the Eleventh Rack in a 11 Rack System	22-101
22-75	Eleven-Rack System Connections Sorted By The Spine Switch Port Location	22-102
22-76	Leaf Switch Connections for the First Rack in a 12 Rack System	22-110
22-77	Leaf Switch Connections for the Second Rack in a 12 Rack System	22-111
22-78	Leaf Switch Connections for the Third Rack in a 12 Rack System	22-112
22-79	Leaf Switch Connections for the Fourth Rack in a 12 Rack System	22-112
22-80	Leaf Switch Connections for the Fifth Rack in a 12 Rack System	22-113
22-81	Leaf Switch Connections for the Sixth Rack in a 12 Rack System	22-114
22-82	Leaf Switch Connections for the Seventh Rack in a 12 Rack System	22-115
22-83	Leaf Switch Connections for the Eighth Rack in a 12 Rack System	22-115
22-84	Leaf Switch Connections for the Ninth Rack in a 12 Rack System	22-116
22-85	Leaf Switch Connections for the Tenth Rack in a 12 Rack System	22-117
22-86	Leaf Switch Connections for the Eleventh Rack in a 12 Rack System	22-118
22-87	Leaf Switch Connections for the Twelfth Rack in a 12 Rack System	22-118
22-88	Twelve-Rack System Connections Sorted By The Spine Switch Port Location	22-119
22-89	Leaf Switch Connections for the First Rack in a 13 Rack System	22-128
22-90	Leaf Switch Connections for the Second Rack in a 13 Rack System	22-129
22-91	Leaf Switch Connections for the Third Rack in a 13 Rack System	22-130
22-92	Leaf Switch Connections for the Fourth Rack in a 13 Rack System	22-130
22-93	Leaf Switch Connections for the Fifth Rack in a 13 Rack System	22-131

22-94	Leaf Switch Connections for the Sixth Rack in a 13 Rack System	22-132
22-95	Leaf Switch Connections for the Seventh Rack in a 13 Rack System	22-133
22-96	Leaf Switch Connections for the Eighth Rack in a 13 Rack System	22-133
22-97	Leaf Switch Connections for the Ninth Rack in a 13 Rack System	22-134
22-98	Leaf Switch Connections for the Tenth Rack in a 13 Rack System	22-135
22-99	Leaf Switch Connections for the Eleventh Rack in a 13 Rack System	22-136
22-100	Leaf Switch Connections for the Twelfth Rack in a 13 Rack System	22-136
22-101	Leaf Switch Connections for the Thirteenth Rack in a 13 Rack System	22-137
22-102	Thirteen-Rack System Connections Sorted By The Spine Switch Port Location	22-138
22-103	Leaf Switch Connections for the First Rack in a 14 Rack System	22-148
22-104	Leaf Switch Connections for the Second Rack in a 14 Rack System	22-148
22-105	Leaf Switch Connections for the Third Rack in a 14 Rack System	22-149
22-106	Leaf Switch Connections for the Fourth Rack in a 14 Rack System	22-150
22-107	Leaf Switch Connections for the Fifth Rack in a 14 Rack System	22-151
22-108	Leaf Switch Connections for the Sixth Rack in a 14 Rack System	22-151
22-109	Leaf Switch Connections for the Seventh Rack in a 14 Rack System	22-152
22-110	Leaf Switch Connections for the Eighth Rack in a 14 Rack System	22-153
22-111	Leaf Switch Connections for the Ninth Rack in a 14 Rack System	22-154
22-112	Leaf Switch Connections for the Tenth Rack in a 14 Rack System	22-154
22-113	Leaf Switch Connections for the Eleventh Rack in a 14 Rack System	22-155
22-114	Leaf Switch Connections for the Twelfth Rack in a 14 Rack System	22-156
22-115	Leaf Switch Connections for the Thirteenth Rack in a 14 Rack System	22-157
22-116	Leaf Switch Connections for the Fourteenth Rack in a 14 Rack System	22-157
22-117	Fourteen-Rack System Connections Sorted By The Spine Switch Port Location	22-158
23-1	Leaf Switch Connections for the First Rack in a Two-Rack System	23-8
23-2	Leaf Switch Connections for the Second Rack in a Two-Rack System	23-9
23-3	Leaf Switch Connections for the First Rack in a Three-Rack System	23-11
23-4	Leaf Switch Connections for the Second Rack in a Three-Rack System	23-12
23-5	Leaf Switch Connections for the Third Rack in a Three-Rack System	23-12
23-6	Leaf Switch Connections for the First Rack in a Four-Rack System	23-14
23-7	Leaf Switch Connections for the Second Rack in a Four-Rack System	23-15
23-8	Leaf Switch Connections for the Third Rack in a Four-Rack System	23-15
23-9	Leaf Switch Connections for the Fourth Rack in a Four-Rack System	23-16
23-10	Leaf Switch Connections for the First Rack in a Five-Rack System	23-19
23-11	Leaf Switch Connections for the Second Rack in a Five-Rack System	23-19
23-12	Leaf Switch Connections for the Third Rack in a Five-Rack System	23-20
23-13	Leaf Switch Connections for the Fourth Rack in a Five-Rack System	23-20

23-14	Leaf Switch Connections for the Fifth Rack in a Five-Rack System	23-21
23-15	Leaf Switch Connections for the First Rack in a Six-Rack System	23-24
23-16	Leaf Switch Connections for the Second Rack in a Six-Rack System	23-24
23-17	Leaf Switch Connections for the Third Rack in a Six-Rack System	23-25
23-18	Leaf Switch Connections for the Fourth Rack in a Six-Rack System	23-25
23-19	Leaf Switch Connections for the Fifth Rack in a Six-Rack System	23-26
23-20	Leaf Switch Connections for the Sixth Rack in a Six-Rack System	23-26
23-21	Leaf Switch Connections for the First Rack in a Seven-Rack System	23-30
23-22	Leaf Switch Connections for the Second Rack in a Seven-Rack System	23-31
23-23	Leaf Switch Connections for the Third Rack in a Seven-Rack System	23-31
23-24	Leaf Switch Connections for the Fourth Rack in a Seven-Rack System	23-32
23-25	Leaf Switch Connections for the Fifth Rack in a Seven-Rack System	23-32
23-26	Leaf Switch Connections for the Sixth Rack in a Seven-Rack System	23-33
23-27	Leaf Switch Connections for the Seventh Rack in a Seven-Rack System	23-34
23-28	Leaf Switch Connections for the First Rack in a Eight-Rack System	23-37
23-29	Leaf Switch Connections for the Second Rack in a Eight-Rack System	23-38
23-30	Leaf Switch Connections for the Third Rack in a Eight-Rack System	23-39
23-31	Leaf Switch Connections for the Fourth Rack in a Eight-Rack System	23-39
23-32	Leaf Switch Connections for the Fifth Rack in a Eight-Rack System	23-40
23-33	Leaf Switch Connections for the Sixth Rack in a Eight-Rack System	23-40
23-34	Leaf Switch Connections for the Seventh Rack in a Eight-Rack System	23-41
23-35	Leaf Switch Connections for the Eighth Rack in a Eight-Rack System	23-41
24-1	Leaf Switch Connections for the First Rack in a Two-Rack System	24-9
24-2	Leaf Switch Connections for the Second Rack in a Two-Rack System	24-10
24-3	Leaf Switch Connections for the First Rack in a Three-Rack System	24-11
24-4	Leaf Switch Connections for the Second Rack in a Three-Rack System	24-11
24-5	Leaf Switch Connections for the Third Rack in a Three-Rack System	24-12
24-6	Leaf Switch Connections for the First Rack in a Four-Rack System	24-13
24-7	Leaf Switch Connections for the Second Rack in a Four-Rack System	24-13
24-8	Leaf Switch Connections for the Third Rack in a Four-Rack System	24-14
24-9	Leaf Switch Connections for the Fourth Rack in a Four-Rack System	24-14
24-10	Leaf Switch Connections for the First Rack in a Five-Rack System	24-15
24-11	Leaf Switch Connections for the Second Rack in a Five-Rack System	24-16
24-12	Leaf Switch Connections for the Third Rack in a Five-Rack System	24-16
24-13	Leaf Switch Connections for the Fourth Rack in a Five-Rack System	24-17
24-14	Leaf Switch Connections for the Fifth Rack in a Five-Rack System	24-17
24-15	Leaf Switch Connections for the First Rack in a Six-Rack System	24-18

24-16	Leaf Switch Connections for the Second Rack in a Six-Rack System	24-19
24-17	Leaf Switch Connections for the Third Rack in a Six-Rack System	24-19
24-18	Leaf Switch Connections for the Fourth Rack in a Six-Rack System	24-20
24-19	Leaf Switch Connections for the Fifth Rack in a Six-Rack System	24-20
24-20	Leaf Switch Connections for the Sixth Rack in a Six-Rack System	24-21
24-21	Leaf Switch Connections for the First Rack in a Seven-Rack System	24-22
24-22	Leaf Switch Connections for the Second Rack in a Seven-Rack System	24-22
24-23	Leaf Switch Connections for the Third Rack in a Seven-Rack System	24-23
24-24	Leaf Switch Connections for the Fourth Rack in a Seven-Rack System	24-23
24-25	Leaf Switch Connections for the Fifth Rack in a Seven-Rack System	24-24
24-26	Leaf Switch Connections for the Sixth Rack in a Seven-Rack System	24-24
24-27	Leaf Switch Connections for the Seventh Rack in a Seven-Rack System	24-25
24-28	Leaf Switch Connections for the First Rack in a Eight-Rack System	24-26
24-29	Leaf Switch Connections for the Second Rack in a Eight-Rack System	24-26
24-30	Leaf Switch Connections for the Third Rack in a Eight-Rack System	24-27
24-31	Leaf Switch Connections for the Fourth Rack in a Eight-Rack System	24-28
24-32	Leaf Switch Connections for the Fifth Rack in a Eight-Rack System	24-28
24-33	Leaf Switch Connections for the Sixth Rack in a Eight-Rack System	24-29
24-34	Leaf Switch Connections for the Seventh Rack in a Eight-Rack System	24-29
24-35	Leaf Switch Connections for the Eighth Rack in a Eight-Rack System	24-30
26-1	Recovery Appliance Repair Categories	26-1
26-2	Oracle Server LED Status Descriptions	26-3
26-3	Sun InfiniBand Switch LED Status Descriptions	26-5
26-4	Replacement Parts for Oracle Server X8 and X8M Database Servers	26-8
26-5	Replacement Parts for Oracle Exadata Storage Server X8M and X8 High Capacity and XT Servers	26-12
26-6	Replacement Parts for Cisco Switch	26-15
26-7	Replacement Parts for InfiniBand Switches	26-16
26-8	Replacement Parts for Power Distribution Units	26-17
26-9	Replacement Cables for Recovery Appliance	26-18

Recovery Appliance Preface

Welcome to *Oracle Zero Data Loss Recovery Appliance* .

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

Conventions

The following text conventions are used in this document:

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

Changes in *Zero Data Loss Recovery Appliance Owner's Guide* Release 23.1

The following are features in RA 23.1.

- Data Security
 - Database encryption best practices. Data is encrypted at the database-level. Data is automatically decrypted for the application. The encryption keys are only accessible by privileged databases.
 - Space-efficient, encrypted backups. Compression + Encryption + Incremental Forever for TDE and non-TDE database.
- Space Management and Operations
 - Auto-tune reserved space for compliance backups
 - Standby database registration to Recovery Appliance Catalog

The following are changes in *Zero Data Loss Recovery Appliance Owner's Guide* for RA 23.1.

- Data Security

The [Ransomware and Cybersecurity](#) section describes the role of the Recovery Appliance in preventing its dangers.
- RA23 details are provided in the [Recovery Appliance Rack Components](#) and [In-Rack Cabling Tables for Recovery Appliance RA23](#).
- Documents flows for discovery for Oracle Enterprise Manager Cloud Control
- This software release addresses US government regulation SEC 17a-4(f) regarding recovery window compliance and legal holds
- New or updated commands for RACLI.
 - [racli add keystore](#)
 - [racli add remote_syslog](#)
 - [racli alter adr_parameter](#)
 - [racli configure export](#)
 - [racli enable archived_log_grouping](#)
 - [racli disable archived_log_grouping](#)
 - [racli disable rpc](#)
 - [racli enable rpc](#)
 - [racli list env](#)
 - [racli list export](#)
 - [racli list remote_syslog](#)

- [racli remove osb_pieces](#)
- [racli remove remote_syslog](#)
- [racli status archived_log_grouping](#)
- Deprecated RACLI commands
 - [racli add vpc_user](#)
 - [racli list vpc_user](#)
 - [racli remove vpc_user](#)
 - [racli patch appliance](#)
 - [racli upgrade appliance](#)

1

Introduction to the Recovery Appliance

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

This chapter contains these sections:

- [About the Recovery Appliance](#)
- [Recovery Appliance Rack Components](#)
- [Compute Server Components](#)
- [Storage Server Components](#)
- [About Tape Backup Infrastructure](#)
- [Restrictions on Modifying Recovery Appliance](#)

About the Recovery Appliance

The Oracle Zero Data Loss 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 the 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 in the recovery window 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.

Ransomware and Cybersecurity

The Oracle Zero Data Loss Recovery Appliance is engineered for database ransomware protection. It has four key technology pillars:

- **Database Protection** includes real-time transaction protection and end-to-end ransomware protection and immutability.
- **Recovery Assurance** includes continuous backup validation, database protection monitoring, as well as high-speed, fast database restore capabilities through a dedicated network.
- **Resilient Architecture** built on a compute and storage servers foundation, which stems from Oracle Exadata engineered systems design methodology. The user model has a separation of duties; the roles for databases, the Recovery Appliance, and for any related appliances are segregated from each other. No one user can access other systems which they are not privileged to do so.

Immutable Backups prevents the backups themselves on a compromised system to be purged or deleted by internal processes or external users.

Ransomware: Today's Most Important Data Protection Topic

Ransomware is a type of malicious software, or *malware*, that encrypts a victim's data to make it unusable. A malicious cyber criminal holds the victim's data hostage with possible additional threats of data destruction or a public release until the victim pays the ransom. Ransomware has become the greatest threat to business survival in modern times that traditional High Availability (HA) and disaster recovery scenarios alone can't mitigate.

Databases typically contain the most critical and sensitive information of a company's operations. A few minutes or even a few seconds of a ransomware attack might corrupt hundreds to thousands of transactions. For mission-critical databases, if such attacks cause data to be lost and system downtime, these can have far-ranging impacts throughout the business in terms of revenue, operations, reputation, and even penalties.

In a typical ransomware situation that penetrates IT systems, the attack might start very innocuously as an email, fake software upgrade alerts, or some other kind of employee infiltration. It enters through access points that are known and common to employees.

The questions initially raised for IT: what processes are in place for securing credentials and access to downstream systems and for disallowing non-authorized access through internal/employee facing interfaces? Were these processes followed?

However, once this first level is breached, ransomware tries to steal additional privileged network access credentials both locally and domain-wide, with the goal of being able to access data. After it gets into a system and finds data, it can manipulate or publish that data for extortion purposes.

Once data is compromised, the next questions for IT become: what are the backup and patch strategies? Were these processes successfully followed?

Unfortunately, ransomware attacks can be so sophisticated that they can then move into the backups themselves. Production data is compromised as well as the backup infrastructure through connected NFS file shares or any kind of access point in which backups are made. It means that the last resort for data protection is potentially compromised.

The worst-case scenario is when infiltration passes all of the access levels and compromises both production and backup data, thereby offering no guarantee of data recovery. If the company's backup and recovery strategies were not been tested, were incomplete, or were not

able to be carried out, the company out of desperation to get the production data back resorts to paying the ransom. However, ransom payment does not guarantee that the data can be recovered once it has been compromised.

Best Practices to Minimize Ransomware Risks

According to the Cybersecurity & Infrastructure Security agency, the best practices to minimize ransomware risks are:

- Backup your data, system images, and configurations. Test your backups. Keep the backups offline.
- Utilize multi-factor authentication.
- Update and patch systems.
- Make sure your security solutions are up to date.
- Review and exercise your incident response plan.

The Recovery Appliance has Resiliency and Recoverability from Cyber-Attacks

The Oracle Zero Data Loss Recovery Appliance is designed to be fault-isolated from the production database. If a cyber-attack hits the production database, the Recovery Appliance is not compromised. This solution stems from the following key architectural features:

- **End-to-End Data Validation**

Validation is key to detecting corrupt backup data throughout the backup lifecycle, but it is equally important for detecting cyber-attacked data.

The Recovery Appliance validates all incoming, on-disk, and replicated backups for Oracle block correctness and recoverability. Any backup data maligned by malware or ransomware attack is detected, recorded, and alerted to the administrator. Action can then be taken in conjunction with the DBAs to disconnect the database from the network and investigate further.

Furthermore, replicated backups cannot be deleted or modified by the primary appliance or its administrators. They are independently validated and managed by the replication Recovery Appliance. They are shielded from any effects of attacks done on the primary Recovery Appliance.

As an alternative or supplementary protection strategy, backups can be archived to Oracle Cloud Storage as a secure location for secondary backup copies. This uses Oracle Key Vault as key store for backup encryption keys. All backups remain encrypted in Cloud Storage. Users require access to the Recovery Appliance and Oracle Key Vault to perform restore operations. The Recovery Appliance can also archive backups to fibre-attached tape libraries via Oracle Secure Backup. Tapes can then be shipped to and stored in an off-site, network disconnected, location that is impervious to cyber-attacks.

- **Air-Gapped Vault Backups**

With Recovery Appliance database-aware incremental-forever replication, the vault appliance is configured behind a firewall which has a window open only during certain times of the day. Recovery Appliance replication proceeds during those times to synchronize the vault appliance. When the firewall is closed, replication pauses. Upon the next open sync window, replication resumes. With incremental forever-based replication, only the minimum amount of data is required to maintain full recoverability from the vault. Unlike with general purpose storage appliances, no full backups are transmitted, which thus limits the sync window and possibility of malicious access to the vault.

With cyber vault deployment, you have a physically network-separated copy of the backups, which can be restored at any time, even if production systems are compromised.

Recovery Appliance in the vault are independently managed, both in their access and policies for retention and other backup attributes. It has full reporting available on the Recovery Appliance that is crucial for audit compliance purposes.

- **Separation of Duty**

Access to the system is controlled via strict separation of duty between DBA and Recovery Appliance administrator roles. DBAs are only given Virtual Private Catalog (VPC) user roles to backup and recover their privileged databases. They cannot access, modify, or delete backups on the Recovery Appliance.

Recovery Appliance administrators only have access to manage and monitor the system, but cannot backup, recover, or modify protected databases. The Recovery Appliance does not expose or allow creation of local users, databases, or other services.

- **Limited Network Access**

With regards to network protocols, VPC users can only connect to the appliance through SQL*Net. HTTPS is used for RMAN backup and restore traffic through the Recovery Appliance Backup Module. No other protocols are employed.

The Recovery Appliance enforces network segregation with the support of VLAN tagged networks, allowing backup and restore traffic to be fully isolated and non-routable between protected databases' specific network zones. In this way, any possibly affected backups would not be exposed to the rest of the enterprise.

- **Superior Resiliency**

The Recovery Appliance itself offers superior resiliency capabilities against cyber-attacks, when compared with traditional backup devices. As an Oracle Engineered System built on Exadata hardware and storage, the Recovery Appliance inherits a resilient architecture for reducing surface of attack on compute and storage servers. This includes:

- hardened password policies
- OS and DB user auditing
- firewall support
- Oracle ILOM (Integrated Lights Out Management)

Real-Time Monitoring & Audit Reporting Enterprise Manager provides 360 degree view of database recoverability and system metrics. This includes audit report of all system access and activities, such as changing retention policies or deleting database backups, which could be evidence of malicious user.

Recovery With No Data Loss

In the event that a database server is attacked and its backups must be recovered to a different server, the Recovery Appliance's real-time redo transport allows recovery to the very last transaction prior to the attack occurrence. This is especially important for cyber-attacks, such as ransomware, where paying the perpetrators does not always mean your data comes back in pristine condition. With the Recovery Appliance, you recover the database to a separate, safe location with no data loss and don't have to pay the ransom.

Recovery Appliance Rack Components

The Oracle Zero Data Loss Recovery Appliance consists of compute servers, storage servers, and the network components to connect to your network. Oracle offers a flexible configuration for the 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.

- [RA23](#)
- [RA21](#)
- [X8M](#)

RA23

[Table 1-1](#) lists the components of a RA23 rack that can support a total of 17 storage servers, which is a usable capacity under normal redundancy of 92 TB per storage server, or minimum (3 storage servers) 274 TB and full rack total 1.57 PB.

In high redundancy configurations, the usable capacity is 60 TB per storage server, or minimum (3 storage servers) 174 TB and full rack total 1 PB.

Table 1-1 Components of a Oracle Zero Data Loss Recovery Appliance RA23 Rack

Quantity	Component
2	Oracle Zero Data Loss Recovery Appliance RA23 compute servers
3 to 17	Oracle Zero Data Loss Recovery Appliance RA23 storage servers
2	36-Port 100 Gb/s RoCE switches
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Management Ethernet switch for administration.

RA21

[Table 1-2](#) lists the components of a RA21 rack that can support a total of 18 storage servers, which is a usable capacity under normal redundancy of 70 TB per storage server, or minimum (3 storage servers) 207 TB and full rack total 1.26 PB.

In high redundancy configurations, the usable capacity is 44 TB per storage server, or minimum (3 storage servers) 129 TB and full rack total 794 TB.

Table 1-2 Components of a Oracle Zero Data Loss Recovery Appliance RA21 Rack

Quantity	Component
1	Oracle Zero Data Loss Recovery Appliance RA21 compute servers
3 to 18	Oracle Zero Data Loss Recovery Appliance RA21 storage servers
2	36-Port Cisco Nexus 9336c Ethernet Switches
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco 9000 series Ethernet switch
2	32 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports. Optional with OSB tape.

X8M

[Table 1-3](#) lists the components of a Recovery Appliance X8M rack that can support a total of 18 storage servers.

Table 1-3 Components of a Oracle Zero Data Loss Recovery Appliance X8M Rack

Quantity	Component
2	Oracle Zero Data Loss Recovery Appliance X8M compute servers
3 to 18	Oracle Zero Data Loss Recovery Appliance X8M storage servers
2	36-Port Cisco Nexus 9336c Ethernet Switches
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco 9000 series Ethernet switch
2	32 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports. Optional with OSB tape.

- [X8](#)
- [X7](#)
- [X6](#)
- [X5](#)
- [X4](#)

X8

[Table 1-4](#) lists the components of a Recovery Appliance X8 rack that can support a total of 18 storage servers.

Table 1-4 Components of a Recovery Appliance X8 Rack

Quantity	Component
2	Oracle Server X8 compute servers
3 to 18	Oracle Storage Server X8 storage servers
2	36-port Sun Datacenter InfiniBand Switch
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco 9000 series Ethernet switch
2	32 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports. Optional with OSB tape.

X7

[#GUID-238808B2-5E4E-4DFF-B1B7-688E8B8E8C92/COMPONENTSOFARECOVERYAPPLIANCEX7RAC-36AC788C](#) lists the components of a Recovery Appliance X7 rack that can support a total of 18 storage servers.

Table 1-5 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	36-port Sun Datacenter InfiniBand Switch
2	Redundant 15 kVA PDUs (single phase or three phase, high voltage or low voltage)
1	48-port Cisco 9000 series Ethernet switch
2	32 Gbps port Oracle fiber channel adapter with 2 x 16 Gb fiber channel ports. Optional with OSB tape.

X6

Table 1-6 lists the components of a Recovery Appliance X6 rack that can support a total of 18 storage servers.

Table 1-6 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

X5

Table 1-7 lists the components of a Recovery Appliance X5 rack.

Table 1-7 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)
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

X4

Table 1-8 lists the components of a Recovery Appliance X4 rack.

Table 1-8 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

Compute Server Components

- [RA23](#)
- [RA21](#)
- [X9M](#)
- [X8M](#)

RA23

The Oracle Zero Data Loss Recovery Appliance RA23 compute server is based on the Oracle Exadata Database Server.

Compute Server	Description
Storage	2x NVMe-Flash SSD (hot-swappable)

Compute Server	Description
Network	• 2 x Dual Port 10/25 Gb Ethernet SFP28 – 2 x 10/25 Gb optical (ingest) – 2 x 10/25 Gb optical (replication) • 2 x Dual Port 100 Gb Ethernet QSFP28 – 2 x 100 Gb optical (ingest) – 2 x 100 Gb optical (replication) • 1 x Quad Port 10 Gb Ethernet RJ45 – 2 x 10 Gb copper (ingest) – 2 x 10 Gb copper (replication) • For ingest network, maximum of – 2 x 10 Gb Ports, – 2 x 25 Gb Ports, or – 2 x 100 Gb Ports • For replication network, maximum of – 2 x 10 Gb Ports, – 2 x 25 Gb Ports, or – 2 x 100 Gb Ports • Ingest and Replication can be different, for example 100Gb for ingest and 25Gb for replication. • Optional Ports – Sun Storage Dual 32 Gb Fibre Channel PCIe Universal HBA, QLogic for tape connectivity • Standard Ports – 2 x 100 Gb QSFP28 RoCE Fabric Ports – 1 x 1 Gb copper Ethernet Port (mgmt) – 1 x ILOM Ethernet Port

All servers include redundant hot swappable fans and power supplies.

RA21

The Oracle Zero Data Loss Recovery Appliance RA21 compute server is based on the Oracle Exadata Database Server.

[Table 1-9](#) lists the components of an individual RA21 compute server.

Table 1-9 RA21 Compute Server Components

Compute Server	Description
Storage	2x NVMe-Flash SSD (hot-swappable)

Table 1-9 (Cont.) RA21 Compute Server Components

Compute Server	Description
Network	- Dual Port 25 Gb Ethernet SFP28 or Quad Port 10 Gb Ethernet RJ45 (ingest or replication) 2 x 10/25 Gb optical, or 2 x 10 Gb copper (ingest) 2 x 10 Gb copper (replication) - Dual Port 25 Gb Ethernet SFP28 (ingest or replication) 2 x 10/25 Gb optical - For ingest network, maximum of 2 x 10 Gb Ports, or 2 x 25 Gb Ports - For replication network, maximum of 2 x 10 Gb Ports, or 2 x 25 Gb Ports - Optional Ports - Sun Storage Dual 32 Gb Fibre Channel PCIe Universal HBA, QLogic for tape connectivity - Standard Ports 2 x 100 Gb QSFP28 RoCE Fabric Ports 1 x 1 Gb Copper Ethernet Port (mgmt) 1 x ILOM Ethernet Port

X9M

The Oracle Zero Data Loss Recovery Appliance X9M is based on the Oracle Exadata Database Server X9M. [Table 1-10](#) lists the components of an individual X9M compute server.

Table 1-10 X9M Compute Server Components

Quantity	Description
2	32-core Intel Xeon 8358 processors (2.6 GHz)
8	512 GB RAM
2	100Gb/s RDMA Network Fabric Card
1	1 GbE Base-T Ethernet port for Integrated Lights Out Manager (ILOM) for remote management
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

Table 1-10 (Cont.) X9M Compute Server Components

Quantity	Description
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
2	Redundant Hot-Swappable Fans
1	Disk Controller HBA with 2 GB cache

X8M

Table 1-11 lists the components of an individual X8M compute server.

Table 1-11 Oracle Zero Data Loss Recovery Appliance X8M Compute Server Components

Quantity	Description
2	24-core Xeon 8160 processors (2.4 GHz)
8	384 GB RAM
2	100Gb/s RDMA Network Fabric Card
1	10/100/1000BASE-T Ethernet port for Integrated Lights Out Manager (ILOM) for remote management
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
2	Redundant Hot-Swappable Fans
1	Disk Controller HBA with 2 GB cache

- X8
- X7

- X6
- X5
- X4

X8

Table 1-12 lists the components of an individual Oracle X8 compute server.

Table 1-12 Oracle X8 Compute Server Components

Quantity	Description
2	24-core Xeon 8160 processors (2.4 GHz)
8	384 GB RAM
2	QDR InfiniBand (40 Gb/s) ports
1	10/100/1000BASE-T Ethernet port for Integrated Lights Out Manager (ILOM) for remote management
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
2	Redundant Hot-Swappable Fans
1	Disk Controller HBA with 2 GB cache

X7

Table 1-13 lists the components of an individual Oracle X7-2 compute server.

Table 1-13 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

Table 1-13 (Cont.) Oracle X7-2 Compute Server Components

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

X6

Table 1-14 lists the components of an Oracle X6-2 compute server.

Table 1-14 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)

X5

Table 1-15 lists the components of an Oracle X5-2 compute server.

Table 1-15 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

Table 1-15 (Cont.) Oracle Server X5-2 Components

Quantity	Description
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)

X4

[Table 1-16](#) lists the components of a Sun X4-2 compute server.

Table 1-16 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)

Storage Server Components

- [RA23](#)
- [RA21](#)
- [X8M](#)
- [X7](#)
- [X6](#)
- [X5](#)

- X4

RA23

Oracle Zero Data Loss Recovery Appliance Storage Server RA23 is based on Oracle Exadata Storage Server High Capacity (HC).

Refer to [Compute Server Components](#) for RA23 Server details.

Storage Server	Description
Storage	12 x 22 TB 7,200 RPM disks 2 x NVMe Flash Acceleration Cards
Network	• 2 x 100 Gb QSFP28 RoCE Fabric ports • 1 x 1 Gb copper Ethernet port (mgmt) • 1 x ILOM Ethernet port

All servers include redundant hot swappable fans and power supplies.

RA21

Oracle Zero Data Loss Recovery Appliance Storage Server RA21 is based on Oracle Exadata Storage Server High Capacity (HC).

Refer to [Compute Server Components](#) for RA21 details.

X8M

[Table 1-17](#) lists the components of an X8 and X8M storage server.

Table 1-17 Oracle Zero Data Loss Recovery Appliance X8 and X8M Server Storage Components

Quantity	Description
2	16-core Intel Xeon 5218 Processors (2.3GHz)
8	192 GB RAM
12	14 TB High Capacity SAS disks
4	6.4 TB flash accelerator PCIe cards
1	Disk controller HBA with 2 GB cache
2	X8M only: 100Gb/s RDMA Network Fabric Card, 1 dual-port (PCIe 3.0), both ports active
2	X8 only: InfiniBand 4X QDR (40 Gb/s) InfiniBand ports, 1 dual-port PCIe 3.0 Host Channel Adapter (HCA), all ports active
2	Embedded Gigabit Ethernet ports
1	10/100/1000BASE-T Ethernet port for Integrated Lights Out Manager (ILOM) for remote management
2	Redundant Hot-Swappable Power Supplies
2	Redundant Hot-Swappable Fans

X7

[Table 1-18](#) lists the components of an Oracle X7–2L storage server.

Table 1-18 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

X6

Table 1-19 lists the components of an Oracle X6-2L storage server.

Table 1-19 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)
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

X5

Table 1-20 lists the components of an Oracle Server X5-2L storage server.

Table 1-20 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)

Table 1-20 (Cont.) Oracle Server X5-2L Components

Quantity	Description
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

X4

[Table 1-21](#) lists the components of a Sun Server X4-2L storage server.

Table 1-21 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
1	Ethernet port for ILOM for remote management

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](#)
- [Third-Party Tape Systems](#)

Oracle Recommended Stack

Recovery Appliance is engineered to work with industry leading tape infrastructure components, including those described in [Table 1-22](#).

Table 1-22 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.

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.

Recovery Appliance allows the use of third-party media management software that offers RMAN-integrated backup module with its client software. Only the *client* software component of a third-party media management product for copy to tape operations may be installed on the Recovery Appliance compute servers.

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.
- Refer to [My Oracle Support note 2014361.1](#) for restrictions and guidance on third party software installation.

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.

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:

- [Site Requirements for Recovery Appliance and Oracle Exadata Storage Expansion Rack](#)
- [Preparing Your Networks for Recovery Appliance](#)
- [Setting Up Auto Service Request](#)
- [Using Oracle Exadata Deployment Assistant](#)

2

Site Requirements for Recovery Appliance and Oracle Exadata Storage Expansion Rack

This chapter describes the site requirements for Recovery Appliance, sometimes referenced as "ZDLRA Rack".

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*

Ensuring That the Site is Ready

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

Task 1 Review Site Requirements

Review the site requirements in this chapter to understand the requirements for ZDLRA Rack.

Task 2 Understand the Networking Requirements and Options

Review [Preparing Your Networks for Recovery Appliance](#) to discern the networking requirements and configuration decisions that you must make.

Task 3 Run Oracle Exadata Deployment Assistant (OEDA)

Run OEDA to create the necessary configuration files.
See [Using Oracle Exadata Deployment Assistant](#).

Task 4 Configure the Network

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

Task 5 Prepare the Site Based on Requirements

Prepare the site based on the requirements, such as, install the network cables and power supplies, prior to the arrival of ZDLRA Rack.
Review the procedures in [#unique_49](#)

General Environmental Requirements

The following sections describe the general environmental requirements for ZDLRA Racks.

General Environmental Requirements for ZDLRA Rack X6 and Later

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

Related Topics

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

Space Requirements

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

In addition to the rack height, 914 mm (36 inches) of space is required above the rack height for maintenance access, as described in [Table 2-3](#).

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.

Related Topics

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

Space Requirements for Racks up to Recovery Appliance X6

All racks up to Recovery Appliance X6-2 use the same hardware rack, and have the same space requirements.

The space requirements are as follows:

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

The minimum ceiling height for the cabinet is 2912 mm (114.65 inches), measured from the true floor or raised floor, whichever is higher. This includes an additional 914 mm (36 inches) of space required above the rack height for maintenance access, as described in [Table 2-3](#). 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.

Receiving, Unpacking, and Access Route Requirements

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

Use the following package dimensions for ZDLRA Rack:

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

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

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

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

**Caution:**

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

Table 2-1 Access Route Requirements

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

Rack Weights for Recovery Appliance

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

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

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

Table 2-2 Shipping Weight

Recovery Appliance	Full Rack	Minimum Configuration
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)

Related Topics

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

Maintenance Access Requirements

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

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

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

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

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

Table 2-3 Maintenance Access Requirements for ZDLRA Rack (All Models)

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



Note:

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

Flooring Requirements

ZDLRA Rack may be installed on raised floor or solid floor environments.

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

Table 2-4 Floor Load Requirements for ZDLRA Rack

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

**Note:**

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

Electrical Power Requirements

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

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

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

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

PDU Power Requirements

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

The following are the PDUs that Oracle Zero Data Loss Recovery Appliance currently supports or has supported in the past:

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

**Note:**

The following specifications are for a single PDU. Every ZDLRA Rack configuration contains two PDUs.

Table 2-5 Specifications for each Low-Voltage 15 kVA Single Phase PDU for North America, South America, Japan and Taiwan

Specification	Value
Number of Inputs	3
Source Voltage	200-240V
Frequency	50/60 Hz

Table 2-5 (Cont.) Specifications for each Low-Voltage 15 kVA Single Phase PDU for North America, South America, Japan and Taiwan

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

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



Note:

The following specifications are for a single PDU. Every ZDLRA Rack configuration contains two PDUs.

Table 2-6 Specifications for each Low-Voltage 22 kVA Single Phase PDU for North America, South America, Japan and Taiwan

Specification	Value
Number of Inputs	3
Source Voltage	200-240V
Frequency	50/60 Hz
Max Line Current	36.8A
Phase Current	N/A
Nominal Power Rating	22kVA
Recommended Circuit Breaker	50A
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Max Group Current	20A
Output Voltage	200-240V
Data Center Receptacle	1ph 2W + ground 30A NEMA L6-30

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

**Note:**

The following specifications are for a single PDU. Every ZDLRA Rack configuration contains two PDUs.

Table 2-7 Specifications for each Low-Voltage 15 kVA Three Phase PDU for North America, South America, Japan and Taiwan

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

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

**Note:**

The following specifications are for a single PDU. Every ZDLRA Rack configuration contains two PDUs.

Table 2-8 Specifications for each Low-Voltage 24 kVA Three Phase PDU for North America, South America, Japan and Taiwan

Specification	Value
Number of Inputs	2
Source Voltage	200-220V
Frequency	50/60 Hz
Max Line Current	34.6A
Phase Current	20A
Nominal Power Rating	24kVA
Recommended Circuit Breaker	50A

Table 2-8 (Cont.) Specifications for each Low-Voltage 24 kVA Three Phase PDU for North America, South America, Japan and Taiwan

Specification	Value
Outlets	42 x C13, 6 x C19
Outlet Groups	6
Max Group Current	20A
Output Voltage	200-240V
Data Center Receptacle	3ph 3W + ground 60A IEC60309

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



Note:

The following specifications are for a single PDU. Every ZDLRA Rack configuration contains two PDUs.

Table 2-9 Specifications for each High-Voltage 15 kVA Single Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

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



Note:

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

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

**Note:**

The following specifications are for a single PDU. Every ZDLRA Rack configuration contains two PDUs.

Table 2-10 Specifications for each High-Voltage 22 kVA Single Phase PDUs for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

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

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

**Note:**

The following specifications are for a single PDU. Every ZDLRA Rack configuration contains two PDUs.

Table 2-11 Specifications for each High-Voltage 15 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

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

Table 2-11 (Cont.) Specifications for each High-Voltage 15 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

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

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



Note:

The following specifications are for a single PDU. Every ZDLRA Rack configuration contains two PDUs.

Table 2-12 Specifications for each High-Voltage 24 kVA Three Phase PDU for Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan

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

Facility Power Requirements

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

Use dedicated AC breaker panels for all power circuits that supply power to the PDU. When planning for power distribution requirements, balance the power load between available AC

supply branch circuits. In the United States of America and Canada, ensure that the overall system AC input current load does not exceed 80 percent of the branch circuit AC current rating.



Note:

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

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

Circuit Breaker Requirements

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

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

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



Note:

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

Electrical Grounding Guidelines

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

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

Temperature and Humidity Requirements

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

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

**Note:**

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

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

Table 2-13 Temperature, Humidity and Altitude Requirements

Condition	Operating Requirement	Non-operating Requirement	Optimum
Temperature	5 to 32 degrees Celsius (41 to 89.6 degrees Fahrenheit)	-40 to 70 degrees Celsius (-40 to 158 degrees Fahrenheit).	For optimal rack cooling, data center temperatures from 21 to 23 degrees Celsius (70 to 74 degrees Fahrenheit)
Relative humidity	10 to 90 percent relative humidity, non-condensing	Up to 93 percent relative humidity.	For optimal data center rack cooling, 45 to 50 percent, non-condensing
Altitude	3048 meters (10000 feet) maximum	12000 meters (40000 feet).	Ambient temperature is reduced by 1 degree Celsius per 300 m above 900 m altitude above sea level

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

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

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

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

Electrostatic discharge (ESD) is easily generated, and hard to dissipate in areas of low relative humidity, such as below 35 percent. ESD becomes critical when humidity drops below 30

percent. It is not difficult to maintain humidity in a data center because of the high-efficiency vapor barrier and low rate of air changes normally present.

Ventilation and Cooling Requirements

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

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

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

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

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


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

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

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

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

Figure 2-1 shows a typical installation of the floor tiles for a full rack.

Figure 2-1 Typical Data Center Configuration for Perforated Floor Tiles

Network Connection and IP Address Requirements for ZDLRA Rack

This section describes the requirements to connect the Oracle Zero Data Loss Recovery Appliance to your existing network infrastructure.

Network Connection Requirements for ZDLRA Rack

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

Network	Connection Requirements
Administration Network	• Mandatory: 1 uplink for the Management Network Switch. • Recommended: 2 uplinks for remote monitoring of the power distribution units (PDUs). • Optional: 1 uplink for ILOM on each database server or storage server if using a separate ILOM network.
Client Network	• For bonded network configurations: 2 uplinks for each database server. • For non-bonded network configurations: 1 uplink for each database server. Note: Non-bonded network configurations are not supported on Recovery Appliance X7 and later systems.

Network	Connection Requirements
Additional Network(s)	- For bonded network configurations: 2 uplinks for each database server and additional network. - For non-bonded network configurations: 1 uplink for each database server and additional network. Note: Non-bonded network configurations are not supported on Recovery Appliance X7 and later systems.
RDMA Network Fabric (Private Network)	No uplinks are required. The RDMA Network Fabric is completely self-contained within Recovery Appliance.

DNS Configuration for ZDLRA Rack

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

The assistant-generated file defines the SCAN as a single name with three IP addresses on the client access network. The three SCAN addresses provide service access for clients to Oracle Zero Data Loss 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.

IP Address Requirements for ZDLRA Rack

The Recovery Appliance requires a dedicated allocation of host names and IP addresses. The number of required IP addresses primarily depends on the size of the system.

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

You should configure the IP addresses in your existing networks after you run OEDA. All IP addresses in the configuration file must be unassigned at the time of initial configuration. In addition, all IP addresses must be statically-assigned IP addresses, not dynamically-assigned (DHCP) addresses.

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

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

Rack-Level IP Address Requirements

The following table outlines the rack-level IP address requirements for Recovery Appliance. These addresses are required regardless of the configuration options that are selected for the system.

Network	IP Address Requirements
Administration Network	1 IP address for the administration network interface on each compute server or storage server. 1 IP address for the ILOM network interface on each database server or storage server. 1 IP address for the management interface on each RDMA Network Fabric switch. 1 IP address for the management interface on the Management Network Switch switch. 1 IP address for the management interface in each power distribution unit (PDU).
RDMA Network Fabric (Private Network)	2 IP addresses for each database server or storage server.

Related Topics

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

Dedicated Backup Network between Oracle Exadata and Oracle Zero Data Loss Recovery Appliance

This section describes an Oracle convention for a dedicated backup network between Oracle Exadata and Oracle Zero Data Loss Recovery Appliance RA23. Its purpose is to facilitate Oracle Exadata customers using 10G networks to take advantage of the 100GE NICs option. The networking default configuration on RA23 remains 2 x 25 GE ports for ingest and 2 x 25GE ports for replication per compute server. A small benefit is that the dedicated network does increase network bandwidth for backup and restore traffic to the Oracle Zero Data Loss Recovery Appliance (ZDLRA).

The convention calls for two 100GE switches for the rack.

Table 2-14 Ethernet Cabling for the Recovery Appliance RA21 Racks

R23 Connection	Exadata X10M	Exadata X9M and earlier
100 GE	25GE NICs connected via Split Cables, or - Direct Connection from Exadata Database Server X10M nodes 100GE NICs to RA23 TOR switches	

Table 2-14 (Cont.) Ethernet Cabling for the Recovery Appliance RA21 Racks

R23 Connection	Exadata X10M	Exadata X9M and earlier
25 GE	Connection from Exadata Database Server X10M node's 25GE NICs to RA23 TOR switches through: • 25/10 SFPs, or • Split cables	Connection from Exadata Database Server X9M node's 25GE NICs to RA23 TOR switches through: • 25/10 SFPs, or • Split cables
10 GE	TOR Switch is not recommended.	Use appropriate 10/40 split cables

* Using 100GE NICs in Exadata Database Server X10M nodes is not necessary.

Using the 10 GE port of the Exadata Database Server X10M with the ZDLRA TOR switches is inefficient, and not using them for this purpose keeps them available should older Exadata systems need to be connected.

Exadata Database Server X10M nodes can connect their 25GE NICs with 25/100GE Splitter Cables (1x100GE to 4x25GE).


25/100 GbE requires OM4 cables and splitters and (Q)SFP28 transceivers. 10 GbE requires OM3 cables and splitters and (Q)SFP+ transceivers.

If the Exadata and RA23 are in adjacent racks and 5 meters is sufficient, use the 7118366 Copper Splitter Cable assembly: 5 meters, QSFP28 to 4 SFP28 part to connect the database machine nodes to the RA23 ToR switches.

For longer distances, you must use optical connections. There are three parts to this:

- QSFP28 transceivers, part 7604159, to connect the cable to the switch. Two are required for every four database server nodes.
- Optical splitter cables to run to the Exadata rack. Again, two are required for every four database server nodes, and the quantity chosen is per controller.
- Exadata has 25 GbE transceivers that are ordered with Exadata and not with RA23. Two are required for each database server node in the Exadata.

Figure 2-2 ZDLRA RA23 100GE to Exadata X10



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](#)
- [Registering Recovery Appliance in the Domain Name System](#)
- [Factory IP Address Settings](#)
- [Port Assignments When Using a Firewall](#)

Overview of Network Requirements

In addition to the compute and storage servers, Recovery Appliance includes equipment to connect the system to your network. The network connections allow clients to connect to the compute servers and also enables remote system administration.

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

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

- **Administration Network:** Also known as the management network, this required network connects to your existing management network infrastructure, and is used for administrative work on all components of Recovery Appliance. By default, the administration network connects the compute servers, storage servers, server Integrated Lights Out Manager (ILOM) interfaces, and RDMA Network Fabric switches to the Management Network Switch in the rack. One uplink is required from the Management Network Switch to your management network.

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

 Notes:

- Separate uplinks to your management network are also recommended for remote monitoring of each power distribution unit (PDU). This configuration enables you to easily differentiate between system outages caused by PDU failure as opposed to failure of the Management Network Switch.
- A properly secured configuration requires full isolation of the administration network from all other networks.

- **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 the 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.

- **Private Network:** Also known as the RDMA Network Fabric, storage network, or interconnect. This network connects the compute servers and storage servers. Oracle Database uses this network for Oracle RAC cluster interconnect traffic and for accessing data on the Oracle Exadata Storage Servers. The private network is automatically configured during installation. It is non-routable, fully contained in Recovery Appliance, and does not connect to your existing networks.

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

Previously, the private network was built using InfiniBand technology. RoCE Network Fabric uses different switches and cables from those used by InfiniBand Network Fabric.

- **Replication network:** The optional replication network uses available ports not used by the administration and the ingest network. It 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:

 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.

The replication network must not be used for the purpose of ingesting backups.

- **Fiber 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.

Ingest and replication networks can be configured active/passive or active/active bonding.


- Active / Passive Bonding - `BONDING_OPTS="mode=active-backup miimon=100 downdelay=2000 updelay=5000 num_grat_arp=100"`
- Active / Active Bonding - `BONDING_OPTS="mode=802.3ad miimon=100 downdelay=200 updelay=200 lacp_rate=1 xmit_hash_policy=layer3+4"`

Ingest can be Active/Active with Replication Active/Passive, or vice-versa. Or both can have the same bonding.

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"](#) for information about how a fibre channel SAN network is configured for backups to tape in a Recovery Appliance environment.

The following diagram displays how the various Recovery Appliance components connect to the different networks.



About the Network Components and Interfaces

Each compute server has varying network components and interfaces:

-
- [RA23](#)
 - [RA21](#)
 - [X8M](#)
 - [X8_X7](#)


RA23

Each compute server in the RA23 configuration consists of the following network components and interfaces:

- 2 x Dual Port 10/25 Gb Ethernet SFP28
 - 2 x 10/25 Gb optical (ingest)
 - 2 x 10/25 Gb optical (replication)
- 2 x Dual Port 100 Gb Ethernet QSFP28
 - 2 x 100 Gb optical (ingest)
 - 2 x 100 Gb optical (replication)
- 1 x Quad Port 10 Gb Ethernet RJ45
 - 2 x 10 Gb copper (ingest)
 - 2 x 10 Gb copper (replication)
- For ingest network, maximum of
 - 2 x 10 Gb Ports,
 - 2 x 25 Gb Ports, or
 - 2 x 100 Gb Ports
- For replication network, maximum of
 - 2 x 10 Gb Ports,
 - 2 x 25 Gb Ports, or
 - 2 x 100 Gb Ports
- Ingest and Replication can be different, for example 100Gb for ingest and 25Gb for replication.
- Optional Ports
 - Sun Storage Dual 32 Gb Fibre Channel PCIe Universal HBA, QLogic for tape connectivity
- Standard Ports
 - 2 x 100 Gb QSFP28 RoCE Fabric Ports
 - 1 x 1 Gb copper Ethernet Port (mgmt)

- 1 x ILOM Ethernet Port

Figure 3-1 Oracle Zero Data Loss Recovery Appliance RA23 Networking




RA21

Each compute server in the RA21 configuration consists of the following network components and interfaces:

- Ethernet ingest and replication network connectivity, :
 - 2 x Dual 10/25G Network Cards, or
 - 2 x Quad 10G Network Cards, or
 - 1 x Quad 10G Network Card and 1 x Dual 10/25G Network Card
- 2 x QSFP2B RoCE Fabric Ports
- 1 Ethernet port for Serial MGT remote management
- 1 Ethernet port for ILOM MGT (Oracle Integrated Lights Out Manager) remote management
- 1 Ethernet port for HOST MGT remote management
- Optional: Dual 32G HBA (Tape) card that can be field installed. This slot is not available for other network cards.

Figure 3-2 Oracle Zero Data Loss Recovery Appliance RA21 Backplane External Network Connectivity




- Maximum of two 10G (or 25G) ports for ingest and two 10G (or 25G) ports for replication, per compute server
- Maximum of four 10G (or 25G) for each network in bonded LACP configuration, per rack
- Replication can be used as separate ingest network ([MOS Note 2126047.1](#))
- VLAN tagging supported on ingest network ([MOS Note 2047411.1](#))
- Slot 2 cannot be used for additional network card.

X8M

Each compute server in the X8M configuration consists of the following network components and interfaces:

- Ethernet ingest and replication network connectivity, :
 - 2 x Dual 10/25G Network Cards, or
 - 2 x Quad 10G Network Cards, or
 - 1 x Quad 10G Network Card and 1 x Dual 10/25G Network Card
- 2 x QSFP2B RoCE Fabric Ports
- 1 Ethernet port for `Serial MGT` remote management
- 1 Ethernet port for `ILOM MGT` (Oracle Integrated Lights Out Manager) remote management
- 1 Ethernet port for `HOST MGT` remote management
- Optional: Dual 32G HBA (Tape) card that can be field installed. This slot is not available for other network cards.

Figure 3-3 Oracle Zero Data Loss Recovery Appliance X8M Backplane External Network Connectivity



- Maximum of two 25G ports for ingest and two 25G ports for replication, per compute server

X8_X7

Each compute server in the X8 and 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 3-4 Oracle Zero Data Loss Recovery Appliance X8 and X7 Backplane 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.

Example of Network Connections for Recovery Appliance

Figure 3-5 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 3-5 Network Diagram for Recovery Appliance



Connecting Recovery Appliance Rack Components to the Networks


Figure 3-6 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 RDMA Network Fabric 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 RDMA Network Fabric network connects the switches to the compute servers and the storage servers.

Figure 3-6 Network Connections to the Recovery Appliance Rack Components



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 management system. See ["About Tape Backup Infrastructure"](#) for the differences in support provided by Recovery Appliance.

Oracle Recommended Stack

When you use the Oracle compatible 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 3-7](#) provides an overview of the network connections when using an Oracle tape system.


Figure 3-7 Recovery Appliance Connection to an Oracle Tape System



Third-Party Tape Systems

When you use a third-party tape system, the backups to tape use the ingest network. This is the same network that the local protected databases use to backup to Recovery Appliance. [Figure 3-8](#) provides an overview of the network connections when using a third-party tape system.

Figure 3-8 Recovery Appliance Connection to a Third-Party Tape System



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](#)" for instructions on when and how to configure VLAN tagging

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 Oracle Exadata Deployment Assistant](#)
- *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

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](#)" for additional information about the `checkip.sh` script.

[Table 3-1](#) 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

Table 3-1 (Cont.) Factory IP Addresses for Recovery Appliance

Rack Unit	Component	Management Network Addresses	InfiniBand Active Bonded IP Addresses	Oracle ILOM IP Addresses
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	RDMA Network Fabric switch	Not applicable	Not applicable	192.168.1.203
U21	Ethernet switch	Not applicable	Not applicable	192.168.1.200
U20	RDMA Network Fabric 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

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 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
NA	Compute servers, storage servers, and InfiniBand ILOMs	LDAP over UDP/TCP	389	Outgoing LDAP (Lightweight Directory Access Protocol) (user configurable)

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

Source	Target	Protocol	Port	Application
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)
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

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

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

Proxy Servers

The Recovery Appliance does **not** support the use of proxy servers.

The client systems for protected databases should not force http(s) proxy settings in environment variables for RMAN sessions. The SBT library will consider those proxy settings that are specified as part of the environment setup, and therefore will use the corresponding http proxy setting. In other words, the SBT library would send backups to the proxy server, which may not be configured properly to support backups and may be unaware of the Recovery Appliance.

The proxy server environment variables must **not** be used at all for the Recovery Appliance.

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](#)
- [Getting Ready to Install ASR](#)
- [Installing ASR Manager](#)
- [Verifying the ASR Manager Installation](#)
- [About the Trap Destinations on Recovery Appliance](#)
- [Troubleshooting ASR](#)

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>

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](#) 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 shows the network connections between ASR and Recovery Appliance.

Figure 4-1 Auto Service Request Network Connections



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>

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.

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.

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

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

Example 4-2 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.
```

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 Oracle Exadata Deployment Assistant

This chapter describes how to use the Oracle Exadata Deployment Assistant, which is not only used for Exadata but also Zero Data Loss Recovery Appliance and other Oracle Engineered Systems. 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.

Overview of Oracle Exadata Deployment Assistant

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

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

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

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

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

Oracle Exadata System Software release 19.1.0 introduced the Web-based interface for OEDA. However, the overall process, described here, is the same regardless of whether you use the Web-based or Java-based user interface.

- Before your engineered system arrives, do the following:
 1. Work with your network and database administrators to evaluate the current network settings, such as current IP address use and network configuration. OEDA supports IPv6 addresses.
 2. Define the settings for the rack, such as network configuration and backup method.
 3. Download the latest version of OEDA from Oracle Technology Network.
 4. Unzip and run the configuration script on a supported platform, and select the type of system to configure.
 5. Go through every page in OEDA and supply values for all required fields. You cannot advance to the next page if you do not supply all of the required values. You must provide naming details and DNS, NTP, and networking details.
 6. At the end of the dialogue with OEDA, configuration files are generated on the client. The files are also listed at the bottom of the `InstallationTemplate.html` file that is generated by OEDA. Depending on your engineered system and configuration, OEDA generates all or some of the following files:

- `databasemachine.xml`
- `CustomerName-rackname.xml`

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

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

7. Review the `InstallationTemplate.html` file to check the entire configuration and verify all information was entered correctly.
- Shortly before your engineered system arrives, or is scheduled to be configured, validate the network configuration, as directed by Oracle. See [Verifying the Network Configuration Prior to Configuring the Rack](#).
 - After your engineered system arrives, the configuration files are copied to a database server, and the validation and installation is completed. See [#unique_99](#).

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

Considerations and Requirement for Using OEDA

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

- Recovery Appliance ships with the Oracle Linux operating system installed on the servers.
- The rack prefix is used to generate host names for all components. For example, if you specify `dbm0` for the rack prefix:
 - The database server host names are like `dbm0db01`
 - The storage server host names are like `dbm0ce101`
 - The host names should be 16 characters or less, because other names are created by combining this host name with certain designations (for cell disk and catalog) and have 32 character limitations that are exceeded when host names are larger than 16 characters.
 - The InfiniBand Network Fabric switch names are like `dbm0sw-iba1`
 - The RoCE Network Fabric switch names are like `dbm0sw-rocea1`

Note:

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

If you have multiple racks, then, each rack prefix should be a unique value that identifies the rack. Oracle recommends using `dbm01` for the first rack, `dbm02` for the second, `dbm03` for the third, and so on.

- If you configure Recovery Appliance with storage management using Oracle Automatic Storage Management (Oracle ASM), the backup method and redundancy protection levels

you use determine how to size the Oracle ASM disk groups that are created during installation.

- If you configure internal backups, then the database backups are created on disk in the Fast Recovery Area in the RECO disk group for the database. The Fast Recovery Area also contains Archived Redo Logs and Flashback Log Files. The division of disk space between the DATA disk group and the RECO disk group is 40% and 60%, respectively.
- If you configure external backups, then database backups are created on disk or tape media that is external to currently deployed Recovery Appliance. The Fast Recovery Area contains only objects such as Archived Redo Logs and Flashback Log Files.
- A valid time zone name is required when installing Recovery Appliance. Time zone values provided for Recovery Appliance and Oracle Linux comes from the Time Zone Database. You must specify a valid time zone name for each server using the `TZ` environment variable. The variable value is of the form *Area/Location*. For example, *America/New_York* is a valid `TZ` value; `EST`, `EDT`, and `UTC-4` are invalid values.

Related Topics

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

Getting Started with the OEDA Browser-based User Interface

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



Note:

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

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



Note:

The OEDA Web-based interface is supported only on Chrome and Firefox browsers.

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

1. Download the latest OEDA release from [My Oracle Support note 888828.1](#) and copy the archive to the system where you want to run the OEDA application server.
2. On the system where you want to run the OEDA application server, extract the contents of the downloaded compressed file. The extracted contents is written to a directory based on the operating system, such as `linux-x64`, `macosx-x64`, or `windows-i586`. This is referred to as the `OEDA_HOME` directory.
3. Launch the OEDA application server (located in the `OEDA_HOME` directory).

- On Linux, OSX, and UNIX run:

```
installOedaServer.sh [-p port-number] [-g]
```

- On Microsoft Windows run:

```
installOedaServer.cmd [-p port-number] [-g]
```

In the command:

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

If not specified, the default port number is 7072. However, you can specify any unused port number. It is not recommended to use port numbers less than 1024.

- `-g` — optionally specifies that the OEDA application server accepts connections on all available network interfaces.

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

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

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

4. Access the OEDA Web-based user interface.

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

```
http://host-name:port-number/oeda
```

In the URL:

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

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

Using the Browser-based Version of Oracle Exadata Deployment Assistant

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

Have the following available before running the deployment assistant:

- IP addresses
- Machine names

- DNS and NTP information

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

- **About:** View version information for OEDA.
- **New:** Start a new configuration.
- **Import:** Import an existing configuration file. The file format is XML.
- **Save:** Save the current configuration to an XML file.
- **Diagnostics:** Creates a ZIP file that can be used to troubleshoot and diagnose issues with the OEDA Web interface.
- **Preferences:** Specify your preferences for how the Web interface behaves.
- **Validate (OEDA Step 1):** Check the current configuration for completeness and correctness.
- **CheckIp:** Checks the IP addresses.
- **Installation Template:** Generates the Installation Template report, which summarizes the current configuration.
- **OEDA Keyboard Shortcut Keys:** Describes keyboard shortcuts for navigating the OEDA Web interface.

The following list describes the configuration flow in the OEDA Web interface. Use the **Help** menu option to display details about the options and controls provided in the OEDA Web interface.

1. Specify the hardware for the deployment on the **Select Hardware** page. This page also captures customer information and various rack-level configuration options.
2. Enter the administration network and private network information on the **Rack Networks** page.
3. Specify the configuration of database server operating system users and groups on the **Users** page. You can choose between Default, Role Separated, or a Custom configuration.
4. Define the clusters on the **Clusters** page.
Each cluster definition specifies the database and storage servers that are associated with the cluster.
5. Define the ASM disk group layout on the **Diskgroups** page.
6. Specify the Oracle Database home information for each Oracle home on the **Create Database Home** page.
7. Specify databases and their attributes on the **Create Database** page.

 Note:

To create databases on Exadata, Oracle recommends using Oracle Exadata Deployment Assistant (OEDA) or the OEDA Command Line Utility (OEDACLI), as both utilities implement the latest Exadata best-practice configuration settings for your Exadata system configuration. The Oracle Database Configuration Assistant (DBCA) is also available in the Oracle Database software bundle. However, DBCA is not recommended for new database creation on Exadata because it does not integrate with the OEDA XML configuration file and may not implement the recommended Exadata-specific configuration settings.

8. Enter the client and backup network information on the **Cluster Networks** page.
If an ingest network or replication network is needed, specify their information on the **Cluster Networks** page.
9. Configure the email alerts, SNMP alerts, Auto Service Request, Oracle Config Manager, and Enterprise Manager alerting on the **Alerting** page.
10. Answer the configuration questionnaire and provide additional information about the deployment in the text field on the **Comments** page. Information entered in this field appears at the bottom of the `InstallationTemplate.html` file.
11. Click **Save and Download** to create the configuration files. You are prompted to choose a location in which to save a ZIP file to your local computer.
12. Copy the ZIP file to a directory on the first database server in Oracle Exadata Rack. Use a subdirectory of `/u01`, such as the `/u01/oeda` directory. These files are used by the field engineer to configure the Oracle Exadata Rack.

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:

- [#unique_104](#)
- [Expanding a Recovery Appliance Rack with Additional Storage Servers](#)
- [Configuring a Recovery Appliance Rack](#)
- [Installing the Recovery Appliance Software](#)

6

Unpacking and Installing ZDLRA Rack

Oracle Rack Cabinet 1242 is the rack cabinet used for all ZDLRA Rack models since Oracle Exadata X7.

Use the following resources for information on unpacking and installing ZDLRA Rack:

- For an outline of the recommended procedures, see the *Oracle Rack Cabinet 1242 Getting Started Guide* at https://docs.oracle.com/cd/E85660_01/html/E85665/index.html.
- For additional details, consult the *Oracle Rack Cabinet 1242 User's Guide* at <https://docs.oracle.com/en/servers/options/rack-cabinet-1242/user-guide/index.html>.

Expanding a Recovery Appliance Rack with Additional Storage Servers

The Oracle Zero Data Loss Recovery Appliance racks have a flexible configuration that allows a minimum of three storage servers, and can be extended based on version.

- **RA23:** a total of 17 storage servers. Usable capacity normal redundancy, 92 TB per storage server, or minimum (3 storage servers) 274 TB and full rack total 1,564 TB. The usable capacity under high redundancy is 60 TB per storage server.
- **RA21:** a total of 18 storage servers. Usable capacity normal redundancy 70 TB per storage server, or minimum (3 storage servers) 207 TB and full rack total 1,260 TB. The usable capacity under high redundancy is 44 TB per storage server.
- **X8M, and X8 through X5:** a total of 18 storage servers.
- **X4:** a total of 14 storage servers.

All Recovery Appliance racks have two compute servers, regardless of the number of storage servers.



Note:

Older storage servers cannot be removed and replaced with newer storage servers while keeping existing Recovery Appliance backups online. If such an exchange is done, a re-image of the Recovery Appliance is required.

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

- [Preparing to Extend Recovery Appliance](#)
- [Precautions and Guidelines](#)
- [Adding Storage Servers to a Rack](#)
- [Making New Storage Available](#)



See Also:

- [In-Rack Cabling Tables for Recovery Appliance RA23](#)
- [In-Rack Cabling Tables for Recovery Appliance RA21](#)
- [In-Rack Cabling Tables for Recovery Appliance X8M](#)
- [In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5](#)
- [In-Rack Cabling Tables for Recovery Appliance X4](#)

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 RDMA Network Fabric Cable Path Lengths](#)
- [Obtaining the Current Configuration Information](#)
- [Generating the Network Configuration Files](#)
- [Moving the Audit and Diagnostic Files](#)
- [Synchronizing Release and Patch Levels](#)
- [Checking the Recovery Appliance Rack](#)

Estimating RDMA Network Fabric Cable Path Lengths

Ensure that you have the required number and lengths of the RDMA Network Fabric 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 RDMA Network Fabric cable under the floor, consider the height of the raised floor when calculating the cable path length.

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.

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 Oracle Exadata Deployment Assistant](#)
- [Installing the Recovery Appliance Software](#)

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.

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.

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](#)
- [InfiniBand Cable Precautions](#)
- [InfiniBand Cable Bundling Guidelines](#)
- [Cable Management Arm Guidelines](#)

General Safety Precautions

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

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

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.

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.

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.

Adding Storage Servers to a Rack

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

- [Adding Storage Servers to a Rack](#)
- [Checking the Recovery Appliance Rack](#)
- [Unpacking and Labeling the New Storage Servers](#)
- [Recovery Appliance RA23, RA21, X8M, and X8 through X5 Population Rules](#)
- [Removing the Doors](#)
- [Adding New Servers](#)
- [Cabling the Storage Servers](#)
- [Closing the Rack](#)



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.

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.

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"](#) for information about cable labels

[Figure 7-1](#) 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.

Recovery Appliance RA23, RA21, X8M, and X8 through X5 Population Rules

- RA23
- RA21


RA23

When adding storage servers to a Recovery Appliance RA23 follow the population rules listed in [Table 7-1](#) and depicted in [Figure 7-2](#). This assumes that the three storage servers of the base rack, minimum configuration are already in place at RU02 through RU06.

Table 7-1 Recovery Appliance RA23 Rack Population Rules

Install Order	Rack Units	Rule
1	RU08 to RU12	Install storage servers 4 through 6 from first available rack slot above existing storage servers at bottom and populate upward.
2	RU29 to RU39	Install storage servers 7 through 12 starting at RU29 and populate upward until RU39.
3	RU27 to RU23	Install storage servers 13 through 15 starting at RU27 and populate downward until RU23.
4	RU18	Install storage server 16 at RU18 between existing storage server and lower RoCE network switch.
5	RU41	Install storage server 17 at RU41, last slot at top of rack

Figure 7-2 RA23 Rack Layout and Population Rules



 See Also:

- [In-Rack Cabling Tables for Recovery Appliance RA23](#)

RA21

When adding servers to a Recovery Appliance RA21, X8M, or X8 through X5 rack, follow the population rules listed in [Table 7-2](#). This assumes that the three storage servers of the base rack are already in place.

Table 7-2 Recovery Appliance RA21, X8M, and X8 through 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 RA21](#)
- [In-Rack Cabling Tables for Recovery Appliance X8M](#)
- [In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5](#)

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 7-3 Removing the Rack Doors



Description of the callouts in [Figure 7-3](#):

- 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 7-4 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.

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 X8, X7, X6, and X5 Rack Layout](#)" or "[Recovery Appliance X4 Rack Layout](#)", as applicable.

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

Figure 7-5 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 7-6 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 to 4 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](#)" for the designated location of each server.
- For Recovery Appliance X6 or X5 racks, load the servers following the population rules. See "[Recovery Appliance RA23, RA21, X8M, and X8 through X5 Population Rules](#)".

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




Figure 7-7 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](#)".

 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

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](#)
- "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

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](#)"

- If you added the server to increase storage on an existing system, proceed to "[Making New Storage Available](#)"

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.

On install by default, a DELTA disk group name is created and the storage location name, which is often identical to the disk group name (DELTA, DELTA1, etc.). However the disk group name and storage location name are two different things. Storage is expanded using the `UPDATE_STORAGE_LOCATION` command.

Note:

The default storage location and the ASM disk group are both called *DELTA*, but are not the same thing.

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.)
- The root password of the new storage cell must still be the factory default.

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. If three (3) or more cells of different size are being added, then a new disk group DELTA (n) is created, a rebalance is initiated, and expect the rebalance to take more time. The griddisks from the cells are automatically added to the ASM diskgroup.

```
# racli add storage_cells --cells=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 --cells=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](#) 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

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](#)
- [Installing the Tape Hardware](#)
- [Verifying the Network Configuration Prior to Configuring the Rack](#)
- [Configuring the RDMA Network Fabric Switch](#)
- [Configuring the Cisco Ethernet Switch](#)
- [Configuring the Power Distribution Units](#)
- [Checking the Health of the Compute Servers](#)
- [Checking the Health of the Storage Servers](#)
- [Verifying the RoCE Network Fabric Configuration](#)
- [Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack](#)



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.

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.

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.

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.

Installing ASR Manager

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

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

Verifying the Network Configuration Prior to Configuring the Rack

Use the `checkip.sh` script to ensure there are no IP address conflicts between the existing network and your new ZDLRA Rack.

The `checkip.sh` script performs a pre-installation check to verify that the IP addresses and host names that you specified in Oracle Exadata Deployment Assistant (OEDA) are defined in

the DNS, that the NTP servers and gateways are available, and that private addresses are not pingable. Running this script before the hardware arrives help to avoid additional delays that would be caused by misconfigured network services, such as Domain Name System (DNS) and NTP.

The `checkip.sh` script is created in a format that matches the operating system of the client on which you ran OEDA. Because this script is run before the engineered system rack has arrived, you typically do not run this script on an engineered system server, but on a client. The client must have access to the same network where the engineered system will be deployed. The script is also available in the ZIP file generated by OEDA.

1. On the client where OEDA was run, copy the `checkip.sh` script generated by OEDA and the XML file `CustomerName_hostname.xml` to the same directory (one directory level up) as the OEDA `config.sh` script.
2. Run the `checkip.sh` script on the client machine or existing server.

Use a command similar to the following, where `configuration_file` is the name of the configuration generated by the Oracle Exadata Deployment Assistant for the rack being installed.

```
# ./checkip.sh -cf configuration_file
```

If the command is run from a Microsoft Windows machine, then the command is `checkip.cmd`.

If this engineered system rack is an addition for an existing installation, then run the `checkip.sh` script from an existing engineered system server. This enables the script to identify in-use IP addresses in the fabric. Not identifying existing IP addresses may cause IP collisions after installation of the new engineered system rack. To create a `checkip.sh` that can run on an existing server, you must run OEDA on a server or client that uses the same operating system as the existing engineered system server. OEDA supports IPv6 addresses.

The output from the script is a file that contains status messages such as `GOOD` or `ERROR`.

If there are conflicts that you are unable to resolve, then work with your assigned Oracle representative to correct the problems.

Configuring the RDMA Network Fabric Switch

You must perform an initial configuration of the RDMA Network Fabric switch.

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](#)
- [Setting the Time on an InfiniBand Switch](#)
- [Checking the Health of an InfiniBand Switch](#)

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
.
.
.
```

8. 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
```

9. 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
.
.
.
```

10. If any of the values are wrong, repeat the `set pendingipparameter` command, and then the `commitpending=true` command.

Setting the Time on an InfiniBand Switch

To set the time on an InfiniBand switch:

1. 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 [Using Oracle Exadata Deployment Assistant](#).

2. 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
```

3. 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](#).

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
  .
  .
  .
```

Setting the Serial Number on a Spine Switch

In a multirack 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
```

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.

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 22.b.

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 Master Manager 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 and 3 until the spine switch is the Subnet Manager Master.
5. Enable Subnet Manager on the leaf switches that were disabled during this procedure.

 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.

Configuring the Cisco Nexus 9336C-FX2 Switch

The RoCE Network Fabric switch supplied with the engineered system rack is minimally configured during installation.

During initial system configuration, you can reset and configure the switch.

1. Connect from the RoCE Network Fabric switch serial console to a laptop or similar device using the available RJ45 cable.
2. Ensure the terminal session is recorded on the laptop by logging the output.
The output can be used as a reference that the switch has been configured correctly.
3. Power on the switch.
4. Log in as the `admin` user.

```
User Access Verification
dbm0sw-rocea0 login: admin
Password: *****
```


 Note:

If you do not have the password for the `admin` user, then contact Oracle Support Services.

5. Erase the existing configuration.

```
dbm0sw-rocea0# write erase
```

```
Warning: This command will erase the startup-configuration.
```

```
Do you wish to proceed anyway? (y/n) [n] y
```

6. Restart the system so you can perform the automated setup.

```
dbm0sw-rocea0# reload
```

```
This command will reboot the system. (y/n)? [n] y
```

```
2017 Aug 31 01:09:00 dbm0sw-rocea0 %$ VDC-1 %$ %PLATFORM-2-  
PFM_SYSTEM_RESET: Manual system restart from Command Line Interface
```

```
CISCO SWITCH Ver7.59  
Device detected on 0:1:2 after 0 msecs  
...
```

7. Switch to normal setup and, when asked if you want to enforce secure password standard, enter `no`, then enter a new password for the `admin` user.

```
Running S93thirdparty-script...
```

```
Populating conf files for hybrid sysmgr ...  
Starting hybrid sysmgr ...  
inserting /isan/lib/modules/klm_cisco_nb.o ... done
```

```
Abort Auto Provisioning and continue with normal setup ? (yes/no) [n]: yes
```

```
---- System Admin Account Setup ----
```

```
Do you want to enforce secure password standard (yes/no) [y]: no
```

```
Enter the password for "admin":  
Confirm the password for "admin":
```

8. When the Basic System Configuration Dialog appears, choose to enter the basic configuration dialog.

```
---- Basic System Configuration Dialog VDC: 1 ----
```

```
This setup utility will guide you through the basic configuration of  
the system. Setup configures only enough connectivity for management  
of the system.
```

Please register Cisco Nexus9000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus9000 devices must be registered to receive entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

9. In the basic configuration, you can use the default inputs until asked to enter the switch name.

In this example, the switch has a name of `test123sw-rocea0`.

```
Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : test123sw-rocea0
```

10. Respond **yes** when asked to configure Out-of-band management configuration, and specify appropriate network addresses when prompted.

```
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
yes
  Mgmt0 IPv4 address : 100.104.10.21
  Mgmt0 IPv4 netmask : 255.255.248.0
  Configure the default gateway? (yes/no) [y]:
  IPv4 address of the default gateway : 100.104.10.1
```

11. Respond **yes** when asked to configure advanced IP options.

```
Configure advanced IP options? (yes/no) [n]: yes
```

12. Respond **yes** when asked to configure static route (this can be changed later).

```
Configure static route? (yes/no) [n]: yes
```

13. Enter the destination prefix and mask, and other values as prompted.

```
Destination prefix : 10.100.100.0
Destination prefix mask : 255.255.255.0
Next hop IPv4 address : 10.100.100.1
```

14. Configure the DNS IPv4 addresses.

```
Configure the DNS IPv4 address? (yes/no) [n]: yes
DNS IP address: 10.100.100.2
```

15. Skip configuring the default domain name (this will be configured later).

```
Configure the default domain name? (yes/no) [n]: no
```

16. Accept the default responses until asked to configure SSH and the NTP server.

```
Enable the telnet service? (yes/no) [n]: no
Enable the ssh service? (yes/no) [y]: yes
  Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
  Number of rsa key bits <1024-2048> [1024]: 1024

Configure the NTP server? (yes/no) [n]: yes
  NTP server IPv4 address : 10.100.100.3
```

17. Accept the default responses until asked to specify the CoPP system profile. Enter `strict`.

```
Configure default interface layer (L3/L2) [L2]:
Configure default switchport interface state (shut/noshut) [noshut]:
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]:
strict
```

18. After reviewing the configuration, save the configuration.

```
The following configuration will be applied:
no password strength-check
switchname test123sw-rocea0
ip route 100.104.8.0 255.255.248.0 100.104.10.1
vrf context management
ip route 0.0.0.0/0 100.104.10.1
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
ntp server 100.104.10.1
system default switchport
no system default switchport shutdown
copp profile strict
interface mgmt0
ip address 100.104.10.21 255.255.248.0
no shutdown

Would you like to edit the configuration? (yes/no) [n]:

Use this configuration and save it? (yes/no) [y]: yes

[#####] 100%
Copy complete.
```

19. Enable the scp server feature on the switch.

```
test123sw-rocea0# feature scp-server
```

20. Save the running configuration to flash.

```
test123sw-rocea0# copy running-config startup-config
[#####] 100%
Copy complete.
```

21. Apply the golden configuration on the switch.

- Starting with Oracle Exadata System Software release 20.1.0, use the procedure described in Applying Golden Configuration Settings on RoCE Network Fabric Switches, in *Oracle Exadata Database Machine Maintenance Guide*.
 - Otherwise, use the following procedure to apply the golden configuration on the switch:
- a. Delete the configuration file on the switch for the target configuration.

 Note:

If you do not remove the file you are replacing, then when you attempt to overwrite the file you will get a 'permission denied' error.

Log in to the switch, enter configuration mode, then run a command similar to the following:

```
test123sw-rocea0# delete bootflash:roce_leaf_switch.cfg
Do you want to delete "/roce_leaf_switch.cfg" ? (yes/no/abort) [y] y
test123sw-rocea0#
```

- b. Log in to a server that has SSH access to the switch, and contains the latest RDMA Network Fabric patch ZIP file.

To find the available RDMA Network Fabric patches, search for '**RDMA network switch**' in My Oracle Support document 888828.1. Download and use the latest patch for your Oracle Exadata System Software release.

- c. Unzip the RDMA Network Fabric patch ZIP file and change directories to the location of the patchmgr utility.
- d. Locate the golden configuration files in the RDMA Network Fabric patch bundle.

The files are located within the `roce_switch_templates` directory.

The golden configuration files are as follows:

- Single rack leaf (leaf): `roce_leaf_switch.cfg`
- Multi-rack spine (mspine): `roce_spine_switch_multi.cfg`
- Multi-rack leaf (mleaf): `roce_leaf_switch_multi.cfg`
- Single rack leaf with Secure Fabric support (sfleaf): `roce_sf_leaf_switch.cfg`
- Multi-rack leaf with Secure Fabric support (msfleaf):
`roce_sf_leaf_switch_multi.cfg`
- Single rack leaf configured with 23 host ports (leaf23):
`roce_leaf_switch_23hosts.cfg`
- Multi-rack leaf configured with 23 host ports (mleaf23):
`roce_leaf_switch_23hosts_multi.cfg`
- Multi-rack leaf configured with 14 inter-switch links (mleaf_ul4):
`roce_leaf_switch_14uplinks_multi.cfg`
- Multi-rack leaf configured with 14 inter-switch links and with Secure Fabric support (msfleaf_ul4): `roce_sf_leaf_switch_14uplinks_multi.cfg`
- Multi-rack leaf configured with 23 host ports and 13 inter-switch links (mleaf23_ul3): `roce_leaf_switch_23hosts_13uplinks_multi.cfg`

- e. Copy the golden configuration file to the switch.

In the following example, 100.104.10.21 represents the IP address of the switch you are configuring.

```
# scp roce_leaf_switch.cfg admin@100.104.10.21:/
User Access Verification
Password:
roce_leaf_switch.cfg 100% 23KB 23.5KB/s 00:00
```

- f. Apply the golden configuration file on the switch.

Use the `run-script` command while connected directly to the switch.

```
test123sw-rocea0# run-script bootflash:roce_leaf_switch.cfg | grep
'none'
```

 Note:

This command may take up to 1-2 minutes on a single-rack switch and up to 3-4 minutes on a multi-rack switch.

- g. Verify the switch configuration.

Use the `patchmgr` utility on the server that has SSH access to the switch, and contains the latest RDMA Network Fabric patch bundle.

In the following command, `roceswitch.lst` is a file that contains the switch host name or IP address.

```
# ./patchmgr --roceswitches roceswitch.lst --verify-config
```

22. Backup up the switch configuration.

Follow the steps in [Backing Up Settings on the ROCE Switch](#), in *Oracle Exadata Database Machine Maintenance Guide*.

23. Optional: Set the clock, using the same procedure as in [Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch](#).

Related Topics

- [Exadata Database Machine and Exadata Storage Server Supported Versions \(My Oracle Support Doc ID 888828.1\)](#)

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.

Configuring the Cisco Catalyst 4948 Ethernet Switch

The Cisco Catalyst 4948 Ethernet switch supplied with ZDLRA Rack is minimally configured during installation.

The minimal configuration disables IP routing, and sets the following:

- Host name
- IP address setup
- Subnet mask
- Default gateway
- Domain name
- Name server
- NTP server
- Time
- Time zone

Before configuring the switch, note the following:

- The Cisco Ethernet switch should not be connected until the running configuration has been verified, and any necessary changes have been made by the network administrator.
- The Cisco Ethernet switch should not be connected to the customer network until the IP addresses on all components have been configured in ZDLRA Rack. This is to prevent any duplicate IP address conflicts which are possible due to the default addresses set in the components when shipped.

Note that the Cisco 4948E-F switch supports multiple uplinks to the customer network by utilizing ports 49 - 52. This is a more complicated switch setup due to the redundant connectivity, and should be performed by the customer's network administrator.

The following procedure describes how to configure the Cisco Ethernet switch. Configuration should be done with the network administrator.

1. Connect a serial cable from the Cisco switch console to a laptop or similar device. An Oracle supplied rollover cable is pre-installed on the Cisco serial console port. Obtain the appropriate adapter and connect it at the end of the rollover cable. An Oracle P/N 530-3100 RJ45-DB9 adapter as used on ILOM ports will also work, connected at the end of the network cable.
2. Ensure the terminal session is recorded on the laptop by logging the output. The output can be used as a reference that the switch has been configured correctly. The default serial port speed is 9600 baud, 8 bits, no parity, 1 stop bit, and no handshake.

```
Switch con0 is now available
Press RETURN to get started.
```

3. Change to the enable mode.

```
Switch> enable
Password: *****
Switch#
```

 Note:

If you do not have the password, then contact Oracle Support Services.

4. Check the current version on the switch.

```
Switch# show version
Cisco IOS Software, Catalyst 4500 L3 Switch Software (cat4500e-
IPBASEK9-M), Version 15.2(3)E2, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2014 by Cisco Systems, Inc.
Compiled Tue 11-Mar-14 18:28 by prod_rel_team

ROM: 12.2(44r)SG12
zdlralsw-ip uptime is 1 minute
System returned to ROM by reload
System image file is "bootflash:cat4500e-ipbasek9-mz.152-3.E2.bin"
Hobgoblin Revision 22, Fortoointe Revision 1.40
...

Configuration register is 0x2102

Switch#
```

The version of the Cisco 4948E-F switch firmware purchased and shipped by Oracle with Recovery Appliance X6 is IPBASEK9-MZ, which includes telnet and ssh support. Currently the full release version string is cat4500e-ipbasek9-mz.152-3.E2.bin.

5. Configure the network for a single VLAN. The following example assumes you are using IPv4 addressing.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface vlan 1
Switch(config-if)# ip address 10.7.7.34 255.255.255.0
Switch(config-if)# end
Switch# *Sep 15 14:12:06.309:%SYS-5-CONFIG_I:Configured from console by
console
Switch# write memory
Building configuration...
Compressed configuration from 2474 bytes to 1066 bytes [OK ]
```

6. If IP routing is required on the switch, then leave the IP routing setting as the default, and configure the default gateway. Replace 10.7.7.1 with the IP address of the gateway for the installation:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip route 0.0.0.0 0.0.0.0 10.7.7.1
Switch(config)# end
Switch# *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:

The system host name is used as the prompt 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:

```
ralsw-ip#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ralsw-ip(config)#ip domain-name example.com
ralsw-ip(config)#ip name-server 10.7.7.3
ralsw-ip(config)#ip name-server 198.51.100.5
ralsw-ip(config)#ip name-server 10.8.160.1
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 ]
```

If you do not have DNS service available, you must still set the domain-name so that you can configure the SSH keys.

9. (Optional) Set the password.

```
ralsw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ralsw-ip(config)# enable password password
ralsw-ip(config)# enable secret password
ralsw-ip(config)# end
ralsw-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. Verify telnet access is disabled. Telnet is not secure, and should not be enabled unless there is a compelling reason. To enable telnet, set a password. To disable it, remove the password.

```
ralsw-ip#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ralsw-ip(config)#line vty 0 15
ralsw-ip(config)#login
% Login disabled on line 1, until 'password' is set
% Login disabled on line 2, until 'password' is set
...
% Login disabled on line 16, until 'password' is set
ralsw-ip(config)#end
```

If the login command returns output as shown above, then telnet access has been disabled. If instead you get a prompt, then telnet access is not yet disabled so should be disabled now.

```
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 ]
```


11. To configure a secure shell (SSH) on the Ethernet switch:

```

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]

```

12. Set the clock and time zone. The switch keeps internal time in Coordinated Universal Time (UTC) format.

- To use UTC, use the following command:

```
no clock timezone global configuration
```

- To use a time zone, use the following command:

```
clock timezone zone hours-offset [minutes-offset]
```

In the preceding command, *zone* is the time zone to display when standard time is in effect, *hours-offset* is the hours offset from UTC, and *minutes-offset* is the minutes offset from UTC.

- Daylight savings time (or summer time) is disabled by default. To set summer time hours, use the following command:

```
clock summer-time zone recurring [week day monthhh:mm week day month \
hh:mm[offset]]
```

In the preceding command, *zone* is the time zone to be displayed when summer time is in effect (EDT, for example), *week* is the week of the month (1 to 5 or last), *day* is the day of the week (Sunday, Monday, ...), *month* is the month (January, February, ...), *hh:mm* is the hours and minutes in 24-hour format, and *offset* is the number of minutes to add during summer time. The default offset is 60 minutes.

- To manually set the clock to any time use the following command, where the time specified is relative to the configured time zone:

```
clock set hh:mm:ss month day year
```

In the preceding command, *hh:mm:ss* is the time in 24-hour format, *day* is the day by date in the month, *month* is the name of the month, and *year* is the 4-digit year.

The ordering of commands is important when setting the local time and time zone. For example, to set the local time to US Eastern time:

```
ralsw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ralsw-ip(config)# clock timezone EST -5
ralsw-ip(config)# clock summer-time EDT recurring
ralsw-ip(config)# end
ralsw-ip# clock set 21:00:00 August 09 2018
ralsw-ip# write memory
Building configuration...
Compressed configuration from 3784 bytes to 1465 bytes [OK ]
ralsw-ip# show clock
21:00:06.643 EST Mon Aug 9 2018
```

13. After setting the local time zone, you can configure up to two NTP servers. Replace the IP addresses used in this example with the values for the installation:

```
ralsw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ralsw-ip(config)# ntp server 10.7.7.32 prefer
ralsw-ip(config)# ntp server 198.51.100.19
ralsw-ip(config)# end
*Sep 15 14:51:08.665:%SYS-5-CONFIG_I:Configured from console by console
ralsw-ip# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes [OK ]
ralsw-ip# show ntp status
<output will vary per network>
.
ralsw-ip# show clock
21:00:23.175 EST Mon Aug 9 2018
```

The NTP server is synchronized to local time when you connect the Cisco switch to the network and it has access to NTP.

Symbols that precede the `show clock` display indicate that the time is the following:

- * Not authoritative
- . Authoritative, but NTP is not synchronized.
- Authoritative (blank space).

14. Verify the Ethernet configuration using the following command:

```
ralsw-ip# show running-config
Building configuration...
Current configuration : 3923 bytes
```

```

!
version 15.2
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
.
.
.

```

 Note:

If any setting is incorrect, then repeat the appropriate step. To erase a setting, enter `no` in front of the same command. For example, to erase the default gateway, use the following commands:

```

ralsw-ip#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
ralsw-ip(config)# no ip default-gateway 10.7.7.1
ralsw-ip(config)# end
ralsw-ip#
*Sep 15 14:13:26.013: %SYS-5-CONFIG_I: Configured from console by
console
ralsw-ip(config)# write memory
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes[OK]

```

15. Save the current configuration.

```

ralsw-ip#copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
Compressed configuration from 2654 bytes to 1189 bytes[OK]

```

16. Exit from the session using the following command:

```

ralsw-ip# exit

ralsw-ip con0 is now available

Press RETURN to get started.

```

17. Disconnect the cable from the Cisco console.

The Cisco switch must not be connected to the management network at this stage. The switch will be connected later after Oracle has configured the systems with the necessary IP addresses and you have worked with the field service engineer to make any additional changes necessary for connecting to the network.

18. To check the Cisco switch, attach a laptop computer to port 48, and ping the IP address of the internal management network to check the configuration.

Do not connect the switch to the management network.

Configuring the Cisco Nexus 93108-1G or 9348 Ethernet Switch

The Cisco Nexus 93108-1G or 9348 Ethernet switch supplied with ZDLRA Rack is minimally configured during installation.

Note that the Cisco Nexus 93108-1G or 9348 switch supports multiple uplinks to the customer network by utilizing the QSFP+ ports. This is a more complicated switch setup due to the redundant connectivity, and should be performed by the customer's network administrator.

Whether you are configuring the switch for the first time, or configuring a replacement switch, use the following procedures:

Performing the Initial Switch Configuration for the Cisco Nexus 93108-1G or 9348 Ethernet Switch

During the initial configuration, you reset the switch and use the Basic System Configuration Dialog to configure the switch.

Before configuring the switch, note the following:

- The Cisco Ethernet switch should not be connected until the running configuration has been verified, and any necessary changes have been made by the network administrator.
- The Cisco Ethernet switch should not be connected to the customer network until the IP addresses on all components have been configured in ZDLRA Rack. This is to prevent any duplicate IP address conflicts which are possible due to the default addresses set in the components when shipped.

Configuration should be done with the network administrator.

1. Connect from the Cisco switch serial console to a laptop or similar device using the available RJ45 cable.
2. Ensure the terminal session is recorded on the laptop by logging the output.

The output can be used as a reference that the switch has been configured correctly. The default serial port speed is 9600 baud, 8 bits, no parity, 1 stop bit, and no handshake.

3. Power on the switch.
4. Log in as the `admin` user.

```
User Access Verification
exadatax7-adm0 login: admin
Password: *****
```

Note:

If you do not have the password for the `admin` user, then contact Oracle Support Services.

5. Erase the existing configuration.

```
exadatax7-adm0# write erase
```

Warning: This command will erase the startup-configuration.

Do you wish to proceed anyway? (y/n) [n] **y**

6. Restart the system so you can perform the automated setup.

```
exadatax7-adm0# reload
```

This command will reboot the system. (y/n)? [n] **y**

```
2017 Aug 31 01:09:00 exadatax7-adm0 %$ VDC-1 %$ %PLATFORM-2-
PFM_SYSTEM_RESET: Manual system restart from Command Line Interface
```

```
CISCO SWITCH Ver7.59
Device detected on 0:1:2 after 0 msec
...
```

7. Switch to normal setup and, when asked if you want to enforce secure password standard, enter no, then enter a new password for the admin user.

```
Running S93thirdparty-script...
```

```
Populating conf files for hybrid sysmgr ...
Starting hybrid sysmgr ...
inserting /isan/lib/modules/klm_cisco_nb.o ... done
```

Abort Auto Provisioning and continue with normal setup ? (yes/no) [n]: **yes**

```
---- System Admin Account Setup ----
```

Do you want to enforce secure password standard (yes/no) [y]: **no**

```
Enter the password for "admin":
Confirm the password for "admin":
```

8. When the Basic System Configuration Dialog appears, choose to enter the basic configuration dialog.

```
---- Basic System Configuration Dialog VDC: 1 ----
```

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Please register Cisco Nexus9000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus9000 devices must be registered to receive entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): **yes**

9. In the basic configuration, you can use the default inputs until asked to enter the switch name.

In this example, the switch has a name of `test123sw-adm0`.

```
Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : test123sw-adm0
```

10. Respond `no` when asked to configure Out-of-band management configuration.

```
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
no
```

11. Respond `yes` when asked to configure advanced IP options.

```
Configure advanced IP options? (yes/no) [n]: yes
```

12. Respond `no` when asked to configure static route (this will be configured later).

```
Configure static route? (yes/no) [n]: no
```

13. Enter the destination prefix and mask, and other values as prompted.

```
Destination prefix : 10.100.100.0
Destination prefix mask : 255.255.255.0
Next hop IPv4 address : 10.100.100.1
```

14. Skip configuring the DNS IPv4 addresses (this will be configured later).

```
Configure the DNS IPv4 address? (yes/no) [n]: no
```

15. Skip configuring the default domain name (this will be configured later).

```
Configure the default domain name? (yes/no) [n]: no
```

16. Accept the default responses until asked to configure SSH and the NTP server.

```
Enable the telnet service? (yes/no) [n]: no
Enable the ssh service? (yes/no) [y]: yes
Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
Number of rsa key bits <1024-2048> [1024]: 1024

Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : 10.100.100.3
```

17. Accept the default responses until asked to specify the CoPP system profile. Enter `lenient`.

```
Configure default interface layer (L3/L2) [L2]:
Configure default switchport interface state (shut/noshut) [noshut]:
```

```
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]:
lenient
```

18. After reviewing the configuration, save the configuration.

The following configuration will be applied:

```
no password strength-check
switchname test123sw-adm0
...
```

Would you like to edit the configuration? (yes/no) [n]:

Use this configuration and save it? (yes/no) [y]: **yes**

```
[#####] 100%
Copy complete.
```

19. Add the VLAN 1 IP address.

```
test123sw-adm0(config)# feature interface-vlan
test123sw-adm0(config)# interface vlan 1
test123sw-adm0(config-if)# ip address 10.100.100.110/24
test123sw-adm0(config-if)# no shutdown
test123sw-adm0(config-if)# exit
```

20. Set the spanning tree port type for ports 1-47.

```
test123sw-adm0(config)# interface E1/1-47
test123sw-adm0(config-if)# spanning-tree port type edge
test123sw-adm0(config-if)# exit
```

21. Set switchport on all 48 ports and set port 48 to a network port (instead of a host port).

```
test123sw-adm0(config)# interface E1/1-48
test123sw-adm0(config-if)# switchport
test123sw-adm0(config-if)# exit
test123sw-adm0(config)# interface E1/48
test123sw-adm0(config-if)# spanning-tree port type network
test123sw-adm0(config-if)# ip route 0.0.0.0/0 10.100.100.1
```

22. Configure the DNS information.

```
test123sw-adm0(config)# ip domain-name example.com
test123sw-adm0(config)# ip name-server 10.100.100.2
test123sw-adm0(config)# exit
```

23. Save the current configuration.

```
test123sw-adm0# copy running-config startup-config
[#####] 100%
Copy complete.
```

24. Optional: Set the clock, as described in the next topic.

Setting the Clock on the Cisco 93108-1G or 9348 Ethernet Switch

After you have performed the initial configuration, you can adjust the time used by the switch.

1. Log in as the `admin` user.
2. View the current time.

```
test123sw-adm0(config)# show clock  
20:44:52.986 UTC Thu Aug 31 2017  
Time source is NTP
```

3. Set the timezone appropriately.

```
test123sw-adm0(config)# clock timezone PST -8 0
```

4. View the modified time.

```
test123sw-adm0(config)# show clock  
12:46:22.692 PST Thu Aug 31 2017  
Time source is NTP
```

5. Save the configuration.

```
test123sw-adm0# copy running-config startup-config  
[#####] 100%  
Copy complete.
```

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

Configuring the Power Distribution Units

The power distribution units (PDUs) are configured with static IP addresses to connect to the network for monitoring.

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.

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.

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.

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_identifier = 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"](#) 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.

Checking the Health of the Storage Servers

A Recovery Appliance X5 and higher versions have 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](#)".

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](#)".

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. X5 has 96 GB, while X8 has 384 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.

Verifying the RoCE Network Fabric Configuration

This procedure describes how to verify the RoCE Network Fabric configuration.

1. Verify the proper `oracle-rdma-release` software versions are being used on the database servers.

```
[root@dbm01adm08 ~]# rpm -qa |grep oracle-rdma-release
oracle-rdma-release-0.11.0-1.el7ora.x86_64
```

The `oracle-rdma-release` software and adapter firmware versions are automatically maintained on the Recovery Appliance storage servers.

2. Check the adapter firmware versions on the database servers.

Use the `CheckHWnFWProfile` script to check firmware versions for the RDMA Network Fabric adapters.

```
# /opt/oracle.SupportTools/CheckHWnFWProfile -action list
```

3. Visually check all the RDMA Network Fabric cable connections within the rack.
The port lights should be on, and the LEDs should be on. Do not press each connector to verify connectivity.
4. Complete the steps described in My Oracle Support Doc ID 2587717.1

Related Topics

- [Verify RoCE Cabling on Oracle Exadata Database Machine X8M-2 and X8M-8 Servers \(My Oracle Support Doc ID 2587717.1\)](#)

Verifying the InfiniBand Network Fabric Network

This procedure describes how to verify the InfiniBand Network Fabric network.

1. Visually check all the RDMA Network Fabric cable connections within the rack. The port lights should be on, and the LEDs should be on. Do not press each connector to verify connectivity.
2. Log in as the `root` user on any component in the rack.
3. Verify the InfiniBand Network Fabric topology using the following commands:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify-topology [-t rack_size]
```

The following example shows the output when the network components are correct.

```
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch.....[SUCCESS ]
Are any external switches connected to each other.....[SUCCESS ]
Are any hosts connected to spine switch.....[SUCCESS ]
```



```
Check if all hosts have 2 CAs to different switches.....[SUCCESS ]
Leaf switch check:cardinality and even distribution.....[SUCCESS ]
Check if each rack has an valid internal ring.....[SUCCESS ]
```

In the preceding command, *rack_size* is the size of the rack. The *-t rack_size* option is needed if the rack is Recovery Appliance Half Rack or Recovery Appliance Quarter Rack. Use the value *halfrack* or *quarterrack*, if needed.

The following example shows the output when there is a bad RDMA Network Fabric switch to cable connection:

```
#./verify-topology
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch.....[SUCCESS ]
Are any external switches connected to each other.....[SUCCESS ]
Are any hosts connected to spine switch.....[SUCCESS ]
Check if all hosts have 2 CAs to different switches.....[SUCCESS ]
Leaf switch check:cardinality and even distribution.....[SUCCESS ]
Check if each rack has an valid internal ring.....[ERROR ]
```

Switches 0x21283a87cba0a0 0x21283a87b8a0a0 have 6 connections between them.
They should have at least 7 links between them

The following example shows the output when there is a bad RDMA Network Fabric cable on a database server:

```
#./verify-topology
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch.....[SUCCESS ]
Are any external switches connected to each other.....[SUCCESS ]
Are any hosts connected to spine switch.....[SUCCESS ]
Check if all hosts have 2 CAs to different switches.....[ERROR ]
Node db01 has 1 endpoints.(Should be 2)
Port 2 of this node is not connected to any switch
-----fattree End Point Cabling verification failed-----
Leaf switch check:cardinality and even distribution.....[ERROR ]
Internal QDR Switch 0x21283a87b8a0a0 has fewer than 4 compute nodes
It has only 3 links belonging to compute nodes [SUCCESS ]
Check if each rack has an valid internal ring.....[SUCCESS ]
```

The following example shows the output when there is a bad connection on the switch and the system:

```
#./verify-topology
[DB Machine Infiniband Cabling Topology Verification Tool ]
Is every external switch connected to every internal switch.....[SUCCESS ]
Are any external switches connected to each other.....[SUCCESS ]
Are any hosts connected to spine switch.....[SUCCESS ]
Check if all hosts have 2 CAs to different switches.....[ERROR ]

Node burxdb01 has 1 endpoints.(Should be 2)
Port 2 of this node is not connected to any switch
-----fattree End Point Cabling verification failed-----
Leaf switch check:cardinality and even distribution.....[ERROR ]
```

```
Internal QDR Switch 0x21283a87b8a0a0 has fewer than 4 compute nodes
It has only 3 links belonging to compute nodes.....[SUCCESS ]
Check if each rack has an valid internal ring.....[ERROR ]
```

```
Switches 0x21283a87cba0a0 0x21283a87b8a0a0 have 6 connections between them
They should have at least 7 links between them
```

Setting the Subnet Manager Master on Oracle Exadata Database Machine Full Rack and Oracle Exadata Database Machine Half Rack

Recovery Appliance X3-2 systems and Recovery Appliance X2-2 systems have three Sun Datacenter InfiniBand Switch 36 switches. Starting with Recovery Appliance X4-2, Recovery Appliance Systems have two Sun Datacenter InfiniBand Switch 36 switches.



Note:

This procedure does not apply to Recovery Appliance X8M racks with RoCE Network Fabric.

The switch located in rack unit 1 (U1) is referred to as the spine switch. The other two switches are referred to as the leaf switches. The location of the leaf switches is as follows:

- Recovery Appliance Two-Socket Systems (X3-2 and later): rack unit 20 (U20) and rack unit 22 (U22)
- Recovery Appliance X2-2 racks: rack unit 20 (U20) and rack unit 24 (U24)
- Recovery Appliance Eight-Socket Systems (X2-8 and later) Full Racks: Rack unit 21 (U21) and rack unit 23 (U23)

The spine switch is the Subnet Manager Master for the InfiniBand Network Fabric subnet. The Subnet Manager Master has priority 8, and can be verified using the following procedure:

1. Log in to the spine switch as the `root` user.
2. Run the `setsmpriority list` command.

The command should show that `smpriority` has a value of 8. If `smpriority` has a different value, then do the following:

- a. Use the `disablesm` command to stop the Subnet Manager.
- b. Use the `setsmpriority 8` command to set the priority to 8.
- c. Use the `enablesm` command to restart the Subnet Manager.

The leaf switches are the Standby Subnet Managers with a priority of 5. This can be verified using the preceding procedure, substituting a value of 5 in the `setsmpriority` command above.

**Note:**

Recovery Appliance Half Rack with Sun Fire X4170 Oracle Database Servers include two Sun Datacenter InfiniBand Switch 36 switches, which are set to priority 5.

To determine the Subnet Manager Master, log in as the `root` user on any InfiniBand Network Fabric switch, and run the `getmaster` command. The location of the Subnet Manager Master is displayed. The following is an example of the output from the `getmaster` command:

```
# getmaster
20100701 11:46:38 OpenSM Master on Switch : 0x0021283a8516a0a0 ports 36 Sun
DCS 36
QDR switch dm01sw-ib1.example.com enhanced port 0 lid 1 lmc 0
```

The preceding output shows the proper configuration. The Subnet Manager Master is running on spine switch `dm01sw-ib1.example.com`.

If the spine switch is not the Subnet Manager Master, then do the following procedure to set the Subnet Manager Master:

1. Use the `getmaster` command to identify the current location of the Subnet Manager Master.
2. Log in as the `root` user on the leaf switch that is the Subnet Manager Master.
3. Disable Subnet Manager on the switch. The Subnet Manager Master relocates to another switch.
4. Use the `getmaster` command to identify the current location of the Subnet Manager Master. If the spine switch is not Subnet Manager Master, then repeat steps 2 and 3 until the spine switch is the Subnet Manager Master.
5. Enable Subnet Manager on the leaf switches that were disabled during this procedure.

**Note:**

- If the InfiniBand Network Fabric network consists of four or more racks cabled together, then only the spine switches should run Subnet Manager. The leaf switches should have Subnet Manager disabled on them.
- Recovery Appliance Half Racks with Sun Fire X4170 Oracle Database Servers, and Recovery Appliance Quarter Racks have two Sun Datacenter InfiniBand Switch 36 switches, and both are set to priority 5. The master is the one with the lowest GUID.

**See Also:**

- ["Enable the Subnet Manager"](#) in *Sun Datacenter InfiniBand Switch 36 User's Guide*
- ["Disable the Subnet Manager"](#) in *Sun Datacenter InfiniBand Switch 36 User's Guide*
- *Oracle Exadata Database Machine System Overview* for hardware component information
- Cabling tables in *Oracle Exadata Database Machine System Overview*

9

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](#)
- [Software Installation Prerequisites](#)
- [Reclaiming Disk Space on the Compute Servers](#)
- [Installing the Software on Recovery Appliance](#)
- [Install Utility Syntax](#)
- [Setting Up Enterprise Manager to Administer Recovery Appliance](#)



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.

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. 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 Oracle Exadata Deployment Assistant](#) also contains this script. You can run this script from any system on the same network as Recovery Appliance. Some platforms use `install.cmd`.

- The OneCommand ZIP file that contains Oracle Exadata Deployment Assistant for use in [Using Oracle Exadata Deployment Assistant](#) 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 install appliance** : 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`.

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 Oracle Exadata Deployment Assistant](#).
- 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`.)
- The `checkip` script ran successfully with no remaining errors.
- All racks of a multirack installation are cabled together. See [Extending a Rack by Adding Another Rack](#).
- All components of all new racks are configured. See [Configuring a Recovery Appliance Rack](#).
- 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"](#).
- The Oracle tape library is installed and operational. A tape library is an optional component of Recovery Appliance.

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.

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.

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.

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, the software ZIP file, and the GI/RDBMS Gold image into subdirectory under `/u01/ra_install/linux-x64`. See "[Using Oracle Exadata Deployment Assistant](#)" 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.

- b. Extract the Recovery Appliance software from patch ZIP file into /u01/ra_install/linux-x64/WorkDir.
- c. Extract the GI/RDBMS Gold image files from the GI/RDBMS Gold image ZIP file. This substep and the substep above must occur before the ra_preinstall.pl script is run.
- d. 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

- e. 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
```

```
ra_preinstall.pl --config_xml=/u01/ra_install/linux-x64/
ExadataConfigurations/Example_Inc-ra01.xml
--oeda_dir=/u01/ra_install/linux-x64 --oracle_uid=5555 --
oinstall_gid=5555
--dba_gid=5556 --raadmin_uid=6557 --raadmin_gid=6557
```

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](#)".

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
```

Note:

You are prompted for the `raadmin` password when it gets to `--step=7`. Remote `root` and `oracle` access are removed during `--step=7`. Later this access can be re-enabled with [racli enable ssh](#).

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](#)" 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"](#).

- 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](#)
- [Changing OSB Tape Backup Passwords](#)
- [Default User Accounts for Oracle Zero Data Loss Recovery Appliance](#)

Install Utility Syntax

The `install` utility has the following syntax:

```
./install.sh -cf config_filename.xml option1 option2...
```

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 9-1](#) describes the options.

Table 9-1 Options of the Install Utility

Option	Description
-cf config_filename.xml	Identifies the XML configuration file. Required.
-h	Describes the script options.
-l	List the steps needed to install and configure the software for the configuration described in the XML file.
-r {n-N n}	Runs steps <i>n</i> through <i>N</i> , or step <i>n</i> .
-s <i>n</i>	Runs step <i>n</i> .
-u {n-N n}	Undoes steps <i>n</i> through <i>N</i> , or step <i>n</i> .

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
.
```

:

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](#)" 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.

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.

Deploying the Enterprise Manager Agents

If your Recovery Appliance is in compliance mode, you must enable SSH access using [racli enable ssh](#) before you can deploy the Enterprise Manager Agents. The process for enabling SSH access requires quorum approval. After the agents have been deployed on all the compute servers, disable SSH access with [racli disable ssh](#).

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 might indicate that the `root.sh` script could not be run due to insufficient privileges. You can 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. If the `root.sh` script was not automatically run, 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. If your Recovery Appliance was in compliance mode before starting these steps and you enabled SSH access in order to deploy the Enterprise Manager, you should disable SSH access now with [racli disable ssh](#).

Discovering the Cluster

Perform the following steps to discover the Recovery Appliance cluster.



Note:

If the Recovery Appliance is TLS-enabled, make sure the TCPS port and protocol are specified during the discovery of the cluster, the Recovery Appliance database, and the Recovery Appliance target.

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.

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.



Note:

If the Recovery Appliance is TLS-enabled, make sure the TCPS port and protocol are specified during the discovery of the cluster, the Recovery Appliance database, and the Recovery Appliance target.

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.

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. These steps apply to Recovery Appliance not in TLS mode.

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** page appears.

5. Select the **Discover new Recovery Appliance hardware components as targets** option and click on **Discover Targets**.

The Recovery Appliance **Hardware Discovery: Infiniband Discovery** page appears.

6. 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.

7. Click **Continue With Recovery Appliance Discovery**.

The **Recovery Appliance Discovery: Properties** page appears.

8. 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**.)

9. In the Recovery Appliance **Monitoring Credentials** section, specify the database user credentials that will be used to monitor the Recovery Appliance.

This can be a named user with the `RA$ADMIN` role or a named user that has the `RA$MON` role.

Click **More Details** to see detailed information about the credentials.

10. In the **Host Credentials** section, provide the credentials for a user that has permission to access the Oracle Home of the Recovery Appliance database.

11. Click **Next**.

The Recovery Appliance Discovery: Oracle Secure Backup Domain page appears.

12. 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**.

13. Click **Next**.

The Recovery Appliance Discovery: Review page appears.

14. Review the target discovery selections, and click **Submit**.

The Recovery Appliance target is created.

Monitoring TLS-enabled Targets

Enterprise Manager (EM or Cloud Control) can monitor targets that use TCPS.

A wallet needs to be setup at the EM repository. This wallet should contain the certificates needed to connect to the TLS-enabled target. In addition, OMS wallet specific properties have to be set. EM has one wallet, and this wallet is used by EM to communicate with any TLS-enabled target.



Note:

EM agents that monitor TLS-enabled targets also need to have wallets created with the necessary certificates and agent properties. Refer to [Monitoring TLS-enabled Targets with EM Agents](#)

1. Check to see if your EM instance has a wallet and its values set.

```
$OMS_HOME/bin/emctl get property -sysman_pwd <sysmanPwd> -name
em.targetauth.db.pki.TrustStore
```

If a wallet exists, the location of the wallet is output.

If a location exists, it is assumed you have the wallet password.

Verify other values are already set for the wallet:

```
<OMS HOME>/bin>emctl get property -sysman_pwd <sysmanPwd> -name
em.targetauth.db.pki.TrustStoreType
<OMS HOME>/bin>emctl get property -sysman_pwd <sysmanPwd> -name
em.targetauth.db.pki.TrustStorePassword
```

2. If a wallet is not found, create a password protected wallet.

The location of the wallet: `<oms_wallet_location>`.

The password used to protect the wallet: `<oms_wallet_password>`.

```
orapki wallet create -wallet <oms_wallet_location> -auto_login -pwd
<oms_wallet_password>
```

Listing of <oms_wallet_location> must contain a cwallet and an ewallet

3. Set the wallet type, location, and password OMS properties.

```
<OMS_HOME>/bin>emctl set property -sysman_pwd sysman -name
em.targetauth.db.pki.TrustStoreType -value PKCS12

<OMS_HOME>/bin>emctl set property -sysman_pwd sysman -name
em.targetauth.db.pki.TrustStore -value
<oms_wallet_location>/ewallet.p12

<OMS_HOME>/bin>emctl set property -sysman_pwd
sysman -name em.targetauth.db.pki.TrustStorePassword -value
<oms_wallet_password>
```

Validate the entries using the get property call.

4. Add certificates to the wallet from the monitored targets and restart the OMS (if a new wallet was created).

```
orapki wallet add -wallet <oms_wallet_location> -trusted_cert -cert
<certFile>

orapki wallet display -wallet <oms_wallet_location> -complete
```

Validate the certificate is shown correctly.

The EM wallet needs the trusted and signed certificates from Recovery Appliance.

5. If a new wallet was created, restart the OMS.

Monitoring TLS-enabled Targets with EM Agents

Enterprise Manager (EM or Cloud Control) Agents can monitor targets that use TCPS.

A wallet needs to be setup at the EM repository. This wallet should contain the certificates needed to connect to the TLS-enabled target. In addition, Agent specific wallet properties have to be set.

1. Create a wallet for the EM agent.

```
orapki wallet create -wallet <agent_wallet_location> -auto_login -pwd
<agent_password_wallet>
```

2. Add certificates to this wallet.

```
orapki wallet add -wallet <agent_wallet_location> -trusted_cert -cert
<certFile>
```

3. Set the wallet properties for the Agent.

```
<AGENT_HOME>/bin>emctl setproperty agent -name connectionTrustStoreType -
value PKCS12

<AGENT_HOME>/bin>emctl setproperty agent -name
connectionTrustStoreLocation -value <agent_wallet_location>/ewallet.p12
```

```
<AGENT_HOME>/bin>emctl setproperty agent -name  
connectionTrustStorePassword -value <agent_password_wallet>
```

4. Restart the agent.
5. Repeat from the beginning for any database being monitored that uses TLS (TCPS).
Add the database certificate to the wallet for the agent on the database's host.

Discovering TLS-Enabled Recovery Appliance

Enterprise Manager (Cloud Control) requires a few extra steps to discover TLS-enabled Recovery Appliances and TLS-enabled databases.

- The Recovery Appliance should be configured to use TCPS (TLS). Refer to [TLS Overview](#).
- The EM wallet or the EM Agent wallet needs to be updated, or created if it didn't already exist.
 - Refer to: [Monitoring TLS-enabled Targets](#)
 - Refer to [Monitoring TLS-enabled Targets with EM Agents](#)

Discovering a Recovery Appliance using TLS requires that the TCPS port and protocol be specified when discovering the cluster, the Recovery Appliance database, and the Recovery Appliance target.

The **ZDLRA** discovery wizard has an option to upload the Recovery Appliance certificate to EM. The certificate is saved in EM and used when protected databases are configured to backup to this Recovery Appliance. Certificates not added during discovery can also be added after discovery with **Recovery Appliance->Target Setup->TLS Trust Certificate..**

Migrating TCP to TCPS Recovery Appliance

If the Recovery Appliance was already discovered in EM using TCP:

1. Create the EM wallet on both the OMS and the agent. Add the Recovery Appliance certificate(s) to the wallet. Set the EM properties.
2. Edit the port/protocol properties for the cluster and cluster database target associated with the Recovery Appliance.
3. Edit the port/protocol properties for the Recovery Appliance target itself. Upload the Recovery Appliance certificates to EM. They are used when protected databases are configured.

From the **Cluster Target** home page in EM, select **Cluster->Target Setup->Monitoring Configuration**

4. Update **Scan Port** to be the TCPS port.
5. From the **Cluster Database** home page in EM, select **Cluster Database->Target Setup->Monitoring Configuration**
6. In the **Instances** section, edit each instance and update the **Port** and **Connection Protocol**.
7. From the **Recovery Appliance** home page in EM, select **Recovery Appliance->Target Setup->Monitoring Configuration**
8. Update all ports and protocols, including ones for the backup scan and replication scans (if needed).

Discovery and Monitoring of Database Targets using TLS

1. Discover the database in EM, making sure to specify the TCPS port and protocol.
Refer to Discovering and Adding Database Targets.
2. If this is a cluster database, specify the TCPS port for the underlying **Cluster SCAN** port.
3. While discovering the **Cluster** database, change the port for the individual database instances to be TCPS ports.
4. If the database has already been discovered in EM using TCP, migrate this to use TCPS.
 - a. Ensure that the OMS and Agent wallets have been configured as above.
 - b. Ensure that OMS wallet and agent wallet(s) have the DB certificate.
 - c. Change the monitoring configuration for the database. Note that if this is a cluster database, you should change the SCAN port for the underlying cluster and set the TCPS port for all the instances of the cluster database.

Configuring the protected database to backup to the TLS-enabled Recovery Appliance

Prerequisites:

- Add the database to the Recovery Appliance using the **Recovery Appliance->Protected Databases** page.
- Add the Recovery Appliance certificate to EM using the **Recovery Appliance->Target Setup->TLS Trusted Certificate** menu item
- 1. Go to the **Database** home page in EM. Navigate to the **Database->Availability->Backup and Recovery->Configure Backup** page.
- 2. Select Recovery Appliance as the destination and specify the database host credentials.
- 3. Select the **Recovery Appliance**, **VPC** user and **Protocol** to use. The Protocol field offers the choice of TCP and TCPS protocols ONLY if the Recovery Appliance is in dual mode.

The **Configure Backup** workflow detects if the database already has an existing wallet and populate the location of the wallet if one is found.

If the wallet is a password-protected wallet, specify the generic password credential needed to open the wallet. A generic password credential can only be created using **emcli** today.

```
$ emcli create_named_credential -auth_target_type='<system>'
-cred_type=GenericPassword -cred_name="<credName>" -
attributes="GENERIC_PASSWORD:<walletPassword>"
```

Run this command as-is with the exception of these two variables.

- **credName** is the name of the generic named credential you are creating.
 - **walletPassword** is the password for the database wallet.
4. Select other options as desired and then click **Submit**.

A deployment procedure is submitted to configure the database. A link to this procedure is displayed in the confirmation box.

You can also navigate to the **Enterprise->Provisioning and Patching->Procedure Activity** menu item to see the deployment procedure execution details.

5. If the database has already been configured to backup to a Recovery Appliance, the **Configure Backup** page appears when you go to **Availability->Backup and Recovery->Configure Backup**.

Provide the host credential and then invoke the **Change Configuration** action in the **Actions** menu on the right-hand corner to:

- Change the Recovery Appliance details, or
- Change any of the backup options (protocol, enable/disable real time redo, parallelism).

Scheduling Backups from the Datas to the Recovery Appliance

After successful configuration of backups, go to the **Availability->Backup and Recovery->Schedule Backup** menu item from the database home page.

On the **Schedule Backup** page, specify the host credentials for the database host and select how often you'd like the backups to be sent. The suggested backup strategy for backups to the Recovery Appliance is to send incremental backups daily.

Part III

Security and Maintenance of Recovery Appliance

Part III provides information about security and maintenance procedures for Recovery Appliance. It contains the following chapters:

- [User Security on Recovery Appliance](#)
- [Keeping the Recovery Appliance Secure](#)
- [Maintaining the Recovery Appliance Hardware](#)
- [Reconfiguring the Hardware](#)

User Security on Recovery Appliance

Increase the security of your data and system by limiting user access and developing strong password security policies.

Default User Accounts for Oracle Zero Data Loss Recovery Appliance

The following table lists the default users and passwords for the Oracle Zero Data Loss Recovery Appliance components. All default passwords should be changed after installation of the Recovery Appliance.

Table 10-1 Default Users and Passwords for Oracle Zero Data Loss Recovery Appliance

Component	User Name and Password
Compute servers	Operating system users: • root/welcome1 • oracle/Welcome\$ • dbmadmin/welcome • dbmmonitor/welcome • raext/(locked and blocked from SSH access) • railm/(locked and blocked from SSH access) • rasys/(locked and blocked from SSH access) • rasec/(locked and blocked from SSH access) • Password for the GRUB boot loader: sos1Exadata Database users:  Note: Only local connections are allowed for externally authenticated users. • SYS/Welcome\$ • SYSTEM/Welcome\$ • raext/(externally authenticated) • ralim/(externally authenticated) • rasec/(externally authenticated) • rasys/change^Me2 OSB tape backup application users: • admin/welcome1 • oracle/welcome1 • encryption key wallet/welcome1

Table 10-1 (Cont.) Default Users and Passwords for Oracle Zero Data Loss Recovery Appliance

Component	User Name and Password
Storage servers	• root/welcome1 • celladmin/welcome • cellmonitor/welcome • CELLDIAG CELLDIAG is an Exadata storage software user, not an operating system user. 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. • Password for the GRUB boot loader: sos1Exadata
RoCE Network Fabric	• root/welcome1
InfiniBand Network Fabric switches	• root/welcome1 • nm2user/changeme • ilom-admin/ilom-admin • ilom-operator/ilom-operator
Ethernet switches	admin/welcome1 Note: Secure the <code>enable mode password</code> and <code>secret</code> values for the admin user.
Power distribution units (PDUs)	• admin/welcome1 The password for the admin user is admin if you reset the PDU to factory default settings.
Compute server ILOMs	• root/welcome1 • MSUser Management Server (MS) uses this account to manage ILOM and reset it if it detects a hang. Do not modify this account. This account is to be used by MS only. Each time MS starts up, it deletes the previous MSUser account and re-creates the account with a randomly generated password. 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.
Storage server ILOMs	• root/welcome1 • MSUser See the description above for details about this user.
InfiniBand Network Fabric ILOMs	• 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"](#) to learn how to change the passwords for the Recovery Appliance components.

User Roles for the Recovery Appliance

The Recovery Appliance introduces roles for named user accounts and limits operations available to those roles to improve security and logging.

The Recovery Appliance has the following security roles that have changed or are new in software release 21.1, and provide more options to meet audit and security requirements.

- The `rasys` account is the original administrator, root-level account formerly needed to perform operations on the Recovery Appliance. Named users `db_user` with roles and responsibilities replace the usage of `rasys` for day-to-day operations.

The `rasys` account is now an internal user account. It remains the owner of the RMAN catalog, the Recovery Appliance metadata schema, and all user-facing views. It is used during deployment, patch, and upgrade by Oracle Support. The usage of `rasys` is restricted and available only for approved tasks and for break-glass operations.

**Note:**

"Break glass" is any time where the API's do not allow access to the data needed. This might be:

- If we need to set a config parameter which is an underscore.
- If we need access to a trace file that is not accessible.
- If we need to run an internal API (`dbms_ra_int.delete_backup_piece`).

- The `db_user` is a role for new named user who can perform limited operations depending on user types.
 - `admin`: this `db_user` user type replaces the usage of `rasys` for configuration and day-to-day Recovery Appliance management operations. This account can manipulate the database and issue SQL Plus commands.
 - `vpc`: this `db_user` user type is for Virtual Private Catalog (VPC) user activities on the Recovery Appliance. It is required to be in the wallet client side to allow access for backing up and restoring.

- **monitor:** this `db_user` user type is intended for OEM applications like Enterprise Manager and job functions that are read-only for monitoring incidents and the status of the Recovery Appliance.
- The `admin_user` account is a role for new named users who manage the Recovery Appliance from an operation's perspective. It permits operating system level operations on the Recovery Appliance that previously required `root` access. However `admin_user` is not `root`.
- The `sys` account is the super user for Oracle databases, and can change any schema in the database. Remote `sys` access is now disabled and can be selectively enabled for approved tasks and for break-glass operations.

Immutability and Role Management

The enforcement of immutability requires restricting and fully controlling access privileges. The data-to-day Recovery Appliance administration is limited to `admin_user` accounts and to documented comments RACLI and DBMS_RA API comments. The `root` and `rasys` accounts are highly restricted, require a quorum of two other admin users to enable, and can be rejected with a single denial from an admin user.

All operations for auditing purposes are logged and assigned to identifiable users. The command history of `admin_user` OS are provided in the `syslog`. The Recovery Appliance administration commands issued through the API are logged in `RA_API_HISTORY`.

RACLI Non-Root User

Allows a non-privileged user to execute RACLI commands.

The Recovery Appliance in release 21.1 has become more secure by limiting `root` access to the Recovery Appliance. It introduces the `raadmin` group, whose members can execute RACLI commands and thus perform system management that previously required `root` access.

This change aligns the Recovery Appliance with LDAP and Name Services Requests and improves auditing. At the same time, privileged remote access (`root SSH`) is removed for better security.

Most Recovery Appliance management tasks can be performed through non-privileged access to RACLI.

Creating an `admin_user`

Issue the following command from the compute server by providing an appropriate system user name for `<user_name>`.

```
racli add admin_user --user_name=<user_name>
```

This adds an admin user to the `raadmin` group. This admin user is created if it is not found in the `passwd` database. The logic prompts you to enter a user password.

- `racli list admin_user`
Lists all of the users who are in the `raadmin` group and can execute RACLI commands.
- `racli alter admin_user --user_name=<user_name>`

Changes the password for the provided `<user_name>`. The logic prompts you to enter a user password.

- `racli remove admin_user --user_name=<user_name>`

Removes the provided `<user_name>` from the `passwd` database. The `<user_name>` has to be a member of the `raadmin` group.

Securing the Operations of the Recovery Appliance

The following steps harden the Recovery Appliance by reducing exposure to powerful users, like `root` and `rasys` and allowing improved auditing of maintenance actions. Although this procedure is optional for many installations and applications, establishing and using secure users is required for operations to be compliant with various regulatory mandates.

For purposes of example, the sample commands have three fictive users: `bob`, `sue`, and `jim`.

1. Create named users and assign them `db_user` with user type `admin` with administration rights.

The `db_user` user type `admin` replaces the usage of `rasys` for configuration and day-to-day Recovery Appliance management operations. This account can issue certain SQLPlus commands within its assigned privileges.

```
racli add db_user --user_type=admin --user_name=bob
racli add db_user --user_type=admin --user_name=sue
```

In this example, `bob` and `sue` are given `--user_type=admin` for administration rights.

Note:

The `db_user` user type `admin` has limits of privileges, and cannot be used as `sysdba` in SQLPlus.

2. Create `ssh` users for the Recovery Appliance.

The `admin_user` account is a role for new named users who manage the Recovery Appliance from an operation's perspective. It permits operating system level operations on the Recovery Appliance that previously required `root` access, however `admin_user` is not `root`.

```
racli add admin_user --user_name=bob
racli add admin_user --user_name=jim
racli add admin_user --user_name=sue
```

In this example, `bob`, `sue` and `jim` are given `admin_user` with administration rights.

3. Disable `ssh` access for `root` and `oracle`.

```
racli disable ssh
```

4. **Disable root access for root, oracle, and raadmin.**

```
racli disable root_access
```

5. **Disable rasys access.**

 Note:

Make sure that you have the `db_user` user type `admin` accounts and `admin_user` accounts before disabling `rasys` access.

```
racli disable rasys_user
```

6. **Disable sys remote access.**

```
racli disable sys_remote_access
```

7. **Validate the time service.**

Refer to Changing the CHRONY Servers.

8. **Validate that the Recovery Appliance is in compliance.**

```
racli run check --check_name=check_ra_compliance
```

The above should return `TRUE`. The `check_ra_compliance` validates:

- `ssh` access for `root` and `oracle` is disabled on all nodes.
- `rasys` access is disabled.
- `sys` remote access is disabled.
- Time service is enabled.
- Two or more `admin_users` for the Recovery Appliance have been established.
- Two or more `db_users` who are `admin` have been established.

If any of the above items are not completed, `check_ra_compliance` fails, because one or more security gaps still exist on the Recovery Appliance.

At the completion of the above steps:

- The initial set of administrative users have been configured.
- An audit trail of actions by administrative users is now possible.
- Various commands are restricted to users with the proper permissions.
- Certain commands are restricted to *quorum* operations requiring approval of others to finally be run.

Kerberos Authentication and the Recovery Appliance

The Recovery Appliance OS user can be vetted with Kerberos authentication.

1. Follow your data center processes for configuring a Recovery Appliance node to authenticate an OS user with Kerberos.

The Kerberos user must belong to the following groups:

- raadmin
- dbmusers
- oinstall

2. Confirm your Kerberos-authenticated user is accessible on all of the Recovery Appliance compute server nodes in the cluster.

```
getent passwd <USER_NAME>
```

This confirms the client configuration is correct for the name services and that the users are present.

3. From the Recovery Appliance, issue the command to add that Kerberos user as an `admin_user`.

```
racli add admin_user --user_name=USER_NAME [--user_uid=USER_ID --  
user_gid=GROUP_ID]
```

--user_name

System user name to add to RACLI admin group.

--user_uid

Set the user identifier for the newly created admin user. Value must be ≥ 1003 .

During the installation of RA 19.x or later, you can define the `raadmin` uid with `ra_preinstall.pl`.

--user_gid

Set the initial login group identifier for the newly created admin user. A group number must refer to an already existing group. Value must be ≥ 1003 .

During the installation of RA 21.1 or later, you can define the `gid` with `ra_preinstall.pl`.

LDAP Authentication and the Recovery Appliance

The Recovery Appliance offers support for LDAP authentication, which grants named users reduced privileges to manage the Recovery Appliance through the RACLI. These user names, whether LDAP users or native OS users, appear in audit logs for the Recovery Appliance. Direct SSH access for `root` and `oracle` users can be removed from Recovery Appliance nodes.

An existing company LDAP infrastructure can be leveraged to allow OS-level LDAP configuration for computer server nodes. This requires shadow/posix user accounts on the LDAP server. The LDAP users need to belong to the `raadmin` group.

The LDAP user can belong to the following groups:

- raadmin - required
- dbmadmin - used for monitoring Exadata.

- `dbmusers` - used for monitoring Exadata.
- `oinstall`

It is important that the group identifier (GID) is standardized for these groups.

- During install you can use `ra_preinstall.pl` to define a specific group identifier.
- During patch/upgrade, you can also specify `raadmin` GID with `ra_preinstall.pl`.



Note:

If you have an existing ZDLRA 21.1 system with a conflicting GID for `raadmin` group, please open a support case so we can review and help update the group identifier.

1. Follow your data center processes for configuring a Recovery Appliance compute server node to authenticate an OS user with LDAP.
2. Confirm your LDAP-authenticated user is accessible on all of the Recovery Appliance compute server nodes in the cluster.

```
getent passwd <USER_NAME>
```

This confirms that the client configuration is correct for the name services and that the users are present.

3. From the Recovery Appliance, issue the command to add that LDAP user as an `admin_user`.

```
racli add admin_user --user_name=USER_NAME [--user_uid=USER_ID --
user_gid=GROUP_ID]
```

`--user_name`

System user name to add to RACLI admin group.

`--user_uid`

Set the user identifier for the newly created admin user. Value must be ≥ 1003 .

During the installation of RA 19.x or later, you can define the `raadmin` uid with `ra_preinstall.pl`.

`--user_gid`

Set the initial login group identifier for the newly created admin user. A group number must refer to an already existing group. Value must be ≥ 1003 .

During the installation of RA 21.1 or later, you can define the `gid` with `ra_preinstall.pl`.

Additional non-standard packages

If you require additional packages which are non-standard please review the two paths forward and pick the one which most aligns to your environment.

1. If you have an existing LDAP client authentication setup procedure that is used on your Exadata systems which include the nonstandard packages.

Continue using the same process to configure LDAP client authentication which you have been successfully using on your other Engineered systems.

Because these are nonstandard, there is a chance you would need to uninstall them during OS Updates and RACLI Updates.

2. If you have never configured LDAP on an Engineered system, and require non-standard Linux packages.
Contact Support to confirm any risks with installing nonstandard Linux packages on the Database/Compute Server nodes.

Quorum

This chapter describes how quorum works when compliance is in operation on the Oracle Zero Data Loss Recovery Appliance.

When compliance is in effect, certain RACLI commands are not just restricted to privileged users but also can be subject to a quorum operation that requires two approvals and no denials from the set of other privileged users. The two tests for validating quorum are:

- Test 1: `TRUE` if there are backups under compliance, legal hold, or other keep control.
- Test 2: `TRUE` if the compliance mode has been enabled.

If Test 1 or Test 2 are `TRUE`, quorum is required. If both tests are `FALSE`, quorum isn't required.

The quorum scenario given below assumes:

- bob, sue, and jim are `db_users` of the system.
- bob and sue are given `db_user --user_type=admin` for administration rights.
- bob, sue and jim are given `admin_user` with administration rights.

The scenario below illustrates quorum operations.

1. Administrator bob is working. He uses his `db_user --user_type=admin` with his `ssh_user` account. He's been adding protected database and trouble shooting incidents.
2. An issue arises with the Recovery Appliance.
3. The action plan from Oracle Support/Development includes tasks that require `rasys` to run.
4. User bob issues the RACLI command to enable the `rasys` login for 6 hours.

```
racli enable rasys_user --expire=6
```

This returns a request identifier that is associated with the user and an increment, such as bob.1.

5. User bob can monitor that status of his request.

```
racli status request --request_id=bob.1
```

6. At least two users who are `admin_user` must approve the request. Users sue and jim use the request identifier and approve the request.

```
(sue) racli approve request --request_id=bob.1
(jim) racli approve request --request_id=bob.1
(bob) racli status request --request_id=bob.1
```

If one `admin_user` denies the request, then the operation (with that request identifier) will not be processed.

7. When the request is approved, user `bob` can proceed with his task of enabling `rasys`, but this time with the request identifier.

```
racli enable rasys_user --request_id=bob.1
```

This particular operation may prompt `bob` for the password to be used for `rasys` while `rasys` is enabled.

8. User `bob` performs the action plan from Oracle Support/Development, logging in as `rasys` with the password specified by `bob` in the command.
9. User `bob` disables `rasys`.

```
racli disable rasys_user
```

This returns a request identifier that is associated with the user and an increment, such as `bob.2`.

10. User `bob` can monitor that status of his request.

```
racli status request --request_id=bob.2
```

11. At least users who are `admin_user` must approve the request. Users `sue` and `jim` use the request identifier and approve the request.

```
(sue) racli approve request --request_id=bob.2
(jim) racli approve request --request_id=bob.2
(bob) racli status request --request_id=bob.2
```

If one `admin_user` denies the request, then the operation (with that request identifier) will not be processed.

12. When the request is approved, user `bob` can proceed with his task of disabling `rasys`, but this time with the request identifier.

```
racli disable rasys_user --request_id=bob.2
```

Default Password Requirements

Oracle Exadata Deployment Assistant (OEDA) implements a default password policy on Oracle Exadata Database Machine.

The last step of OEDA, "Secure Oracle Exadata Database Machine", implements the following password requirements:

- Dictionary words are not valid or accepted.
- Character classes for passwords are uppercase letters, lowercase letters, digits, and special characters.
- Passwords must contain characters from all four character classes. Passwords using only one, two, or three character classes are not allowed.
- The minimum length of a password is eight characters.

- Pass-phrases are allowed. A pass-phrase should contain at least three words, be 16 to 40 characters in length, and contain different character classes.
- A new password cannot be similar to old passwords. There must be at least eight characters in the new password that were not present in the old password.
- A maximum of three consecutive characters of the same value can be used in a password.
- A maximum of four consecutive characters of the same character class can be used in a password. For example, `abcde1#6B` cannot be used as a password because it uses five consecutive lower case letters.

Default Security Settings Implemented by OEDA

Oracle Exadata Deployment Assistant (OEDA) includes a step to implement default security settings on Recovery Appliance.

The last OEDA configuration step implements the following security settings:

- The following password rules apply by default for all operating system users on the compute servers and storage servers:
 - Non-root users must change their password during first login.
 - The password complexity rules depend on the Oracle Linux version in use.
For systems with Oracle Linux 7 or later:
 - * The minimum password length is 8 characters,
 - * The password must contain at least one digit, one uppercase character, one lowercase character, and one other character.
 - * The password must not contain the same character consecutively more than 3 times.
 - * The password must not contain more than 4 consecutive characters from the same class (digits, lowercase letters, uppercase letters, or other characters).
 - * For password changes, the new password must contain a minimum of 8 character changes.
 - For systems with Oracle Linux 6 or earlier, the minimum password length is 5 characters with no additional complexity requirements.
 - The maximum password age is 60 days.
 - The minimum amount of time between password changes is 1 day.
 - Warning alerts are generated 7 days before password expiry.
 - When changing a user password, the new password cannot match any of the 10 previous passwords.
- An operating system user account is locked for 15 minutes after three failed login attempts within a 15-minute period.
- For the `root` user, SSH equivalency is removed for all compute servers and storage servers.

Keeping the Recovery Appliance Secure

This chapter describes policies and procedures to keep Recovery Appliance secure.

Securing the Hardware

After installation of Oracle Zero Data Loss Recovery Appliance, the hardware should be secured.

Hardware can be secured by restricting access to the hardware and recording the serial numbers. Oracle recommends the following practices to restrict access:

- Install Oracle Zero Data Loss Recovery Appliance and related equipment in a locked, restricted-access room.
- Lock the rack door unless service is required on components within the rack.
- Restrict access to hot-pluggable or hot-swappable devices because the components can be easily removed by design. See
- Store spare field-replaceable units (FRUs) or customer-replaceable units (CRUs) in a locked cabinet. Restrict access to the locked cabinet to authorized personnel.
- Mark all significant items of computer hardware, such as FRUs.
- Keep hardware activation keys and licenses in a secure location that is easily accessible to the system managers in the case of a system emergency.
- Record the serial numbers of the components in Oracle Zero Data Loss Recovery Appliance, and keep a record in a secure place. All components in Oracle Zero Data Loss Recovery Appliance have a serial number.

Related Topics

- [How To Obtain The Serial Number Associated With The System Board, Motherboard, Disk Controller, Disks, Infiniband HCA And More Contained In A Cell Or Compute Box \(Exadata-Sun V2 or X2 / 11.2\)? \(My Oracle Support Doc ID 949614.1\)](#)
- [How to Determine the Serial Number of a Datacenter InfiniBand Switch 36 or QDR InfiniBand Gateway InfiniBand Switch \(My Oracle Support Doc ID 1299791.1\)](#)

Getting the Rack Serial Number

Use the `ipmitool` utility to get the serial number for the rack.

When interacting with Oracle Support Services, the CSI number for a rack is based on the rack serial number.

1. Log in to one of the servers in the rack as the `root` user.
2. Use `ipmitool` to get the serial number for the rack.

```
# ipmitool sunoem cli "show /SP system_identifier"
Connected. Use ^D to exit.
-> show /SP system_identifier
```

```
/SP
Properties:
  system_identifier = Exadata Database Machine X2-8xxxxAKyyyy

-> Session closed
Disconnected
```

Getting the Serial Numbers for Rack Components

The CheckHWnFWProfile command can be used to display the serial number of most of the system components.

1. Log in to one of the servers in the rack as the `root` user.
2. On each server in the rack, use `CheckHWnFWProfile` with the `-S` option to display the serial number of the components for that server.

```
# /opt/oracle.SupportTools/CheckHWnFWProfile -S > /tmp/
CheckHWnFWProfile_hostname.txt
```

The result is specific to each server, so the command must be performed on every node. The following is a partial example of the output:

```
Server_Model=ORACLE_SERVER_X8-2L
====START SERIAL NUMBERS====
==Motherboard, from dmidecode==
--System serial--
1904XCA000
--Motherboard serial--
469996N+0000RD01RN
--Chassis serial--
1900XCA000
--Rack serial--
AK00400000
==Infiniband HCA==
ID:      CX354A - ConnectX-3 QSFP
PN:      7046442
EC:      XX
SN:      465000K-1800000000
V0:      PCIe Gen3 x8
==Motherboard, RAM etc from ipmitool==
FRU Device Description : Builtin FRU Device (LUN 0 ID 0)
...
Product Name           : ILOM
Product Version        : 4.0.4.38.a

FRU Device Description : BMC
...
Product Name           : ILOM
Product Version        : 4.0.4.38.a

FRU Device Description : /SYS (LUN 0 ID 3)
```

```

...
Product Part Number   : 8200669
Product Serial        : 1900XCA000

FRU Device Description : DBP (LUN 0 ID 210)

Board Part Number     : 7341141
Board Extra           : Rev 09

FRU Device Description : HDD0 (LUN 0 ID 47)
Device not present (Requested sensor, data, or record not found)

FRU Device Description : HDD1 (LUN 0 ID 48)
Device not present (Requested sensor, data, or record not found)

...

FRU Device Description : MB (LUN 0 ID 4)
Board Mfg Date         : Sun Jan 20 16:57:00 2019
Board Mfg              : Oracle Corporation
...

FRU Device Description : MB/BIOS (LUN 0 ID 5)
...

FRU Device Description : MB/CPLD (LUN 0 ID 8)
Product Manufacturer   : Oracle Corporation
Product Name           : Power Control FPGA
Product Version        : FW:3.9

FRU Device Description : M2R0/SSD0 (LUN 0 ID 211)
Device not present (Requested sensor, data, or record not found)

FRU Device Description : M2R1/SSD0 (LUN 0 ID 212)
Device not present (Requested sensor, data, or record not found)

FRU Device Description : MB/NET0 (LUN 0 ID 43)
Product Manufacturer   : INTEL
Product Name           : 1G Ethernet Controller
...

FRU Device Description : MB/P0 (LUN 0 ID 16)
Product Manufacturer   : Intel
Product Name           : Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz
...

FRU Device Description : MB/P0/D0 (LUN 0 ID 24)
Product Manufacturer   : Samsung
Product Name           : 16384MB DDR4 SDRAM DIMM
...

FRU Device Description : MB/P0/D1 (LUN 0 ID 25)
Device not present (Requested sensor, data, or record not found)

FRU Device Description : MB/P0/D2 (LUN 0 ID 26)
Product Manufacturer   : Samsung

```

```
Product Name          : 16384MB DDR4 SDRAM DIMM
...

FRU Device Description : MB/P1 (LUN 0 ID 17)
Product Manufacturer   : Intel
Product Name           : Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz
...

FRU Device Description : MB/P1/D0 (LUN 0 ID 36)
Product Manufacturer   : Samsung
Product Name           : 16384MB DDR4 SDRAM DIMM
...

FRU Device Description : PS0 (LUN 0 ID 63)
...

FRU Device Description : PS1 (LUN 0 ID 64)
...

FRU Device Description : SP/NET0 (LUN 0 ID 1)
...

FRU Device Description : SP/NET1 (LUN 0 ID 2)
...

FRU Device Description : /UUID (LUN 0 ID 6)
...

FRU Device Description : TOP_LEVEL_CH (LUN 0 ID 251)
Chassis Type           : Rack Mount Chassis
Chassis Part Number    : 8200669
Chassis Serial         : 1900XCA0000
Chassis Extra          : chassis_name:ORACLE SERVER X8-2L

FRU Device Description : TOP_LEVEL_PROD (LUN 0 ID 250)
Product Manufacturer   : Oracle Corporation
Product Name           : Exadata X8-2
Product Part Number    : Exadata X8-2
Product Serial         : AK00430000

====END SERIAL NUMBERS====
```

Getting the Rack Serial Number for a Cisco 9336C or 9348 Switch

Use the `show license host-id` command on the switch to get the serial number.

1. Connect to the switch from a server with SSH equivalency configured, or log in as the admin user.
2. Obtain the serial number for the switch by entering the `show license host-id` command.

The host ID is also referred to as the device serial number.

```
# switch# show license host-id
License hostid: VDH=FLA12345678
```

Use the entire ID that appears after the equal sign (=). In this example, the host ID is FLA12345678.

Getting the Rack Serial Number for a Sun Datacenter InfiniBand Switch 36

Use the `showfruinfo` command on the switch to get the serial number.

1. Log in to the switch as `root`.

```
$ ssh root@switch_name
```

2. Use the `showfruinfo` command to view the serial number for the switch.

```
root@ib-switch-> showfruinfo
Sun_Man1R:
UNIX_Stamp32 : Fri Mar 19 16:29:59 2010
Sun_Fru_Description : ASSY,NM2-GW
Vendor_ID_Code : 11 E1
Vendor_ID_Code_Source : 01
Vendor_Name_And_Site_Location : 4577 CELESTICA CORP. SAN JOSE CA US
Sun_Part_Number : 5111402
Sun_Serial_Number : 0110SJC-1010NG0040
Serial_Number_Format : 4V3F1-2Y2W2X4S
Initial_HW_Dash_Level : 03
Initial_HW_Rev_Level : 50
Sun_Fru_Shortname : NM2 gateway
Sun_Hazard_Class_Code : Y
Sun_SpecPartNo : 885-1655-01
Sun_FRU_LabelR:
Sun_Serial_Number : AK000XXXX2
FRU_Part_Dash_Number : 541-4188-01
```

Getting the Serial Number for a Cisco 4948 Ethernet Switch

Use the `sh inventory` command on the switch to get the serial number.

1. Log in to the Cisco Ethernet switch.
2. Obtain the serial number for the switch and its components by entering the `sh inventory` command.

```
# Switch# sh inventory
NAME: "Switch System", DESCR: "Cisco Systems, Inc. WS-C4948 1 slot switch "
PID: , VID: , SN: FOX0000G0B6
NAME: "Linecard(slot 1)", DESCR: "10/100/1000BaseT (RJ45), 1000BaseX
(SFP)
Supervisor with 48 10/100/1000BASE-T ports and 4 1000BASE-"
PID: WS-C4948 , VID: V09 , SN: FOX0000G0B6
NAME: "Power Supply 1", DESCR: "Power Supply ( AC 300W )"
PID: PWR-C49-300AC , VID: , SN: QCS0000B1XR
NAME: "Power Supply 2", DESCR: "Power Supply ( AC 300W )"
PID: PWR-C49-300AC , VID: , SN: QCS0000B1X5
```

Securing the Software

Frequently, hardware security is implemented through software measures.

Implement the following guidelines to protect hardware and software:

- Change all default passwords when the system is installed at the site. Recovery Appliance uses default passwords for initial installation and deployment that are widely known. A default password could allow unauthorized access to the equipment. Devices such as the network switches have multiple user accounts. Be sure to change all account passwords on the components in the rack.
- Limit use of the `root` super user account. Use `non-root` access when possible. Create and use Integrated Lights Out Manager (ILOM) user accounts for individual users to ensure a positive identification in audit trails, and less maintenance when administrators leave the team or company.
- Restrict physical access to USB ports, network ports, and system consoles. Servers and network switches have ports and console connections, which provide direct access to the system.
- Restrict the capability to restart the system over the network.
- Create named `admin_users` to manage the hosts.
- Disable direct root and oracle access.
- Create named `db_users` for administration and monitoring.
- Disable remote `sys` access.
- Disable remote `rasys` access.

Maintaining a Secure Environment

After security measures are implemented, they must be maintained to keep the system secure.

Software, hardware and user access need to be updated and reviewed periodically. For example, organizations should review the users and administrators with access to Recovery Appliance to verify if the levels of access and privilege are appropriate. Without review, the level of access granted to individuals may increase unintentionally due to role changes or changes to default settings. It is recommended that access rights for operational and administrative tasks be reviewed to ensure that each user's level of access is aligned to their roles and responsibilities.

Refer to User Accounts in the Recovery Appliance Environment.

Organizations are encouraged to utilize tools to detect unauthorized changes, configuration drift, and prepare for security updates. Oracle Enterprise Manager provides an integrated solution for managing operational issues for hardware, deployed applications, and services.

Maintaining Network Security

Follow these guidelines to ensure the security of local and remote access to the system:

- Network switch configuration files should be managed offline, and access to the configuration file should be limited to authorized administrators. The configuration file should contain descriptive comments for each setting. Consider keeping a static copy of the configuration file in a source code control system.

For more information on network switch configuration, refer to the vendor documentation for the network switch.

- Review the client access network to ensure that secure host and Integrated Lights Out Manager (ILOM) settings are in effect. Review the settings periodically to ensure that they remain intact.
- Use only signed certificates from the Certification Authority.
- Set time-outs for extended sessions and set privilege levels.
- Use authentication, authorization, and accounting (AAA) features for local and remote access to a network switch.
- Use the port mirroring/switch port analyzer (SPAN) capability of the switch for intrusion detection system (IDS) access.
- Implement port security to limit access based upon a MAC address (MAC ACL).
- Require users to use strong passwords by setting minimum password complexity rules and password expiration policies.
- Enable logging and send logs to a dedicated secure log host.
- Configure logging to include accurate time information, using NTP and timestamps.
- Review logs for possible incidents and archive them in accordance with the organization's security policy.

Recovering the Recovery Appliance

The Recovery Appliance has a background process that backs up the necessary EM agent, OSB catalog, and Recovery Appliance catalog information (e.g. backup metadata, protected database/VPC users, protection policies, etc.) This process runs every 2 to 8 hours (depending on release and number of racks in your Recovery Appliance) alternating between compute nodes. This process is scheduled to run from the crontab every 12 hours for 2 node cluster or every 6 hours for a 4 node cluster.

In the event of a complete system failure, this backup is used to recover the Recovery Appliance catalog, along with all users and policies at the time the backup was taken.

Goals of the Recovery Appliance Backup/Recovery process:

- Protected databases should be able to resume their backups as soon as the system is back online, with no changes to database-side configuration.
- Protected databases should have access to Recovery Appliance backups (including tape and replicated copies) as soon as the system is back online .

The most recent export bundle is recommended to be copied from the Recovery Appliance to an external disk or tape location periodically (once or twice a day). This ensures that the export bundles are preserved outside of the Recovery Appliance, for future data recovery purposes.

The export file has the naming convention `ra_backup.$date.$$tar.gz`, where `$date` = `$year . $month . $day . '.' . $hour . $min`, and `$$` is the process id of the running process. The location for the export file is `/raacfs/raadmin/ra_export/`.

Example: `/raacfs/raadmin/ra_export/ra_backup.2022630.051.47141.tar.gz`

To recover a Recovery Appliance:

1. For a one-time export of the Recovery Appliance state, start the export process.

```
[root@<HOST>~]# /opt/oracle.RecoveryAppliance/bin/ra_export.sh --manual
Created ZDLRA export /raacfs/raadmin/ra_export/
ra_backup.2019313.1341.333165.tar.gz
```

This covers OSB, EM, and the Recovery Appliance components (non delta pool backup metadata, protected database users, protection policies, etc.)

The automatic exports happen multiple times per day. The export bundles can be copied off-site from this location:

```
/raacfs/raadmin/ra_export/
```

2. After a system outage requiring recovery of the Recovery Appliance, copy the last exported bundle into `/opt/oracle.RecoveryAppliance/import/` onto the fixed or new Recovery Appliance.
3. Run the installation on the fixed or new Recovery Appliance.

```
[root@<HOST>~]# racli install appliance
```

Maintaining the Recovery Appliance Hardware

This chapter describes how to maintain the Recovery Appliance rack components. It contains the following topics:

- [Cautions and Warnings](#)
- [Determining the Server Model](#)
- [Powering On and Off a Recovery Appliance Rack](#)
- [Replacing the Disk Controller Batteries](#)
- [Replacing a Power Distribution Unit](#)
- [Resetting a Non-Responsive Oracle ILOM](#)
- [Maintaining the Compute Servers](#)
- [Reimaging a Compute Server](#)
- [Maintaining the Storage Servers](#)
- [Maintaining the Physical Disks of Storage Servers](#)
- [Maintaining the Flash Disks of Storage Servers](#)
- [Replacing a Disk Controller Battery Backup Unit](#)
- [Using the Storage Server Rescue Procedure](#)



See Also:

[Replacement Units](#)

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](#)".
- 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.

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
```

Powering On and Off a Recovery Appliance Rack

This section includes the following topics:

- [Emergency Power-Off Procedure](#)
- [Shutting Down Recovery Appliance](#)
- [Starting Up Recovery Appliance](#)

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

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.

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.

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.

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. Disable the keystore as part of appliance shutdown. If the keystore is not created, skip this step.

If a keystore is created, this is required to stop the HSM wallet for copy to cloud encryption.

Before appliance shutdown, run:

```
[root@myhost ~]# racli disable keystore
[root@myhost ~]# racli status keystore
      Node: zdlra41
Wallet Type: HSM
      Status: Closed
      Node: zdlra42
Wallet Type: HSM
      Status: Closed
```

2. Log in as `oracle` to either Recovery Appliance compute server.
3. Open a SQL connection to Oracle Database as the `rasys` user:

```
$ sqlplus rasys
```

4. Check the status of the services:

```
SQL> SELECT state FROM ra_server;

STATE
-----
ON
```

5. Shut down Recovery Appliance services:

```
SQL> exec dbms_ra.shutdown;
```

6. Disconnect from Oracle Database:

```
SQL> exit
```

7. 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.

8. 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
```

9. Stop Oracle Database:

```
$ srvctl stop database -d zdlra5
```

10. 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
```

11. Switch to the `root` user.

12. 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.

13. 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.
```

14. Shut down or reboot the hardware as required, in the following order:
 - a. Compute servers
 - b. Storage servers
 - c. Rack and switches

Powering Down the Servers

Before powering down a server, stop the services running on it, as described in ["Shutting Down Recovery Appliance"](#).

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 13-1 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 13-1](#) 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.

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.

Starting Up Recovery Appliance

Turn on the rack components first, then start the software services.

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](#)".

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.

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 `ra01cel01` using IPMI, you use its Oracle ILOM with a command like the following:

```
# ipmitool -H ra01cel01-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>

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 osbadmin
```

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. Enable the keystore as part of appliance startup. This is required step to open the HSM wallet for copy to cloud encryption. After appliance restart, run:

```
[root@myhost ~]# racli enable keystore
[root@myhost ~]# racli status keystore
      Node: zdlra42
Wallet Type: HSM
      Status: Open
      Node: zdlra41
Wallet Type: HSM
      Status: Open
```

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 13-1](#) shows the worst case life expectancy of a battery in Recovery Appliance.

Table 13-1 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

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 milliampere 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.

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.

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.

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.

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. 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>

Replacing a Tape Drive

A tape drive can be replaced while Recovery Appliance is online.

 Note:

The bridging tape drive is located in the top drive slot of the base module. The robot control is a SCSI Medium Changer device that appears as LUN 1 on a bridging tape drive. If the library is partitioned, there must be two tape drives in the base module, and each tape drive provides the robot control for its assigned partition. If the failed drive is the bridging drive, then the SL150 library must be taken offline to the host as replacing the tape drive will cause the host to lose connectivity to the SL150 library. If the tape drive is not a bridging drive, then it can be hot swapped.

To replace a tape drive:

1. Pause the Recovery Appliance `sbt_library`.

```
exec dbms_ra.pause_sbt_library(lib_name=>'ROBOT0');
```

2. Quiesce activity for this tape drive.
3. Log into the SL150 remote interface using your browser.
4. (Optional) Enable the locate library indicator
5. Prepare the tape drive for removal:
 - a. Click *Library* on the left of the remote interface.
 - b. Move the cursor to the drive you need to replace.
 - c. Click the drive icon and select *Remove Drive*.
 - d. Click OK in the confirmation dialog box. The physical LED at the rear of the drive tray lights to indicate the drive is ready for removal.
6. Remove the tape drive:
 - a. Access the back of the library (open the rear door of the rack, if applicable).
 - b. Locate the tape drive with the blue LED (indicates drive is ready for removal).

- c. Verify that the interface cables are labeled. Attach a label if necessary.
- d. Disconnect the cables from the jacks on the left side of the drive tray.
- e. Loosen the captive thumbscrews on the drive tray.
- f. Grasp the drive tray, pull it out of the library drive slot, and set it upright and flat on a static free surface.

 Caution:

ESD damage. Do not touch any electronic components or contacts.

- 7. Remove the replacement drive from its packaging.
- 8. Replace the tape drive:
 - a. Grasp the rear corners of the drive tray.
 - b. Guide the front of the drive tray into the module drive slot.
 - c. Push the drive tray completely into the drive slot.
 - d. Verify that the LEDs are active on the rear of the drive tray.
 - e. Tighten the captive thumbscrews firmly on each side of the drive tray to make sure there is no tray movement in any direction.
- 9. Push the Locate indicator on the robot CRU to extinguish the light, if applicable.
- 10. Confirm that the library recognizes and accounts for the drive from the Web GUI or front Op Panel.
- 11. Make sure the drive port is enabled. View the *Drive Properties* and change drive settings if appropriate.
- 12. Connect the interface and Ethernet cables to the proper jack on the left side of the drive tray.
- 13. Log out of the SL150 remote interface or return the touch screen to Home.
- 14. Run `obtool` to update serial number of the drive object.

```
obtool chdev -S <drive_object_replaced> | obtool chdev -S robot0_tape03
```

- 15. Resume ZDLRA `sbt_library`.

```
exec dbms_ra.resume_sbt_library(lib_name=>'ROBOT0');
```

Many of the above steps come from [How to Remove/Replace a Tape Drive CRU in the StorageTek SL150 Modular Tape Library \[VCAP\] \(Doc ID 1473764.1\)](#).

Save Configuration before Updating Firmware

Preserve the existing Recovery Appliance configuration when updating Oracle ILOM system firmware.

You can update Oracle ILOM system firmware while the host is powered on. The Oracle ILOM firmware image includes firmware for the service processor (SP, Oracle ILOM) and the server's host components (FPGAs). The Oracle ILOM firmware update goes into effect immediately. However, the update of host components is deferred until the affected host is power cycled.

Because Oracle ILOM can be updated while the host is powered on, this feature reduces the total system downtime.

Update the firmware using the Oracle ILOM command-line interface.



Note:

To save the configuration for the Recovery Appliance during a firmware upgrade, answer the prompts with **y** when asked Preserve existing SP configuration (y/n)?

1. Log in to Oracle ILOM with an account with Admin privileges.
2. Load the firmware image from the stored location using the `load -source` command followed by the directory path to the firmware image you want to install. Type:

```
-> load -source protocol://server_IPaddress/<path_to_image>/<image.pkg>
```

Where **protocol** can be: http, https, ftp, tftp, sftp, scp

For example, if you are accessing the server through a tftp server with an IP address of 198.xxx.yyy.123 in a directory called `ilom/jdoe` and with the `<image.pkg>` named `firmware.pkg`, enter the following command:

```
-> load -source tftp://198.xxx.yyy.123/tftpboot/ilom/jdoe/firmware.pkg
```

The following information is displayed:

An upgrade takes several minutes to complete. Oracle ILOM will enter a special mode to load new firmware. No other tasks can be performed in Oracle ILOM until the firmware upgrade is complete and Oracle ILOM is reset.

You can choose to postpone the server BIOS upgrade until the next server power off. If you do not do that, you should perform a clean shutdown of the server before continuing.

3. Answer the following prompts:

```
Are you sure you want to load the specified file? y
```

```
Preserve existing SP configuration (y/n)? y
```



Note:

This prompt preserves your existing Oracle ILOM settings after the firmware update is complete. A **y** preserves the existing Recovery Appliance configuration.

```
Preserve existing BIOS configuration (y/n)? y
```



Note:

This prompt preserves your existing BIOS configuration settings after the firmware upgrade is complete.

```
Delay BIOS upgrade until the next poweroff or reset (y/n)? y
```


Answer "Y" (yes) to the Delay BIOS Upgrade question and, if the host is ON and there are host components to be updated, the host remains powered on and the host component updates are deferred until the next time the host powers off and powers on (next reset/reboot).

Answer "N" (no) to the Delay BIOS Upgrade question and, if the host is ON and there are host components to be updated, the host is forced OFF so that host component updates can be applied immediately. After Oracle ILOM reboots, the host is powered on automatically if it was forced off.

 Note:

If the server has a pending BIOS upgrade, the power reset could take longer to complete. This is expected behavior, as it is necessary to power cycle the server to upgrade the BIOS firmware. If the upgrade includes an FPGA update, the process can take as long as 26 minutes to complete.

4. Wait for the Oracle ILOM status message to confirm that the process is complete.

 Note:

For more details, see [Update System Firmware Using Oracle ILOM](#)

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](#)
- [Resetting Oracle ILOM Using the Remote Console](#)
- [Resetting Oracle ILOM Using IPMITool](#)
- [Resetting Oracle ILOM By Removing Power](#)

 See Also:

Oracle Integrated Lights Out Manager (ILOM) 3.0 documentation at

<http://docs.oracle.com/cd/E19860-01/E21549/bbgiedhj.html#z4000b491400243>

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
```

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**.

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.

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.

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"](#) for information about the LEDs
- ["Parts for Compute Servers"](#) for the repair procedures

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
Firmware state: Online, Spun Up
Firmware state: Online, Spun Up
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.

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"](#)
- ["Downloading the Latest Release of the Cluster Verification Utility"](#)
- ["Removing the Failed Compute Server from the Cluster"](#)
- ["Preparing the USB Flash Drive for Imaging"](#)
- ["Copying the Image to the New Compute Server"](#)
- ["Configuring the Replacement Compute Server"](#)
- ["Preparing the Replacement Compute Server for the Cluster"](#)
- ["Applying Patch Bundles to a Replacement Compute Server"](#)
- ["Cloning the Oracle Grid Infrastructure"](#)
- ["Clone Oracle Database Homes to the Replacement Compute Server"](#)

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.

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.

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 2345667aabbdd.
...
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 "ra01db02"
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

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 d1360
# ./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 `d1360` directory and the `computeImageMaker` file from the working compute server. The directory and file require about 2 GB of disk space.

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](#).

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.

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.

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.

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.

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](#)".

```
# 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

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

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](#)
- [Enabling Network Connectivity Using the Diagnostics ISO](#)

**Note:**

Older storage servers cannot be removed and replaced with newer storage servers while keeping existing Recovery Appliance backups online. If such an exchange is done, a re-image of the Recovery Appliance is required. (Painful!)

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](#)".
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, then reactivate them:

```
CellCLI> ALTER GRIDDISK ALL ACTIVE
```

If you skipped step 2, 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."

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.

Maintaining the Physical Disks of Storage Servers

This section contains the following topics:

- [About System Disks and Data Disks](#)
- [Monitoring the Status of Physical Disks](#)
- [What Happens When Disk Errors Occur?](#)
- [About Detecting Underperforming Disks](#)
- [About Rebalancing the Data](#)
- [Monitoring Hard Disk Controller Write-Through Cache Mode](#)
- [Replacing a Failed Physical Disk](#)
- [Replacing a Faulty Physical Disk](#)
- [Removing an Underperforming Physical Disk](#)
- [Moving All Drives from One Storage Server to Another](#)
- [Removing and Replacing the Same Physical Disk](#)
- [Reenabling a Rejected Physical Disk](#)



See Also:

Oracle Maximum Availability Architecture (MAA) website at <http://www.oracle.com/goto/maa> for additional information about maintenance best practices

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.

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

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](#)" 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 "](#)
- ["Replacing a Faulty Physical Disk"](#)

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 13-2](#) describes the conditions that trigger disk confinement:

Table 13-2 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"](#).

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
.
.
.
```

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.

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:
```

```
GasGuageStatus:  
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
```

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"](#).



See Also:

- ["Parts for Storage Servers"](#)
- *Oracle Database Reference* about the `V$ASM_OPERATION` view

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"](#).

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"](#).

**See Also:**

- ["Parts for Storage Servers"](#)
- *Oracle Database Reference* for information about the `V$ASM_OPERATION` view

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"](#)

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](#)".
 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:

```
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.

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.

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 reenble 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 reenble force
Physical disk hard_disk_name/hard_disk_id was reenabled.
```

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](#)
- [Faulty Status Indicators](#)
- [Identifying Flash Disks in Poor Health](#)
- [Identifying Underperforming Flash Disks](#)
- [When Is It Safe to Replace a Faulty Flash Disk?](#)
- [Replacing a Failed Flash Disk](#)
- [Replacing a Faulty Flash Disk](#)
- [Removing an Underperforming Flash Disk](#)
- [About Write-Back Flash Cache](#)

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" 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

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.

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
```

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 13-2](#) 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 CONFINED 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
.
.
.
```

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.

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](#)".

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](#)" 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

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?"](#) 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"](#).

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.

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](#)".

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](#)".
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, 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, then you must manually add the grid disks back to the Oracle ASM disk group.

About Write-Back Flash Cache

You cannot modify the write-back flash cache settings on Recovery Appliance.

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](#)
- [Replacing a Disk Controller BBU on a Storage Server](#)



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.

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. Otherwise, go to step 7.

6. Verify that the battery state is Operational. This step is required only when the cache policy output from step 5 is not writeback.

```
# /opt/MegaRAID/MegaCli/MegaCli64 -adpbbscmd -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.

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.

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](#)
- [About the Rescue Procedure](#)
- [Rescuing a Server Using the CELLBOOT USB Flash Drive](#)
- [Reconfiguring the Rescued Storage Server](#)
- [Recreating a Damaged CELLBOOT USB Flash Drive](#)

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.

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

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.

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](#)" or "[If the Server Has High Redundancy](#)".
- 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

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.

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

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
```

Reconfiguring the Hardware

This chapter describes how to make basic adjustments. It contains the following topics:

- [Changing Component Passwords](#)
- [Changing the DNS Servers](#)
- [Changing the chrony/NTP Servers](#)
- [Changing the Time Zone Settings](#)

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](#)
- [Changing the Storage Server Passwords](#)
- [Creating New Users and Changing Passwords on a PDU](#)
- [Changing the Oracle ILOM Password](#)
- [Changing an InfiniBand Switch Password](#)
- [Changing the Cisco Ethernet Switch Password](#)
- [Changing the Recovery Appliance Database User Password](#)
- [Changing OSB Tape Backup Passwords](#)



See Also:

[Default User Accounts for Oracle Zero Data Loss Recovery Appliance](#) for a list of default users and passwords of the various Recovery Appliance components.

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`.

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
```

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.

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
```

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.

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
```

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

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
```

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 `rasys`.
2. Issue the following RACLI command.

```
racli alter rasys_user --password
```

3. When prompted, enter and confirm 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.

See Also:

[Default User Accounts for Oracle Zero Data Loss Recovery Appliance](#) for information about the default Recovery Appliance database accounts and passwords.

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 for Oracle Zero Data Loss Recovery Appliance](#) for information about the default OSB tape backup accounts and passwords for Recovery Appliance.

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"](#)
- ["Changing the DNS Server Address on the Cisco Ethernet Switch"](#)
- ["Changing the DNS Server Address on the Compute Server"](#)
- ["Changing the DNS Server on Storage Servers"](#)

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
```

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
```

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.

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.

Changing the chrony/NTP Servers

You can change the configuration settings for the `chrony` implementation of the Network Time Protocol (NTP) servers after the initial setup. All servers and switches in Recovery Appliance must reference the same NPT 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 `chrony`/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"](#)
- ["Setting the CHRONY Server Address on the InfiniBand Switches"](#)
- ["Set the NTP Server Address on the Cisco Ethernet Switch"](#)
- ["Setting the CHRONY Server on the Storage Servers"](#)
- ["Restart Recovery Appliance"](#)



Note:

- You must restart Recovery Appliance after completing these tasks.
- You can configure one or two CHRONY servers for use with Recovery Appliance.

Setting the NTP Server Address on the Compute Servers

To set the `chrony`/NTP server address on a compute server:

1. Log in to a compute server as `root`.
2. Stop the `chrony`/NTP services on the server:

```
service chronyd stop
```
3. Open `/etc/chrony.conf` in a text editor and enter the IP address of the new `chrony`/NTP server.
4. Restart the `chrony`/NTP services on the compute server:

```
service chronyd start
```
5. Repeat this task for the second compute server.

Setting the CHRONY Server Address on the InfiniBand Switches



Note:

InfiniBand switches are used for X8 and earlier compute servers of the Recovery Appliances. X8M and X9M use different network fabric switches.

Do not manually edit the files on the InfiniBand switches.

To set the chrony/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/chrony/server/1 address=ntp_ip1
set /SP/clients/chrony/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.

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 chrony/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
```

Setting the CHRONY Server on the Storage Servers

To set the `chrony`/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.

Restart Recovery Appliance

After changing the servers and switches, restart Recovery Appliance.



See Also:

["Powering On and Off a Recovery Appliance Rack"](#)

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"](#)
- ["Change Time Zone Settings on the Compute Servers"](#)
- ["Change the Time Zone Settings on the InfiniBand Switches"](#)
- ["Change the Time Zone Settings on the Cisco Switch"](#)

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.

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
```

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**.

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

Recovery Appliance Command Line Interface Reference

RACLI is a command line utility for installing and administering the Recovery Appliance system.

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	Adds an object to the Recovery Appliance.
racli add admin_user	Adds a specified ADMIN_USER system user to the RACLI admin group.
racli add certificate	Adds https certification on the Recovery Appliance.
racli add cloud_key	Generates the public and private key files required for OCI access.
racli add cloud_location	Adds a cloud storage object to the Recovery Appliance.
racli add cloud_user	Adds a cloud user to the Recovery Appliance schema.
racli add credential	Adds a user credential to the wallet of the Recovery Appliance.
racli add credential_wallet	Adds a wallet object to the Recovery Appliance.
racli add db_user	Adds a database administration user who can manage and configure the Recovery Appliance.

Command	Description
<code>racli add host</code>	Adds a host to those available on the Recovery Appliance.
<code>racli add keystore</code>	Adds a keystore to the Recovery Appliance.
<code>racli add migration_server</code>	Adds a migration server to the Recovery Appliance.
<code>racli add nfs</code>	Specifies that the Network File System (NFS) should be added to the Recovery Appliance schema.
<code>racli add protected_db</code>	Adds a protected Database to the Recovery Appliance.
<code>racli add rasys_wallet</code>	Creates the rasys user wallet.
<code>racli add remote_syslog</code>	Creates a configuration file to export log files to a remote server.
<code>racli add storage_cells</code>	Prepares new Recovery Appliance storage cells for storage location expansion.
<code>racli add storage_disks</code>	Adds disks with the newly added storage cells to the ASM disk groups for storage location expansion.
<code>racli add tag</code>	Specifies the diagnostic TAG to add to all nodes of the Recovery Appliance.
<code>racli add vpc_user</code>	Creates a new Virtual Private Catalog (VPC) user account in the Recovery Appliance schema.
<code>racli alter</code>	Modifies the properties on objects supported in the Recovery Appliance schema.
<code>racli alter admin_user</code>	Changes password for a <code>admin_user</code> system user.
<code>racli alter adr_parameter</code>	Changes Automatic Diagnostic Repository (ADR) parameters (<code>LOG_FILE_SIZE</code> and <code>LOG_FILE_NUM</code>).
<code>racli alter cloud_key</code>	Modifies the information associated with a cloud key.
<code>racli alter cloud_location</code>	Modifies the properties of cloud storage supported on the Recovery Appliance.
<code>racli alter db_user</code>	Changes password for a database user.
<code>racli alter keystore</code>	Creates a TDE master key for the Recovery Appliance, generates new keys, or changes the password for the keystore on the Recovery Appliance.
<code>racli alter network</code>	Modifies the Recovery Appliance protocol with TLS enabled only disabled.
<code>racli alter parameter</code>	Modifies the specified parameter whether it should be included or ignored as a check.
<code>racli alter passphrase</code>	Modifies password used to access the Recovery Appliance.
<code>racli alter rasys_user</code>	Modifies the Recovery Appliance rasys user account password and management of the password.
<code>racli alter vpc_user</code>	Modifies an existing Recovery Appliance user account and changes its password.
<code>racli approve</code>	Approves the object in the Recovery Appliance schema.
<code>racli approve request</code>	Approves a given request identifier and gives that vote to the quorum.

Command	Description
racli configure	Configure an object from operations performed by the Recovery Appliance.
racli configure export	Configures an NFS path for the Recovery Appliance bundle to export to.
racli create	Creates an object in the Recovery Appliance.
racli create certificate	Creates a certificate used for HTTPS operations.
racli delete	Deletes an object from the Recovery Appliance.
racli delete protection_policy	Deletes a protection policy from the Recovery Appliance.
racli delete replication_server	Deletes a replication server.
racli deny	Denies the object in the Recovery Appliance schema.
racli deny request	Denies a given request identifier and gives that vote to the quorum.
racli disable	Disables an object from operations performed by the Recovery Appliance.
racli disable aide	Disables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.
racli disable archived_log_grouping	Disables archived log grouping on the Recovery Appliance.
racli disable check	Disables a specific check object from check run operations performed on the Recovery Appliance.
racli disable cloud_location	Disables a cloud storage object for operations performed by the Recovery Appliance.
racli disable keystore	Disables the keystore.
racli disable nfs	Disables the specified Network File System (NFS).
racli disable rasys_user	Disables access to the rasys user.
racli disable root_access	Disables sudo to root access for user.
racli disable rpc	Disables remote procedure call (RPC) on all compute nodes.
racli disable ssh	Disables external ssh access to the system for 'root' and 'oracle' users.
racli disable sys_remote_access	Disables remote access to the <code>sys</code> database root user.
racli enable	Enables an object for operations performed by the Recovery Appliance.
racli enable acfs	Enables the specified Automatic Storage Management Cluster (ACFS).
racli enable aide	Enables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.
racli enable archived_log_grouping	Enables archived log grouping on the Recovery Appliance.
racli enable check	Enables a check on the Recovery Appliance.
racli enable cloud_location	Enables a cloud storage object for operations performed by the Recovery Appliance.
racli enable keystore	Opens the keystore so that it can be used.
racli enable nfs	Enables the specified Network File System (NFS).
racli enable rasys_user	Enables access to the rasys user.
racli enable root_access	Enables access to the root user.

Command	Description
<code>racli enable rpc</code>	Enables remote procedure call (RPC) on all compute nodes.
<code>racli enable ssh</code>	Enables external ssh access for 'root' and 'oracle' users.
<code>racli enable sys_remote_access</code>	Enables access to the <code>sys</code> database root user.
<code>racli export appliance</code>	.
<code>racli install</code>	Installs software.
<code>racli install appliance</code>	Installs software to the Recovery Appliance software.
<code>racli install okv_endpoint</code>	Installs the software required for Oracle Key Vault (OKV) end points on the Recovery Appliance.
<code>racli list</code>	Lists information about objects given in the Recovery Appliance schema.
<code>racli list admin_user</code>	List <code>admin_user</code> system users who can execute <code>racli</code> commands.
<code>racli list certificate</code>	Lists details of all HTTPS certificate on the Recovery Appliance.
<code>racli list check</code>	List checks available on the Recovery Appliance.
<code>racli list cloud_key</code>	List details in short (default) or long format about a specified cloud key or all cloud keys.
<code>racli list cloud_location</code>	Lists details about the cloud storage known to the Recovery Appliance.
<code>racli list cloud_user</code>	List details about the specified cloud user, or all cloud users.
<code>racli list compute_nodes</code>	Lists the names of the compute nodes available on the Recovery Appliance.
<code>racli list db_user</code>	List details about the specified database user, or all database users.
<code>racli list env</code>	Lists env parameters based on type.
<code>racli list export</code>	Lists all the Recovery Appliance export bundles.
<code>racli list grid_home</code>	Lists details about <code>grid_home</code> .
<code>racli list host</code>	Lists details about all hosts available to the Recovery Appliance.
<code>racli list ip_addr</code>	Lists details about all IP addresses available to the Recovery Appliance.
<code>racli list listener</code>	Lists details of an individual or all listeners on the Recovery Appliance.
<code>racli list migration_server</code>	List migration server.
<code>racli list nfs</code>	Lists details about one or more Network File System mounts.
<code>racli list okv_endpoint</code>	Lists details about the OKV end points known to the Recovery Appliance.
<code>racli list oracle_home</code>	Lists details about <code>oracle_home</code> .
<code>racli list parameter</code>	Lists the names and value of the init parameters for the Recovery Appliance.
<code>racli list protected_db</code>	List protected databases.
<code>racli list remote_syslog</code>	List Sys Log configuration files.
<code>racli list san</code>	Lists subject alternate name (SAN).
<code>racli list scan</code>	Lists details of scan on the Recovery Appliance.

Command	Description
<code>racli list scan_listener</code>	Lists details of scan_listener on the Recovery Appliance.
<code>racli list storage_cells</code>	Lists names of storage cells available to the Recovery Appliance.
<code>racli list tag</code>	List diagnostic tags available on the Recovery Appliance.
<code>racli remove</code>	Removes objects from the Recovery Appliance schema.
<code>racli remove acfs</code>	Removes the specified Automatic Storage Management Cluster (ACFS) mount from the Recovery Appliance schema.
<code>racli remove admin_user</code>	Removes a <code>admin_user</code> system user from the <code>system passwd</code> database.
<code>racli remove certificate</code>	Removes the specified HTTPS certificate.
<code>racli remove cloud_key</code>	Removes the specified cloud key, but only if no cloud users are associated with the key.
<code>racli remove cloud_location</code>	Removes the specified cloud storage from the Recovery Appliance, unless the cloud archive job template is tied to the library.
<code>racli remove cloud_user</code>	Removes a cloud user.
<code>racli remove db_user</code>	Removes a database user.
<code>racli remove grid_home</code>	Removes the specified grid home.
<code>racli remove host</code>	Removes the specified host.
<code>racli remove keystore</code>	Removes the specified keystore from the Recovery Appliance schema.
<code>racli remove nfs</code>	Removes the specified Network File System (NFS) mount from the Recovery Appliance schema.
<code>racli remove oracle_home</code>	Removes the specified oracle home.
<code>racli remove osb_pieces</code>	Removes all pieces that are no longer in the ZDLRA catalog from the OSB catalog.
<code>racli remove remote_syslog</code>	Removes from the Recovery Appliance schema the ability to send the system log files to a remote location.
<code>racli remove tag</code>	Removes the specified diagnostic TAG from all nodes of the Recovery Appliance.
<code>racli remove vpc_user</code>	Removes the specified user account from the Recovery Appliance metadata schema.
<code>racli run</code>	Performs an action for an object to the Recovery Appliance.
<code>racli run check</code>	Runs the enabled checks for the Recovery Appliance.
<code>racli run diagnostics</code>	Runs and collects the required diagnostics on the Recovery Appliance.
<code>racli run preinstall</code>	Execute preinstallation tasks.
<code>racli start</code>	The action to start a specified service (object).
<code>racli start appliance</code>	Starts all services on the Recovery Appliance.
<code>racli start nfs</code>	Start NFS services of the Recovery Appliance.
<code>racli status</code>	Check status of objects associated with the Recovery Appliance.

Command	Description
<code>racli status aide</code>	Returns the status of the Advanced Intrusion Detection Environment (AIDE).
<code>racli status appliance</code>	Returns a summary of the status of the appliance.
<code>racli status archived_log_grouping</code>	Returns if archived log grouping is enabled or disabled, if enabled it returns the group log count.
<code>racli status cloud_location</code>	Lists the status of a cloud storage object for operations performed by the Recovery Appliance.
<code>racli status compliance</code>	Lists the status (either PASS or FAIL) of all compliance-ready checks.
<code>racli status disk_group</code>	Verifies that the rebalance operation completed after storage cells or storage disks are added to the ASM disk group.
<code>racli status keystore</code>	Returns the status of the keystore in use by the Recovery Appliance.
<code>racli status listener</code>	Returns the status of an individual listener or all listeners on the Recovery Appliance.
<code>racli status nfs</code>	Returns the status of the Network File Service (NFS).
<code>racli status okv_endpoint</code>	Returns the status of the OKV_ENDPOINTS associated with the Recovery Appliance.
<code>racli status request</code>	Shows status of all scan on the Recovery Appliance.
<code>racli status scan</code>	Shows status of all scan on the Recovery Appliance.
<code>racli status scan_listener</code>	Shows status of all scan_listeners on the Recovery Appliance.
<code>racli stop</code>	The action to stop a specified service (object).
<code>racli stop appliance</code>	Stops all services on the Recovery Appliance.
<code>racli stop nfs</code>	Stop NFS services of the Recovery Appliance.
<code>racli uninstall</code>	Backs out a portion of the installation.
<code>racli uninstall appliance</code>	Backout a step during install.
<code>racli uninstall okv_endpoint</code>	Uninstalls the specified OKV end point.
<code>racli update</code>	Update a component of the Recovery Appliance.
<code>racli update aide</code>	Updates the Advanced Intrusion Detection Environment (AIDE).
<code>racli update appliance</code>	Updates the ZDLRA.
<code>racli update parameter</code>	Update init parameters in the Recovery Appliance.
<code>racli update ssh_keys</code>	TBD.
<code>racli update ssh_known_hosts</code>	Maintains a system-wide configuration file <code>/etc/ssh/ssh_known_hosts</code> with host keys for all nodes in the Recovery Appliance cluster.
<code>racli update storage_location</code>	Update Recovery Appliance pool data file with the most current number of storage cells.
<code>racli update support_tools</code>	Updates the AHF and Exachk versions using the staged binaries zip files.
<code>racli version</code>	Returns the ZDLRA software release version in use..

racli add

Adds an object to the Recovery Appliance.

Options

Details

Adds an object to the Recovery Appliance.

racli add admin_user

Adds a specified `ADMIN_USER` system user to the RACLI admin group.

Syntax

```
racli add admin_user --user_name=USER_NAME [--user_uid=USER_ID --  
user_gid=GROUP_ID]
```

Options

--user_name

System user name to add to RACLI admin group.

--user_uid

Set the user identifier for the newly created admin user. Value must be ≥ 1003 .
During the installation of RA 19.x or later, you can define the `raadmin uid` with `ra_preinstall.pl`.

--user_gid

Set the initial login group identifier for the newly created admin user. A group number must refer to an already existing group. Value must be ≥ 1003 .
During the installation of RA 21.1 or later, you can define the `gid` with `ra_preinstall.pl`.

Details

This system user can execute RACLI commands. User is created if not already in the `passwd` database. This command prompts you to enter a user password.

racli add certificate

Adds https certification on the Recovery Appliance.

Syntax

```
racli add certificate { [--signed_cert=VALUE --trusted_cert=VALUE] | [--  
self_signed] }
```

Options

--self_signed

Specifies that the Recovery Appliance should generate the certificates itself. This mode is for testing.

--signed_cert

Specifies the full path of Recovery Appliance certificate to be added.

--trusted_cert

Specifies the full path of root certificate to be added.

Details

When the current certificate is about to expire or has expired, use this command to upload a new valid certificate to the Recovery Appliance wallet with the `--signed_cert` and `--trusted_cert` arguments.

racli add cloud_key

Generates the public and private key files required for OCI access.

Syntax

```
racli add cloud_key --key_name=KEY_NAME [--fingerprint=PUBFINGERPRINT --  
pvt_key_path=PVTKEYFILE]
```

Options

--key_name

The name of the `cloud_key` object to create.

--fingerprint

Specifies the name of the fingerprint for the public key that is paired with the specified private key. The fingerprint is used to authenticate the request.

--pvt_key_path

The name of a private key used to sign BMC API requests. The key file must be in PEM format.

Details

Generates the public and private key files required for OCI access. It also associates fingerprint and pre-existing private key files.

The first syntax form generates the public and private keys required for OCI access. The user takes the generated public key to get fingerprint from OCI console.

```
racli add cloud_key --key_name=KEY_NAME
```

The second syntax form associates fingerprint and private key files that already exist.

```
racli add cloud_key --key_name=KEY_NAME --fingerprint=PUBFINGERPRINT --  
pvt_key_path=PVTKEYFILE
```

racli add cloud_location

Adds a cloud storage object to the Recovery Appliance.

Syntax

```
racli add cloud_location --cloud_user=CLOUD_USER_NAME --host=HOST_URL --  
bucket=OCI_BUCKET_NAME [--enable_archive | --disable_archive] [--  
archive_after_backup=NUMBER:{DAYS|YEARS} --streams=NUMBER --  
proxy_host=HTTP_SERVER --proxy_port=HTTP_PORT --proxy_id=HTTP_USER --  
proxy_pass=HTTP_PASS --import_all_trustcert=X509_CERT_PATH --  
retain_after_restore=NUMBER:HOURS] [-guaranteed={yes|no}] [--immutable --  
temp_metadata_bucket=METADATA_BUCKET_NAME]
```

Options

--archive_after_backup

The number in days or years for period of time before backups are moved to archive state. It takes the form "--archive_after_backup=NUMBER:DAYS" or "--archive_after_backup=NUMBER:YEARS", where NUMBER is a positive integer. If not specified, the default is 2 days.

--bucket

The name of the bucket where the backup will go. Note that the install tool will create the specified bucket if it does not exist.

--cloud_user

Previously configured cloud_user object with all authentication requirements.

--disable_archive

Disables archival of backup data.

--enable_archive

Enables archival of backup data after the time interval specified by the --archive_after_backup parameter. If not specified, by default the archival will be disabled.

--host

Host name for the Oracle Bare Metal Cloud account.

--import_all_trustcert

Import all X509 certificates from Java truststore.

--streams

The number of streams to be associated with this specific cloud_location. Directly translates to the number of sbt_attrib sets created.

--proxy_host

HTTP proxy server. Used for OCI-C support.

--proxy_id

HTTP proxy server username. Used for OCI and stored as part of the alias_http_proxy in location wallet.

--proxy_pass

HTTP proxy server password. Stored as part of the alias_http_proxy in location wallet.

--proxy_port

HTTP proxy server port. Used for OCI-C support.

--retain_after_restore

Specifies the number of hours to retain information after a restore.

It takes the form "--retain_after_restore=NUMBER:HOURS". It has a maximum value of "240:HOURS".

--guaranteed

If --guaranteed=yes, this cloud location will require a copy of a backup piece before it can be deleted. For virtual backups, only one of the Level 0 or Level 1 copies of a backup piece are required.

--immutable

This makes the bucket being added immutable. Must be used with the --temp_metadata_bucket option.

--temp_metadata_bucket

Used for immutable buckets.

Specifies the metadata bucket name for immutable cloud location setup. The metadata bucket must be a mutable bucket (i.e. non immutable) and must exist. It can be created using the OCI console. Each metadata bucket must be unique. You cannot use the same metadata bucket for different immutable buckets.

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 to the Recovery Appliance schema.

```
racli add cloud_location --cloud_user=samu ---host=https://cloud_uri/oci -  
bucket=sam_backup
```

racli add cloud_user

Adds a cloud user to the Recovery Appliance schema.

Syntax

```
racli add cloud_user --user_name=CLOUD_USER_NAME --key_name=CLOUD_KEY_NAME --  
user_ocid=UOCID --tenancy_ocid=TOCID [--compartment_ocid=COCID]
```

Options**--user_name**

The name to be associated with this particular cloud user.

--key_name

The specific cloud key to be associated with this cloud user.

--tenancy_ocid

The tenancy OCID for the Oracle Bare Metal Cloud account.

--user_ocid

The user OCID for the Oracle Bare Metal Cloud account.

--compartment_ocid

The compartment OCID within the tenancy of the Oracle Bare Metal Cloud Account.

Details

Adds a cloud user to the Recovery Appliance schema.

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.

Details

Adds a user credential to the wallet of the Recovery Appliance.

Examples

This example adds a credetinal 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:
```


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

Adds a wallet object to the Recovery Appliance. 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:
```

racli add db_user

Adds a database administration user (`db_user`) who can manage and configure the Recovery Appliance.

Syntax

```
racli add db_user --user_name=USER_NAME --user_type=admin|replication|monitor|  
vpc [ --insecure=VALUE ]  
[ --password_hash ] [ --wallet=VALUE ]
```

Options

--user_name

User name for the database user.

--user_type

Can be `admin`, `replication`, `monitor`, or `vpc`. Admin user can manage and configure the Recovery Appliance. Replication user has permissions to copy and retrieve associate backup files between Recovery Appliances, and associated wallet access. Monitor user has read-only access to Enterprise Manager. VPC user is the virtual private catalog user.

--insecure

This option can only be used with the VPC user_type.

VPC users created with this option will have the DEFAULT profile and will not have password rollover support.

--ip_list

This option is only applicable to cloud systems.

A list of comma separated IP addresses that the database user can connect from

Using this option will automatically assign a database service to the user. The user can then connect to the database from any of the allowed IPs using the name of their assigned service.

--password_hash

Specifies whether a hash password value is provided.

--wallet

Specifies the wallet that stores the password.

racli add host

Adds a host to those available on the Recovery Appliance.

Syntax

```
racli add host --host_name=HOSTNAME --ip=IPADDRESS --domain_name=DOMAINNAME
```

Options**--host_name**

Specifies the name of the host to be added.

--ip

Specifies the IP address for the host being added.

--domain_name

Specifies the domain name for the host being added.

racli add keystore

Adds a keystore to the Recovery Appliance.

Syntax

```
racli add keystore --type=[HSM|SW] [--restart_db]
```

Options**--type**

Specifies the type of keystore to be created. If not specified, default is HSM.

--restart_db

Restarts the database stack after changes to the keystore to ensure they are honored at service restart.

Details

Adds a keystore to the Recovery Appliance. 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
```

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'

Details

Specifies that the Network File System (NFS) should be added to the Recovery Appliance schema.

racli add rasys_wallet

Creates the rasys user wallet.

Syntax

```
racli add rasys_wallet
```

Options

Details

Creates the rasys user wallet.

racli add remote_syslog

Creates a configuration file to export log files to a remote server.

Syntax

```
racli add remote_syslog [--dest=DESTNAME] [--port=PORTNUM] [--  
config_name=CNAME]
```

Options

--dest

Specifies the remote server that will receive log files.

--port

Specifies the port number of remote server that will receive log files.

--config_name

Specifies the name of the configuration file.

racli add storage_cells

Prepares new Recovery Appliance storage cells for storage location expansion.

Syntax

```
racli add storage_cells --cells=cell1,cell2,...,celln [--priority=high]
```

Options

--cells

The comma-separated list of storage cell names to prepare.

--priority

Indicates the priority level for the re-balance. The default is low. Low = 4. High = 32. This command should only be run after re-balance is complete.

Details

Prepares new Recovery Appliance storage cells for storage location expansion.

- 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 --cells=cell11,cell1,...celln2
```

racli add storage_disks

Adds disks with the newly added storage cells to the ASM disk groups for storage location expansion.

DEPRECATED: The functionality of this command is now included in [racli add storage_cells](#).

Syntax

```
racli add storage_disks --cells=cell11,cell12,...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

Adds disks with the newly added storage cells to the ASM disk groups for storage location expansion.

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 --cells=cell11,cell12,...celln
```

racli add tag

Specifies the diagnostic TAG to add to all nodes of the Recovery Appliance.

Syntax

```
racli add tag --tag=myTagName --level=myLevel
```

Options

--tag

Name of the new tag to be added.

--level

The level of the ANDed diagnostics to include in the tag, available levels are:

level 1 System activity report.

level 2 RA dump txt file.

level 4 RA dump datapump.

level 8 ORDERING_WAIT analysis.

level 16 Corruption.

level 32 Gather racli logs from each node.

level 64 Get RA version information from each node.

level 128 Get store and config files from each node.

level 256 Get root and oracle history from each node.

level 512 Get network information from each node.

level 1024 Run obtool gathering useful information.

level 2048 Get osb logs from each node.

level 4096 Get /usr/etc/ob/xcr from each node.

level 8192 Get backup/admin/state from each node.

level 16384 Get all nodes ASM and ZDLRA alert.logs.

The levels can be ANDed together to create a collection.

Details

Specifies the diagnostic TAG to add to all nodes of the Recovery Appliance

Examples

This example adds a new TAG to collect racli logs and Recovery Appliance version information from each node:

```
racli add tag --tag=new_tag --level=196608
```

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 (deprecated)
```

Options

user_name

The name of the Recovery Appliance catalog user account to add.

Details

(deprecated) This command is deprecated. Please use `racli add db_user --user_name=myVpcUserName --type=vpc` to add a vpc user.

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.
- The password on the `DEFAULT` database profile will expire. To avoid the password expiring, change the profile associated with the VPC user.

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:
```

racli alter

Modifies the properties on objects supported in the Recovery Appliance schema.

Options

Details

Modifies the properties on objects supported in the Recovery Appliance schema.

racli alter admin_user

Changes password for a `admin_user` system user.

Syntax

```
racli alter admin_user --user_name=USER_NAME
```

Options

--user_name

System user name to alter.

Details

Prompts to enter a user password.

racli alter adr_parameter

Changes Automatic Diagnostic Repository (ADR) parameters (LOG_FILE_SIZE and LOG_FILE_NUM).

Syntax

```
racli alter adr_parameter --log_file_size=LOGS_SIZE --log_file_num=LOGS_NUMBER
```

Options

--log_file_num

Number of log file segments. Default value is 10.

--log_file_size

Size of each log file segment. Default value is 300MB.

racli alter cloud_key

Modifies the information associated with a cloud key.

Syntax

```
racli alter cloud_key --key_name=KEY_NAME [--fingerprint=PUBFINGERPRINT --  
pvt_key_path=PVTKEYFILE]
```

Options

--key_name

The name of the cloud key object.

--fingerprint

Specifies the name of the fingerprint for the public key that is paired with the specified private key. The fingerprint is used to authenticate the request.

--pvt_key_path

The name of a private key used to sign requests. The key file must be in PEM format.

Details

Modifies the information associated with a cloud key.

A cloud key has to have a fingerprint associated with it in order to later be associated with a user.

racli alter cloud_location

Modifies the properties of cloud storage supported on the Recovery Appliance.

Syntax

```
racli alter cloud_location --location_name=CLOUD_LOCATION_NAME [--  
archive_after_backup=NUMBER:{DAYS|YEARS}] [--streams=NUMBER] [--  
proxy_host=HTTP_SERVER] [--proxy_id=HTTP_USER] [--proxy_pass=HTTP_PASS] [--  
proxy_port=HTTP_PORT] [--retain_after_restore=NUMBER:HOURS]
```

Options

--archive_after_backup

The number in days or years for period of time before backups are moved to archive state. It takes the form "--archive_after_backup=NUMBER:DAYS" or "--archive_after_backup=NUMBER:YEARS", where NUMBER is a positive integer. If not specified, the default is 2 days.

--location_name

The name of the cloud location where the backup will go. The cloud location correlates to OCI object store bucket. Note that the install tool will create the specified bucket if it does not exist.

--proxy_host

HTTP proxy host.

--proxy_id

HTTP proxy server username. Used for OCI and stored as part of the `alias_http_proxy` in location wallet.

--proxy_pass

HTTP proxy server password. Stored as part of the `alias_http_proxy` in location wallet.

--proxy_port

HTTP proxy server port. Used for OCI-C support.

--streams

The number of streams to be associated with this specific cloud_location. Directly translates to the number of `sbt_attrb` sets created.

Integer that specifies number of streams to associate with this specific cloud_location.

--retain_after_restore

Specifies the number of hours to retain information after a restore. It takes the form "--retain_after_restore=NUMBER:HOURS". It has a maximum value of "240:HOURS".

Details

Modifies the properties of cloud storage supported on the Recovery Appliance. This is equivalent to the `sbt_library` used for copy-to-tape. This `sbt_library` is specific to archive-to-cloud.

racli alter db_user

Changes password for a database user.

Syntax

```
racli alter db_user --user_name=USER_NAME --user_type=admin|replication|  
monitor|vpc
```

Options

--user_name

database vault admin user name to alter.

--user_type

Can be admin, replication, monitor, or vpc. Admin user can manage and configure the Recovery Appliance. Replication user has permissions to copy and retrieve associate backup files between Recovery Appliances, and associated wallet access. Monitor user has read-only access to Enterprise Manager. VPC user is the virtual private catalog user.

--enable_stig

This option can only be used with the vpc user type.

Changes the vpc user to the RA_STIG profile that has password rollover support.

--expire_pw_rollover

End the password rollover period for this database user immediately. The database user will not be able to login with the old password anymore.

This option will not prompt the user for a new password.

Details

Prompts to enter a user password.

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 { --inititalize_key | --restart_db | --password | --rekey }
```

Options

--rekey

Generates new TDE master encryption key (rekey) for the Recovery Appliance. Keys for protected databases are created using `dmbs_ra.key_rekey`.

--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.

Details

Creates a TDE master key for the Recovery Appliance, generates new keys, or changes the password for the keystore on the Recovery Appliance.

Examples

The following example restarts the database stack following a keystore change.

```
racli alter keystore --restart_db
```

racli alter network

Modifies the Recovery Appliance protocol with TLS enabled|only|disabled.

Syntax

```
racli alter network --network_type=VALUE [--silent]
```

Options

Details

--service

Indicate the service being modified on the system. Valid value is "ra_server". Cannot be used of --network_type or its arguments.

--network_type

Indicate network type on the system. Cannot be used of --service or its arguments.

--encrypt

Specifies TLS encryption status on the system: "only" means HTTPS encryption; "enable" means dual HTTPS and HTTP; and "disable" means HTTP.

--http_port

Specifies HTTP port number to use. Default port is 8001.

--https_port

Specifies HTTPS port number to use. Default port is 8005.

--tcp_port

Specifies the TCP port number to use. Default port is 1521.

--tcps_port

Specifies TCPS port number to use. Default port is 2484

--rep_tcps_port

Specifies the replication TCPS port number to use. Default port is 2485.

--rep_tcp_port

Specifies the replication TCP port number to use. Default port is 1522.

--silent

When present

Refer to the command line syntax regarding which options can be used at the same time.

The optional parameters `--network_type` and `--service` cannot be used together, and they cannot each other's optional parameters.

When configuring HTTPS (`encrypt=[enable!only]`), optional parameters for changing the default ports for HTTP, TCP, and replication TCP are ignored. Similarly, when disabling HTTPS (`encrypt=disable`), optional parameters for changing ports for HTTPS, TCPS, and replication TCPS are ignored.

Examples

The following example configures TLS on system for dual mode HTTPS and HTTP. It changes the TCPS port number and leaves the other ports set to their default values.

```
racli alter network --service=ra_server --encrypt=enable --tcps_port=2848
```

racli alter parameter

Modifies the specified parameter whether it should be included or ignored as a check.

Syntax

```
racli alter parameter --type=db --parameter_name=NAME --ignore_check={ false | true } [--force]
```

Options

`--type`

The type of the init parameter to be checked. Presently only "db" is supported.

`--parameter_name`

The name of the init parameter to be checked.

`--ignore_check`

Specifies whether to ignore a parameter (true) or to include the parameter (false) in a check of the database.

`--force`

Certain parameters are deemed important for database checks and are managed in an exclusion list to always be checked. The `--force` option bypasses the exclusion list so that the parameter can be ignored.

racli alter passphrase

.

Syntax

```
racli alter passphrase
```

Options

Modifies password used to access the Recovery Appliance.

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.

Details

Modifies password used to access the Recovery Appliance.

- 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.

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

--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

Details

Modifies the Recovery Appliance rasys user account password and management of the password.

Examples

The following example changes the password for the *rasys* 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.

racli alter vpc_user

Modifies an existing Recovery Appliance user account and changes its password.

Syntax

```
racli alter vpc_user --user_name=myVpcUserName (deprecated)
```

Options

Details

(deprecated) This command is deprecated. Please use `racli alter db_user --user_name=myVpcUserName --type=vpc` to alter a vpc user.

--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:
```

racli approve

Approves the object in the Recovery Appliance schema.

Options

racli approve request

Approves a given request identifier and gives that vote to the quorum.

Syntax

```
racli approve request --request_id=VALUE
```

Options

--request_id

Specifies the request identifier to be approved.

racli configure

Configure an object from operations performed by the Recovery Appliance.

Options

racli configure export

Configures an NFS path for the RA bundle to export to.

Syntax

```
racli configure export --nfs_path=/ranfs/smaug/ra_export
```

Options

The NFS must be added using "racli add nfs" prior to running this command.

RA export bundle will be exported to this path in addition to the default location.

--nfs_path

Specifies the path to the NFS share

racli create

Create an object from operations performed by the Recovery Appliance.

Options

racli create certificate

Creates a certificate used for HTTPS operations.

Syntax

```
racli create certificate [--signed_cert_valid=VALUE] [--  
trusted_cert_valid=VALUE] --country=VALUE --state=VALUE --location=VALUE --  
organization=VALUE --organization_unit=VALUE --email_address=VALUE [--  
signed_cert_only | --trusted_cert_only]
```

Options

--signed_cert_valid

Specifies the number of days that the signed certificate remains valid. The default is 365 days (1 year).

Note: **--signed_cert_valid** cannot be larger than **--trusted_cert_valid**.

--trusted_cert_valid

Specifies the number of days that the trust certificate remains valid. The default is 3650 days (10 years).

--country
Specifies the country of origin for the certificate.

--state
Specifies the state of origin for the certificate.

--location
Specifies the location of origin for the certificate.

--organization
Specifies the organization responsible for the certificate.

--organization_unit
Specifies the unit within the organization responsible for the certificate.

--email_address
Specifies the email address for the owner of the certificate.

--signed_cert_only
Enables operation involving only validation of signed certificates, not of trusted certificates.

--trusted_cert_only
Enables operation involving only validation of trusted certificates, not of signed certificates.

racli deny

Denies the object in the Recovery Appliance schema.

Options

racli deny request

Denies a given request identifier and gives that vote to the quorum.

Syntax

```
racli deny request --request_id=VALUE
```

Options

--request_id
Specifies the request identifier to be denied.

racli disable

Disables an object from operations performed by the Recovery Appliance.

Options

Details

Disables an object from operations performed by the Recovery Appliance. Object remains known to the Recovery Appliance for a quick enable later.

racli disable aide

Disables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.

Syntax

```
racli disable aide
```

Options

Details

AIDE is an Exadata utility that creates a database of files on the system and tracks changes. AIDE ensures file integrity and detects system intrusions. AIDE needs to be disabled during system installation and upgrades, and then afterward enabled again.

racli disable archived_log_grouping

Disables archived log grouping on the Recovery Appliance.

Syntax

```
racli disable archived_log_grouping
```

Options

Details

This sets the max number of archive logs grouped into a single backup before being written to tape to 1.

racli disable check

Disables a specific check object from check run operations performed on the Recovery Appliance.

Syntax

```
racli disable check { --check_name=CHECK_NAME }
```

Options

--check_name

Specifies the name of the check to be disabled.

racli disable cloud_location

Disables a cloud storage object for operations performed by the Recovery Appliance.

Syntax

```
racli disable cloud_location [ --location_name=CLOUD_LOCATION_NAME ]
```

Options

--location_name

Specifies the cloud location name to be disabled.

Details

Disables a cloud storage object for operations performed by the Recovery Appliance. The enable action defaults to **all** for cloud_location unless specified. This translates to 'pause' for sbt_library in a ready state.

racli disable keystore

Disables the keystore.

Syntax

```
racli disable keystore
```

Options

Details

Disables the keystore.

racli disable nfs

.

Syntax

```
racli disable nfs --name=myMountName
```

Options

Disables the specified Network File System (NFS).

Details

--name

Name of the NFS mount object.

racli disable rasys_user

Disables access to the `rasys` user.

Syntax

```
racli disable rasys_user [--request_id=req_id]
```

Options

If compliance mode is enabled, a quorum is needed to disable access.

If compliance mode is not enabled, quorum is not needed.

--request_id

The request id for the `rasys` disable request.

racli disable root_access

Disables `sudo` to root access for user.

Syntax

```
racli disable root_access [--request_id=req_id]
```

Options

If compliance mode is enabled, a quorum is needed to disable access.

If compliance mode is not enabled, quorum is not needed.

--request_id

The request id for the disable request.

Details

This command disables `root` access for `root`, `oracle`, and `raadmin`.

This command is `local_only`. This can be run on the local node, but cannot be run on other nodes.

racli disable rpc

Disables remote procedure call (RPC) on all compute nodes.

Syntax

```
racli disable rpc
```

Options**Details**

If OSB is not installed and there are no NFS mounts, RPC is disabled by default.

racli disable ssh

Disables external `ssh` access for the `root` and `oracle` users.

Syntax

```
racli disable ssh
```

racli disable sys_remote_access

Disables access to the `sys` database root user.

Syntax

```
racli disable sys_remote_access [--request_id=req_id]
```

Options

If compliance mode is enabled, a quorum is needed to disable access.

If compliance mode is not enabled, quorum is not needed.

After this command is run, the Recovery Appliance needs to be restarted to complete the disable operation.

--request_id

The request identifier for the `sys` remote access disable request.

racli enable

.

Options

Enables an object for operations performed by the Recovery Appliance.

Details

Enables an object for operations performed by the Recovery Appliance.

racli enable acfs

Enables the specified Automatic Storage Management Cluster (ACFS).

Syntax

```
racli enable acfs --name=myMountName
```

Options

--name

Name of the ACFS mount object.

Details

Enables the specified Automatic Storage Management Cluster (ACFS).

racli enable aide

Enables the Advanced Intrusion Detection Environment (AIDE) on the Recovery Appliance.

Syntax

```
racli enable aide
```

Options

Details

AIDE is an Exadata utility that creates a database of files on the system and tracks changes. AIDE ensures file integrity and detects system intrusions. AIDE needs to be disabled during system installation and upgrades, and then afterward enabled again.

racli enable archived_log_grouping

Enables archived log grouping on the Recovery Appliance.

Syntax

```
racli enable archived_log_grouping
```

Options

Details

This sets the max number of archive logs grouped into a single backup before being written to tape to 128.

racli enable check

Enables a check on the Recovery Appliance.

Syntax

```
racli enable check { --check_name=CHECK_NAME };
```

Options

--check_name

Indicates the name of the check to be enabled on the Recovery Appliance.

Details

Enables a check on the Recovery Appliance.

racli enable cloud_location

Enables a cloud storage object for operations performed by the Recovery Appliance.

Syntax

```
racli enable cloud_location [ --location_name=CLOUD_LOCATION_NAME ]
```

Options

--location_name

Specifies the cloud location name to be enabled.

Details

Enables a cloud storage object for operations performed by the Recovery Appliance. The enable action defaults to **all** for cloud_location unless specified. This translates to 'resume' for sbt_library in a paused state.

racli enable keystore

Opens the keystore so that it can be used.

Syntax

```
racli enable keystore
```

Options

Details

Opens the keystore so that it can be used. This is the OKV endpoint keystore, also known as the OKV shared wallet.

racli enable nfs

Enables the specified Network File System (NFS).

Syntax

```
racli enable nfs --name=myMountName
```

Options

--name

Name of the NFS mount object.

Details

Enables the specified Network File System (NFS).

racli enable rasys_user

Enables access to the `rasys` user.

Syntax

```
racli enable rasys_user [--request_id=req_id] [--expire=HOURS]
```

Options

If compliance mode is enabled, a quorum is needed to gain access.

If compliance mode is not enabled, quorum is not needed.

--expire

Time in hours for how long `rasys` remains enabled. Default is 3 hours.

--request_id

The request id for the access request.

This command will prompt for a new `rasys` password when request is approved and quorum is reached.

racli enable rpc

Enables remote procedure call (RPC) on all compute nodes.

Syntax

```
racli enable rpc
```

racli enable ssh

Enables external `ssh` access for the `root` and `oracle` users.

Syntax

```
racli enable ssh
```

racli enable sys_remote_access

Enables remote access to the `sys` database root user.

Syntax

```
racli enable sys_remote_access [--request_id=req_id]
```

Options

If compliance mode is enabled, a quorum is needed to disable access.

If compliance mode is not enabled, quorum is not needed.

--request_id

The request id for the `sys` remote access enable request.

racli export

Used for exporting data.

Options

Details

Used for exporting data.

racli export appliance

Creates an export bundle for the Recovery Appliance.

Syntax

```
racli export appliance
```

Options

Details

Creates an export bundle for the Recovery Appliance.

racli install

Installs software.

Options

Details

Installs software.

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. Prompts for the `raadmin` password. Starts all Recovery Appliance services installed in the previous steps. Removes remote `root` and `oracle` access during this step.

Details

Installs software to the Recovery Appliance software.

Before you run `racli install appliance`:

- Ensure that all current installation patches are stated 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
```

racli install okv_endpoint

Installs the software required for Oracle Key Vault (OKV) end points on the Recovery Appliance.

Syntax

```
racli install okv_endpoint [--okv_client_location=PATH]
```

Options

--okv_client_location

Specifies the location of the OKV client jar file to use during the installation. It must be available on all compute nodes in the cluster.

Details

Installs the software required for Oracle Key Vault (OKV) end points on the Recovery Appliance. This is run after the 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]
```

racli list

Lists information about objects given in the Recovery Appliance schema.

Options

Details

Lists information about objects given in the Recovery Appliance schema.

racli list admin_user

List `admin_user` system users who can execute `racli` commands.

Syntax

```
racli list admin_user
```

Options

Details

User has to be a member of RACLI admin group, which is `raadmin` by default.

racli list certificate

Lists details of all HTTPS certificate on the Recovery Appliance.

Syntax

```
racli list certificate
```

Options

Details

Use this command to obtain certificate serial numbers needed for "racli remove certificate" operations.

racli list check

List checks available on the Recovery Appliance.

Syntax

```
racli list check { --status=STATUS | --all } [--verbose]
```

Options

--status

List checks by STATUS. Valid STATUS are `enabled` | `disabled`.

--all

List names of all checks.

--verbose

List details of specified checks.

Details

List checks available on the Recovery Appliance.

racli list cloud_key

List details in short (default) or long format about a specified cloud key or all cloud keys.

Syntax

```
racli list cloud_key { --key_name=KEY_NAME | --all } [--long] [--fingerprint]
```

Options

--key_name

Specifies a key name associated with the cloud storage to list details.

--fingerprint

Displays the fingerprint associated with the cloud key.

--all

Lists details about all cloud keys known to the Recovery Appliance.

--long

Lists extended details about one or all cloud keys.

Details

List details in short (default) or long format about a specified cloud key or all cloud keys.

racli list cloud_location

Lists details about the cloud storage known to the Recovery Appliance.

Syntax

```
racli list cloud_location { --location_name=CLOUD_LOCATION_NAME | --all } [--long]
```

Options

--location_name

Specifies a cloud location name to list details.

--all

Lists details about all cloud storage known to the Recovery Appliance.

--long

Lists extended details about one or all containers.

Details

Lists details about the cloud storage known to the Recovery Appliance.

Examples

The following example lists all known cloud_location objects.

```
racli list cloud_location --all
```

racli list cloud_user

List details about the specified cloud user, or all cloud users.

Syntax

```
racli list cloud_user { --user_name=USER_NAME | --all } [--long]
```

Options**--user_name**

The name user wants associated with this particular cloud user.

--all

Lists details about all cloud users known to the Recovery Appliance.

--long

Lists extended details about one or all cloud users.

Details

List details about the specified cloud user, or all cloud users.

racli list compute_nodes

Lists the names of the compute nodes available on the Recovery Appliance.

Syntax

```
racli list compute_nodes
```

Options**Details**

Lists the names of the compute nodes available on the Recovery Appliance.

racli list db_user

List details about the specified database user, or all database users.

Syntax

```
racli list db_user { --user_name=USER_NAME | --all } [--long --  
user_type=admin|monitor|vpc]
```

Options

--user_name

Lists details about the user with the specified user_name.

--all

Lists all database users known to the Recovery Appliance.

--long

Lists extended details about one or all database users.

--user_type

Filters the user results to list only the user(s) with the specified user_type.

racli list env

Lists env parameters based on type.

Syntax

```
racli list env --type=db|asm|local
```

Options

--type

The type of env which has 3 options: db, asm and local.

racli list export

Lists all the RA export bundles.

Syntax

```
racli list export { --all | --default | --nfs }
```

Options

--all

List all the export bundles in both the default /raacfs/raadmin/ra_export location and the configured NFS location, if any.

--default

List all the export bundles in the default /raacfs/raadmin/ra_export location.

--nfs

List all the export bundles in the configured NFS location, if any.

racli list grid_home

Lists details about grid_home.

Syntax

```
racli list grid_home { --active | --inactive | --removed | --all }
```

Options

--active

List details about active grid_home.

--inactive

List details about all inactive grid_home.

--removed

List details about removed grid_home.

--all

List details about all grid_home.

racli list host

Lists details about all hosts available to the Recovery Appliance.

Syntax

```
racli list host
```

Options

racli list ip_addr

Lists details about all IP addresses available to the Recovery Appliance.

Syntax

```
racli list ip_addr
```

Options

racli list listener

.

Options

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.

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.

Details

Lists 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
```


racli list oracle_home

Lists details about oracle_home.

Syntax

```
racli list oracle_home { --active | --inactive | --removed | --all }
```

Options

--active

List details about active oracle_home.

--inactive

List details about all inactive oracle_home.

--removed

List details about removed oracle_home.

--all

List details about all oracle_home.

racli list parameter

Lists the names and value of the init parameters for the Recovery Appliance.

Syntax

```
racli list parameter --type=db [--invalid | --valid | --all] [--long]
```

Options

--type

The type of the init parameter to be listed.

--invalid

Show only invalid init parameter names.

--valid

Show only valid init parameter names.

--all

Show all init parameter names.

--long

Show init parameter value as well as names.

Details

Lists the names and value of the init parameters for the Recovery Appliance.

racli list remote_syslog

List Sys Log configuration files.

Syntax

```
racli list remote_syslog [--config_name=CNAME]
```

Options

--config_name

Specifies the name of the configuration file.

racli list san

Lists subject alternate name (SAN).

Syntax

```
racli list san
```

Options

racli list scan

Lists details of scan on the Recovery Appliance.

Syntax

```
racli list scan
```

Options

racli list scan_listener

Lists details of scan_listener on the Recovery Appliance.

Syntax

```
racli list scan_listener
```

Options

racli list storage_cells

Lists names of storage cells available to the Recovery Appliance.

Syntax

```
racli list storage_cells
```

Options**Details**

Lists names of storage cells available to the Recovery Appliance.

racli list tag

List diagnostic tags available on the Recovery Appliance.

Syntax

```
racli list tag { --tag=<TAG_NAME> | --all }
```

Options**--tag**

List diagnostics performed under the specific tag.

--all

List diagnostics performed by all tags.

Details

List diagnostic tags available on the Recovery Appliance.

racli remove

Removes objects from the Recovery Appliance schema.

Options**Details**

Removes objects from the Recovery Appliance schema. Removed objects become unknown to the Recovery Appliance and cannot be quickly enabled later.

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.

Details

Removes the specified Automatic Storage Management Cluster (ACFS) mount from the Recovery Appliance schema.

racli remove admin_user

Removes a `admin_user` system user from the `system passwd` database.

Syntax

```
racli remove admin_user --user_name=USER_NAME
```

Options

--user_name

System user name to remove.

Details

User has to be a member of RACLI admin group, which is `raadmin` by default.

racli remove certificate

Removes the specified HTTPS certificate.

Syntax

```
racli remove certificate --serial_number=SERIALNUMBER
```

Options

--serial_nuber

Specifies the serial number of the HTTPS certificate to be removed.

Details

However, the certificate is not removed if TLS is enabled and the specified certificate is the last one in its `signed_cert` or `trusted_cert` type.

racli remove cloud_key

Syntax

```
racli remove cloud_key --key_name=KEY_NAME
```

Options

Removes the specified cloud key, but only if no cloud users are associated with the key.

Details

--key_name

The name of the public and private key to remove.

racli remove cloud_location

Removes the specified cloud storage from the Recovery Appliance, unless the cloud archive job template is tied to the library.

Syntax

```
racli remove cloud_location --location_name=LOCATION_NAME [--force]
```

Options

--location_name

Specifies the name of the cloud location to be removed.

--force

If a cloud storage job template is tied to the library, this option forces its removal anyway.

Details

Removes the specified cloud storage from the Recovery Appliance, unless the cloud archive job template is tied to the library.

racli remove cloud_user

Removes a cloud user.

Options

--user_name

The cloud user object name associated with this particular user.

Details

Removes a cloud user. A cloud user associated with a cloud location cannot be removed.

racli remove db_user

Removes a database user.

Syntax

```
racli remove db_user --user_name=USER_NAME --user_type=admin|monitor|vpc
```

Options

--user_name

Database user to remove.

--user_type

Type of the database user to be removed.

racli remove grid_home

Removes the specified grid home.

Syntax

```
racli remove grid_home --grid_home=inactiveGridHome
```

Options

--grid_home

Path of grid_home to be removed.

racli remove host

Removes the specified host.

Syntax

```
racli remove host --host_name=HOSTNAME
```

Options

--host_name

Indicates the name of the host to be removed.

racli remove keystore

Removes the specified keystore from the Recovery Appliance schema.

Syntax

```
racli remove keystore [--restart_db]
```

Options

--restart_db

Restarts the database stack after changes to the keystore to ensure they are honored at service restart.

Details

Removes the specified keystore from the Recovery Appliance schema.

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.

Details

Removes the specified Network File System (NFS) mount from the Recovery Appliance schema.

racli remove oracle_home

Removes the specified oracle home.

Syntax

```
racli remove oracle_home --oracle_home=inactiveOracleHome
```

Options

--oracle_home

Path of oracle_home to be removed.

racli remove osb_pieces

Removes all pieces that are no longer in the ZDLRA catalog from the OSB catalog.

Syntax

```
racli remove osb_pieces
```

Options

racli remove remote_syslog

Removes from the Recovery Appliance schema the ability to send the system log files to a remote location.

Syntax

```
racli remove remote_syslog [--config_name=CNAME]
```

Options

--config_name

Specifies the name of the configuration file.

racli remove tag

Removes the specified diagnostic TAG from all nodes of the Recovery Appliance.

Syntax

```
racli remove tag --tag=TagName
```

Options

--tag

Name of the diagnostic tag to be removed.

Details

Removes the specified diagnostic TAG from all nodes of the Recovery Appliance.

racli remove vpc_user

Removes the specified user account from the Recovery Appliance metadata schema.

Syntax

```
racli remove vpc_user --user_name=myVPCUserName (deprecated)
```


Options

--user_name

The Virtual Private Catalog user for the Recovery Appliance.

Details

(deprecated) This command is deprecated. Please use `racli remove db_user --user_name=myVpcUserName --type=vpc` to remove a vpc user.

- 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
```

racli run

Performs an action for an object to the Recovery Appliance.

Options

Details

Performs an action for an object to the Recovery Appliance.

racli run check

Runs the enabled checks for the Recovery Appliance.

Syntax

```
racli run check { --check_name=CHECK_NAME | --group_name=GROUP_NAME | --all };
```

Options

--all

Indicates that all checks are run.

--check_name

Indicates the name of the check to be run.

--group_name

Indicates the name of the check group to be run.

Details

Runs the enabled checks for the Recovery Appliance.

racli run diagnostics

Runs and collects the required diagnostics on the Recovery Appliance.

Syntax

```
racli run diagnostics { --tag=TAG | --level=LEVEL | --path };
```

Options

--tag

Specifies the name of the tag of diagnostics to be collected, where default tags are:

```

    sar = 'system activity report'
    txt_dump = 'RA dump txt file'
    dp_dump = 'RA dump datapump'
    ord_wait = 'ORDERING_WAIT analysis'
    corruption = 'corruption'
    racli_logs = 'gather racli logs from each node'
    ra_version = 'get RA version information from each node'
    str_cfg = 'get store and config files from each node'
    history = 'get root and oracle history from each node'
    network = 'get network information from each node'
    obtooldiag = 'run obtool gathering useful information'
    osblogs = 'get osb logs from each node'
    osbxcr = 'get /usr/etc/ob/xcr from each node'
    osbstate = 'get backup/admin/state from each node'
    alerts = 'get ASM and RDBMS alert.logs from each node'
    awr = 'generate awr dump for last 7 days'
```

And collection tags:

```

ra_low:  [1] Generates a System Activity Report on the Recovery Appliance
           [2] Gets the version information from each node of the Recovery Appliance
ra_med:  [1] Generates a System Activity Report on the Recovery Appliance
           [2] Performs a text format dump of the Recovery Appliance
           [3] Gets the version information from each node of the Recovery Appliance
ra_high: [1] Generates a System Activity Report on the Recovery Appliance
           [2] Performs a datapump dump of the Recovery Appliance
           [3] Gets the version information from each node of the Recovery Appliance
osb_low: [1] Gets OSB information through obtool on the Recovery Appliance
           [2] Gathers OSB logs from all nodes of the Recovery Appliance
osb_med: [1] Gets OSB information through obtool on the Recovery Appliance
           [2] Gathers OSB logs from all nodes of the Recovery Appliance
           [3] Gathers OSB xcr files from all nodes of the Recovery Appliance
osb_high: [1] Gets OSB information through obtool on the Recovery Appliance
           [2] Gathers OSB logs from all nodes of the Recovery Appliance
           [3] Gathers OSB xcr files from all nodes of the Recovery Appliance
           [4] Gathers OSB state files from the current node of the Recovery Appliance
```

--level

Alternatively, if a tag doesn't exist the level can be specified directly, where:

```

level      1 - system activity report
level      2 - RA dump txt file
level      4 - RA dump datapump
level      8 - ORDERING_WAIT analysis
level     16 - corruption
level     32 - gather racli logs from each node
level     64 - get RA version information from each node
level    128 - get store and config files from each node
level    256 - get root and oracle history from each node
level    512 - get network information from each node
level   1024 - run obtool gathering useful information
level   2048 - get osb logs from each node
level   4096 - get /usr/etc/ob/xcr from each node
level   8192 - get backup/admin/state from each node
level  16384 - get all nodes ASM and ZDLRA alert.logs
level  32768 - generate awr dump for last 7 days

```

The levels are additive, so to gather a system activity report and RA dump txt file you would use `--level=3`

--path

Specifies the location on the Recovery Appliance where the diagnostic bundle will be created, default being `radump`.

Examples

The following example collects racli logs from all nodes and creates a bundle in default location `/radump`

```
racli run diagnostics --tag=racli_logs
```

The following example using the level directly, runs and bundles system activity report and alert.logs and creates a bundle in path `/tmp`.

```
racli run diagnostics --level=16385 --path=/tmp
```

racli run preinstall

Execute preinstallation tasks.

Syntax

```

racli run preinstall [--zdlra_release_file=<FILE>] [--rollback]
[--raadmin_uid=<UID>] [--raadmin_gid=<GID>] [--rasec_uid=<UID>]
[--sec_gid=<GID>] [--rasys_uid=<UID>] [--rasys_gid=<GID>]

```

Options

--rollback

Perform rollback of the preinstall command.

--rollback

Perform rollback of the preinstall command.

--raadmin_uid=<UID>

The new 'raadmin' user id. Value must be `>= 1003` and not equal to 1101 or 1102.

--raadmin_gid=<GID>

The new 'raadmin' group id. Value must be ≥ 1003 and not equal to 1101 or 1102.

--rasec_uid=<UID>

The new 'rasec' user id. Value must be ≥ 1003 and not equal to 1101 or 1102.

--sec_gid=<GID>

The new 'sec' group id. Value must be ≥ 1003 and not equal to 1101 or 1102.

--rasys_uid=<UID>

The new 'rasys' user id. Value must be ≥ 1003 and not equal to 1101 or 1102.

--rasys_gid=<GID>

The new 'rasys' group id. Value must be ≥ 1003 and not equal to 1101 or 1102.

--zdlra_release_file=<FILE>

Specifies ZDLRA release file to use.

racli start

The action to start a specified service (object).

Options

Details

The action to start a specified service (object).

racli start appliance

Starts all services on the Recovery Appliance.

Options

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.

Details

Start NFS services of the Recovery Appliance.

racli status

Check status of objects associated with the Recovery Appliance.

Options

Details

Check status of objects associated with the Recovery Appliance.

racli status aide

Returns the status of the Advanced Intrusion Detection Environment (AIDE).

Syntax

```
racli status aide
```

Options

racli status appliance

Returns a summary of the status of the appliance.

Syntax

```
racli status appliance
```

Options

Details

Returns a summary of the status of the appliance.

racli status archived_log_grouping

Returns if archived log grouping is enabled or disabled, if enabled it returns the group log count.

Syntax

```
racli status archived_log_grouping
```

Options

racli status cloud_location

Lists the status of a cloud storage object for operations performed by the Recovery Appliance.

Syntax

```
racli status cloud_location [ --location_name=CLOUD_LOCATION_NAME ] [--long]
```

Options

--location_name

Specifies the name of the cloud location for status.

--long

Lists extended details about cloud storage known to the Recovery Appliance.

Details

Lists the status of a cloud storage object for operations performed by the Recovery Appliance. This action defaults to **all** for cloud_location unless specified.

racli status compliance

Lists the status (either PASS or FAIL) of all compliance-ready checks.

Syntax

```
racli status compliance [--long]
```

Options

--long

Lists extended details of each compliance-ready check: for the Recovery Appliance.

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

Verifies that the rebalance operation completed after storage cells or storage disks are added to the ASM disk group.

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.

racli status keystore

Returns the status of the keystore in use by the Recovery Appliance.

Syntax

```
racli status keystore
```

Options

Details

Returns the status of the keystore in use by the Recovery Appliance.

racli status listener

Returns the status of an individual listener or all listeners on the Recovery Appliance.

Syntax

```
racli status listener [--listener_name | --all]
```

Options

Details

--all

Specifies that the status of all listeners should be returns. This is the default operation if no arguments are provided.

--listener_name

Specifies the name of the listener to get status information about.

racli status nfs

Returns the status of the Network File Service (NFS).

Syntax

```
racli status nfs --name=myMountName
```

Options

Details

Returns the status of the Network File Service (NFS).

racli status okv_endpoint

Returns the status of the OKV_ENDPOINTs associated with the Recovery Appliance.

Syntax

```
racli status okv_endpoint
```

Options

Details

Returns the status of the OKV_ENDPOINTs associated with the Recovery Appliance.

racli status request

Shows status of all scan on the Recovery Appliance.

Syntax

```
racli status request [--request_id=VALUE]
```

Options

--request_id

Specifies the request identifier for the status report.

racli status scan

Shows status of all scan on the Recovery Appliance.

Syntax

```
racli status scan
```

Options

racli status scan_listener

Shows status of all scan_listeners on the Recovery Appliance.

Syntax

```
racli status scan_listener
```


Options

racli stop

The action to stop a specified service (object).

Options**Details**

The action to stop a specified service (object).

racli stop appliance

Stops all services on the Recovery Appliance.

Options

racli stop nfs

Stop NFS services of the Recovery Appliance.

Syntax

```
racli stop nfs --name=myMountName
```

Options**--name**

Name of the ACFS mount to start.

Details

Stop NFS services of the Recovery Appliance.

racli uninstall

Backs out a portion of the installation.

Options**Details**

Backs out a portion of the installation.

racli uninstall appliance

Backout a step during install.

Syntax

```
racli uninstall appliance --step={ STEP_NUMBER | ALL }
```

Options

--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.

Details

Backout a step during install.

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
```

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.

Details

Uninstalls the specified OKV end point.

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
```

racli update

Update a component of the Recovery Appliance.

Options

Details

Update a component of the Recovery Appliance.

racli update aide

Updates the Advanced Intrusion Detection Environment (AIDE).

Syntax

```
racli update aide
```

Options

Details

AIDE is an Exadata utility that creates a database of files on the system and tracks changes. AIDE ensures file integrity and detects system intrusions.

racli update appliance

Updates the Recover Appliance.

Syntax

```
racli update appliance [--ignore_incidents]
```

Options

--ignore_incidents

Ignore active incidents (not recommended).

Details

This command replaces "racli patch appliance" and "racli upgrade appliance" and determines whether the update should be a patch or an upgrade. It performs all the prechecks. If the prechecks pass, it proceeds with the patch or upgrade steps. If the update fails in the middle and after the identified error is corrected, this command can be rerun and it will continue with the step that failed.

Do not run this command from `:/raacfs/raadmin/tmp/`. because patching will fail when the CRS is stopped.

1. As an admin user or root user, stage the software release zip file.

```
$ scp <ZIP_FILE_NAME> <USER>@firstnode:/raacfs/raadmin/tmp/
```

2. As an admin user or root user, run the ra_preinstall.

```
<USER>@firstnode> # racli run preinstall  
--zdlra_release_file=/raacfs/raadmin/tmp/<ZIP_FILE_NAME>
```

Sta

racli update parameter

Update init parameters in the Recovery Appliance.

Syntax

```
racli update parameter --type=db [--restart_db]
```

Options

--type

The type of init parameter to be updated. Presently only "db" is supported.

Details

Update init parameters in the Recovery Appliance.

racli update ssh_keys

.

Options

racli update ssh_known_hosts

Maintains a system-wide configuration file `/etc/ssh/ssh_known_hosts` with host keys from all of the nodes in the Recovery Appliance cluster.

Syntax

```
racli update ssh_known_hosts
```

Details

This command can help resolve installation issues that arise from SSH KEY inconsistencies.

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

Details

Update RA pool data file with the most current number of storage cells.

racli update support_tools

Updates the AHF and Exachk versions using the staged binaries zip files.

Syntax

```
racli update support_tools {--ahf_file=<AHF_FILE_NAME>} {--  
exachk_file=<EXACHK_FILE_NAME>}
```

Options

--ahf_file

Updates AHF version the the provided AHF version.

--exachk_file

Updates Exachk version the the provided Exachk version.

Details

Updates the TFA and Exachk versions using the staged binaries zip files.

racli version

Provides version information for software running on the Recovery Appliance.

Syntax

```
racli version [--history] [--cache] [--all_nodes] [--json] [--force]
```

Options

--all_nodes

Displays the version from all nodes in cluster.

--cache

Retrieve version values from existing cache file.

--force

Retrieve real-time data of versions and updates `raa_versions` table if database is online.

--history

Displays all active and previous versions. A blank line is expected when there is no history data to report.

--json

Display data in JavaScript Object Notation (JSON).

Details

Provides version information for software running on the Recovery Appliance.

The version information provided includes:

exadata image

The exadata image currently installed on the Recovery Appliance. A blank line is expected when there is no history data to report.

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.

Part V

Cabling Information for Oracle Zero Data Loss Recovery Appliance

Oracle Zero Data Loss Recovery Appliance has different cabling tables depending on the hardware model.

16

In-Rack Cabling Tables for Recovery Appliance RA23

This appendix identifies the location of the various components in an Oracle Zero Data Loss Recovery Appliance RA23 full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

Cable Color Coding


The cables used in Recovery Appliance are color coded as follows:

- **Black:** RDMA Network Fabric 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

Recovery Appliance RA23 Rack Layout

The [Figure 16-1](#) figure shows the rack configuration and the sequence for populating the rack from minimum to the maximum (full) rack configuration. The "RA23 HC (*number*)" is the insertion sequence when expanding storage from the minimum RA23 HC (3) to full RA23 HC (17).

Figure 16-1 RA23 Rack Population



Oracle Integrated Lights Out Manager Cabling Connections RA23

Table 16-1 shows the cable connections from the servers to the Oracle ILOM switches in an Oracle Zero Data Loss Recovery Appliance RA21 and RA23 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 16-1 ILOM Cabling for the Recovery Appliance RA23 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
U16	Compute server	30
U14	Compute 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.

Administrative Ethernet Port Cabling Tables RA23

Figure 17-2 shows the cable connections from the servers to the gigabit Ethernet switch in an Oracle Zero Data Loss Recovery Appliance RA23 full rack. The port numbers are the same in racks with fewer storage servers.

The Management Network Switch is located in rack unit 21 on Recovery Appliance RA23. Ethernet cables are blue.

In the following tables, *Un* denotes the unit height in the rack, where *n* is the number, such as U20. The cables used for the Management Network Switch are blue in color. Cabling connections marked with * are required for all rack configurations, and represent the minimum rack configuration.

Table 16-2 Ethernet Cabling for the Recovery Appliance RA23 Racks

From Rack Unit ¹	Network Management	To U21 Ethernet Port	Component
U39	NET 0	1	Storage server

Table 16-2 (Cont.) Ethernet Cabling for the Recovery Appliance RA23 Racks

From Rack Unit ¹	Network Management	To U21 Ethernet Port	Component
U39	NET MGT	2	Storage server
U37	NET 0	3	Storage server
U37	NET MGT	4	Storage server
U35	NET 0	5	Storage server
U35	NET MGT	6	Storage server
U33	NET 0	7	Storage server
U33	NET MGT	8	Storage server
U31	NET 0	9	Storage server
U31	NET MGT	10	Storage server
U29	NET 0	11	Storage server
U29	NET MGT	12	Storage server
U27	NET 0	13	Storage server
U27	NET MGT	14	Storage server
U41	NET 0	15	Storage server
U41	NET MGT	16	Storage server
U25	NET 0	17	Storage server
U25	NET MGT	18	Storage server
PDU A*	NET MGT	19	
PDU B*	NET MGT	20	
U23	NET 0	21	Storage server
U23	NET MGT	22	Storage server
U18	NET 0	25	Storage server
U18	NET MGT	26	Storage server
U16*	NET 0	29	Compute server
U16*	NET MGT	30	Compute server
U14*	NET 0	31	Compute server
U14*	NET MGT	32	Compute server
U12	NET 0	33	Storage server
U12	NET MGT	34	Storage server
U10	NET 0	35	Storage server
U10	NET MGT	36	Storage server
U08	NET 0	37	Storage server
U08	NET MGT	38	Storage server
U06*	NET 0	39	Storage server
U06*	NET MGT	40	Storage server
U04*	NET 0	41	Storage server
U04*	NET MGT	42	Storage server

Table 16-2 (Cont.) Ethernet Cabling for the Recovery Appliance RA23 Racks

From Rack Unit ¹	Network Management	To U21 Ethernet Port	Component
U02*	NET 0	43	Storage server
U02*	NET MGT	44	Storage server
U22*	NET MGT 0	45	RDMA Network Fabric switch
U20*	NET MGT 0	46	RDMA Network Fabric switch
U01*	NET MGT 0	47	
	External	48	

¹ Un is the unit location in the rack, where *n* is the number.

* These marked components are required for all rack configurations, and together represent the minimum rack configuration.

Single-Phase Power Distribution Unit Cabling Connections RA23

Table 17-3 shows the cable connections from the power distribution units (PDUs) to the rack units in an Oracle Zero Data Loss Recovery Appliance RA23 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 16-3 Single-Phase PDU Cabling for the Recovery Appliance RA23 Full Racks

Rack Unit ¹	PDU-A & PSU0	PDU-B & PSU1	Cable Length
U42	G4-3	G1-3	3 meters
U41	G4-6	G1-0	3 meters
U40	G5-5	G0-1	3 meters
U39	G1-3	G4-3	3 meters
U37	G1-2	G4-4	3 meters
U35	G1-5	G4-1	3 meters
U33	G2-6	G3-0	2 meters
U31	G3-3	G2-0	3 meters
U29	G3-3	G2-3	3 meters
U27	G3-2	G2-4	3 meters
U26	G5-6	G0-0	3 meters
U25	G4-1	G1-5	3 meters
U24	G4-0	G1-6	3meters
U23	G5-1	G0-5	3meters
U22*	G1-6	G4-0	2 meters
U21*	G0-4	G5-2	2 meters

Table 16-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance RA23 Full Racks

Rack Unit ¹	PDU-A & PSU0	PDU-B & PSU1	Cable Length
U20*	G0-3	G5-3	2 meters
U19	G5-3	G0-3	3 meters
U18	G4-2	G1-4	3 meters
U16*	G1-4	G4-2	2 meters
U14*	G2-1	G3-5	3 meters
U12	G2-0	G3-6	3 meters
U10	G2-3	G3-3	3 meters
U8	G2-4	G3-2	3 meters
U6*	G0-6	G5-0	3 meters
U4*	G0-2	G5-4	3 meters
U2*	G0-5	G5-1	3 meters
U1*	G2-5	G3-1	3 meters

¹ Un is the unit location in the rack, where *n* is the number.

* These items are required for all Recovery Appliances.

Three-Phase Power Distribution Unit Cabling Connections RA23

Table 17-4 shows the cable connections from the power distribution units (PDUs) to the rack units in an Oracle Zero Data Loss Recovery Appliance RA23 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 16-4 Three-Phase PDU Cabling for the Recovery Appliance RA23 Full Racks

Rack Unit ¹	PDU-A & PSU0	PDU-B & PSU1	Cable Length
U42	G4-1	G1-5	3 meters
U41	G3-1	G0-5	3 meters
U40	G5-4	G2-2	3 meters
U39	G1-4	G4-2	3 meters
U37	G2-1	G5-5	3 meters
U35	G0-4	G3-2	3 meters
U33	G2-2	G5-4	3 meters
U31	G1-2	G4-4	3 meters
U29	G0-1	G3-5	3 meters
U27	G2-0	G5-6	3 meters
U26	G5-1	G2-5	3 meters
U25	G1-1	G4-5	3 meters

Table 16-4 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance RA23 Full Racks

Rack Unit ¹	PDU-A & PSU0	PDU-B & PSU1	Cable Length
U23	G0-2	G3-4	3 meters
U22*	G1-6	G4-0	2 meters
U21*	G2-6	G5-0	2 meters
U20*	G1-5	G4-1	2 meters
U19	G3-6	G0-0	3 meters
U18	G4-2	G1-4	3 meters
U16*	G1-3	G4-3	3 meters
U14*	G0-5	G3-1	3 meters
U12	G2-5	G5-1	3 meters
U10	G0-3	G3-3	3 meters
U8	G2-3	G5-3	3 meters
U6*	G2-4	G5-2	3 meters
U4*	G0-6	G3-0	3 meters
U2*	G0-0	G3-6	3 meters
U1*	G1-0	G4-6	3 meters

¹ Un is the unit location in the rack, where *n* is the number.

* These items are required for all Recovery Appliances.

RDMA Network Fabric Switch-to-Switch Cabling Connections RA23

For RoCE Network Fabric, [Table 16-5](#) lists the ports and cable lengths for the connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) RDMA Network Fabric cables.

Table 16-5 RDMA Network Fabric Switch-to-Switch Cabling RA23

From RDMA Network Fabric Switch U20 Port ¹	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
04	04	1 meter
05	05	1 meter
06	06	1 meter
07	07	1 meter
30	30	1 meter
31	31	1 meter
32	32	1 meter

Table 16-5 (Cont.) RDMA Network Fabric Switch-to-Switch Cabling RA23

From RDMA Network Fabric Switch U20 Port ¹	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
33	33	1 meter

¹ Un is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Server Cabling Connections RA23

The following tables list the location, ports, and cables for the connections from the RDMA Network Fabric switches to the computer and HC storage servers in a Oracle Zero Data Loss Recovery Appliance RA23 full rack. The port numbers and cable lengths are the same in racks with fewer HC Storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) cables.

Table 16-6 RDMA Network Fabric Switch to Compute Server Cabling for the Recovery Appliance RA23 Full Racks

From RDMA Network Fabric Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U20	22*	U16	2U compute server	PCIe 5, P1	3 meters
U20	23*	U14	2U compute server	PCIe 5, P1	3 meters
U22	22*	U16	2U compute server	PCIe 5, P2	3 meters
U22	23*	U14	2U compute server	PCIe 5, P2	3 meters

¹ Un is the unit location in the rack, where *n* is the number.

* All Recovery Appliances have these entries.

Table 16-7 RDMA Network Fabric Switch to HC Server Cabling for the Recovery Appliance RA23 Full Racks

From RDMA Network Fabric Switch Rack Unit ¹	RoCE Port	To Rack Unit	Device	Port	Cable Length
U20	15	U41	HC Storage server	PCIe 10, P1	3 meters
U20	8	U39	HC Storage server	PCIe 10, P1	3 meters
U20	9	U37	HC Storage server	PCIe 10, P1	3 meters
U20	10	U35	HC Storage server	PCIe 10, P1	3 meters
U20	11	U33	HC Storage server	PCIe 10, P1	3 meters
U20	12	U31	HC Storage server	PCIe 10, P1	2 meters
U20	13	U29	HC Storage server	PCIe 10, P1	2 meters

Table 16-7 (Cont.) RDMA Network Fabric Switch to HC Server Cabling for the Recovery Appliance RA23 Full Racks

From RDMA Network Fabric Switch Rack Unit ¹	RoCE Port	To Rack Unit	Device	Port	Cable Length
U20	14	U27	HC Storage server	PCIe 10, P1	2 meters
U20	15	U26	HC Storage server	PCIe 10, P1	2 meters
U20	16	U25	HC Storage server	PCIe 10, P1	2 meters
U20	17	U24	HC Storage server	PCIe 10, P1	2 meters
U20	18	U23	HC Storage server	PCIe 10, P1	2 meters
U20	19	U19	HC Storage server	PCIe 10, P1	2 meters
U20	20	U18	HC Storage server	PCIe 10, P1	2 meters
U20	24	U12	HC Storage server	PCIe 10, P1	2 meters
U20	25	U10	HC Storage server	PCIe 10, P1	2 meters
U20	26	U8	HC Storage server	PCIe 10, P1	2 meters
U20	27	U6	HC Storage server	PCIe 10, P1	2 meters
U20	28	U4	HC Storage server	PCIe 10, P1	2 meters
U20	29	U2	HC Storage server	PCIe 10, P1	2 meters
U20 and U22 should mirror one another					
U22	15	U41	HC Storage server	PCIe 10, P2	3 meters
U22	8	U39	HC Storage server	PCIe 10, P2	3 meters
U22	9	U37	HC Storage server	PCIe 10, P2	3 meters
U22	10	U35	HC Storage server	PCIe 10, P2	3 meters
U22	11	U33	HC Storage server	PCIe 10, P2	3 meters
U22	12	U31	HC Storage server	PCIe 10, P2	2 meters
U22	13	U29	HC Storage server	PCIe 10, P2	2 meters
U22	14	U27	HC Storage server	PCIe 10, P2	2 meters
U22	15	U26	HC Storage server	PCIe 10, P2	2 meters
U22	16	U25	HC Storage server	PCIe 10, P2	2 meters
U22	17	U24	HC Storage server	PCIe 10, P2	2 meters
U22	18	U23	HC Storage server	PCIe 10, P2	2 meters
U22	19	U19	HC Storage server	PCIe 10, P2	2 meters
U22	20	U18	HC Storage server	PCIe 10, P2	2 meters
U22	24	U12	HC Storage server	PCIe 10, P2	2 meters
U22	25	U10	HC Storage server	PCIe 10, P2	2 meters
U22	25	U10	HC Storage server	PCIe 10, P2	2 meters
U22	26	U8	HC Storage server	PCIe 10, P2	2 meters
U22	27	U6	HC Storage server	PCIe 10, P2	2 meters
U22	28	U4	HC Storage server	PCIe 10, P2	2 meters
U22	29	U2	HC Storage server	PCIe 10, P2	2 meters

¹ Un is the unit location in the rack, where n is the number.

* All Recovery Appliances have these entries.

In-Rack Cabling Tables for Recovery Appliance RA21

This appendix identifies the location of the various components in an Oracle Zero Data Loss Recovery Appliance RA21 full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

Cable Color Coding

The cables used in Recovery Appliance are color coded as follows:

- **Black:** RDMA Network Fabric 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

Recovery Appliance RA21 Rack Layout

The [Figure 17-1](#) figure shows in the minimum rack configuration in the middle, the maximum rack configuration on the right, and on the left the sequence "RA21 HC (*number*)" that the racks are inserted when expanding storage from the minimum RA21 HC (3) to full RA21 HC (18).

Figure 17-1 RA21 Rack Population





Figure 17-2 shows the front and rear views of an Oracle Zero Data Loss Recovery Appliance RA21 full rack. The front is shown on the left, and the rear is shown on the right. The red bullet next to a component means that it is required for all Recovery Appliances.

Figure 17-2 Rack Layout of the Recovery Appliance RA21 Full Rack





Note:

The 1U empty space at the bottom of the rack is reserved for an optional spine switch, which is used in multi-rack installations. Do not install any other unit in this space.

Figure 18-2 shows the front and rear views of a an Oracle Zero Data Loss Recovery Appliance RA21 minimum configuration.

Figure 17-3 Rack Layout of the Recovery Appliance RA21 Minimum Configuration

Rear View			Front View		
U42	1U Solid Filler	U42	U42	1U Solid Filler	U42
U41	1U Solid Filler	U41	U41	1U Solid Filler	U41
U40	1U Solid Filler	U40	U40	1U Solid Filler	U40
U39	1U Solid Filler	U39	U39	1U Solid Filler	U39
U38	1U Solid Filler	U38	U38	1U Solid Filler	U38
U37	1U Solid Filler	U37	U37	1U Solid Filler	U37
U36	1U Solid Filler	U36	U36	1U Solid Filler	U36
U35	1U Solid Filler	U35	U35	1U Solid Filler	U35
U34	1U Solid Filler	U34	U34	1U Solid Filler	U34
U33	1U Solid Filler	U33	U33	1U Solid Filler	U33
U32	1U Solid Filler	U32	U32	1U Solid Filler	U32
U31	1U Solid Filler	U31	U31	1U Solid Filler	U31
U30	1U Solid Filler	U30	U30	1U Solid Filler	U30
U29	1U Solid Filler	U29	U29	1U Solid Filler	U29
U28	1U Solid Filler	U28	U28	1U Solid Filler	U28
U27	1U Solid Filler	U27	U27	1U Solid Filler	U27
U26	1U Solid Filler	U26	U26	1U Solid Filler	U26
U25	1U Solid Filler	U25	U25	1U Solid Filler	U25
U24	1U Solid Filler	U24	U24	1U Solid Filler	U24
U23	1U Solid Filler	U23	U23	1U Solid Filler	U23
U22	* RDMA Network Fabric	U22	U22	* 1U Vented Filler	U22
U21	* Ethernet Switch	U21	U21	* 1U Vented Filler	U21
U20	* RDMA Network Fabric	U20	U20	* 1U Vented Filler	U20
U19	1U Solid Filler	U19	U19	1U Solid Filler	U19
U18	1U Solid Filler	U18	U18	1U Solid Filler	U18
U17	* Compute Server	U17	U17	* Compute Server	U17
U16	* Compute Server	U16	U16	* Compute Server	U16
U15	1U Solid Filler	U15	U15	1U Solid Filler	U15
U14	1U Solid Filler	U14	U14	1U Solid Filler	U14
U13	1U Solid Filler	U13	U13	1U Solid Filler	U13
U12	1U Solid Filler	U12	U12	1U Solid Filler	U12
U11	1U Solid Filler	U11	U11	1U Solid Filler	U11
U10	1U Solid Filler	U10	U10	1U Solid Filler	U10
U9	1U Solid Filler	U9	U9	1U Solid Filler	U9
U8	1U Solid Filler	U8	U8	1U Solid Filler	U8
U7	Storage Server	U7	U7	Storage Server	U7
U6		U6	U6		U6

Oracle Integrated Lights Out Manager Cabling Connections RA21

Table 17-1 shows the cable connections from the servers to the Oracle Integrated Lights Out Management switches in an Oracle Zero Data Loss Recovery Appliance RA21 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 17-1 ILOM Cabling for the Recovery Appliance RA21 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.

Administrative Ethernet Port Cabling Tables RA21

Table 17-2 shows the cable connections from the servers to the gigabit Ethernet switch in an Oracle Zero Data Loss Recovery Appliance RA21 full rack. The port numbers are the same in racks with fewer storage servers.

The Management Network Switch is located in rack unit 21 on Recovery Appliance RA21. Ethernet cables are blue.

In the following tables, *Un* denotes the unit height in the rack, where *n* is the number, such as U20. The cables used for the Management Network Switch are blue in color. Cabling connections marked with * are required for all rack configurations, and represent the minimum rack configuration.

Table 17-2 Ethernet Cabling for the Recovery Appliance RA21 Racks

From Rack Unit ¹	Network Management	To U21 Ethernet Port	Component
U39	NET 0	1	Storage server
U39	NET MGT	2	Storage server
U37	NET 0	3	Storage server
U37	NET MGT	4	Storage server
U35	NET 0	5	Storage server
U35	NET MGT	6	Storage server
U33	NET 0	7	Storage server
U33	NET MGT	8	Storage server
U31	NET 0	9	Storage server
U31	NET MGT	10	Storage server
U29	NET 0	11	Storage server
U29	NET MGT	12	Storage server
U27	NET 0	13	Storage server
U27	NET MGT	14	Storage server
U41	NET 0	15	Storage server
U41	NET MGT	16	Storage server
U25	NET 0	17	Storage server
U25	NET MGT	18	Storage server
PDU A*	NET MGT	19	
PDU B*	NET MGT	20	
U23	NET 0	21	Storage server
U23	NET MGT	22	Storage server
U18	NET 0	25	Storage server
U18	NET MGT	26	Storage server
U17*	NET 0	27	Compute server
U17*	NET MGT	28	Compute server
U16*	NET 0	29	Compute server
U16*	NET MGT	30	Compute server
U14	NET 0	31	Storage server
U14	NET MGT	32	Storage server
U12	NET 0	33	Storage server
U12	NET MGT	34	Storage server
U10	NET 0	35	Storage server
U10	NET MGT	36	Storage server

Table 17-2 (Cont.) Ethernet Cabling for the Recovery Appliance RA21 Racks

From Rack Unit ¹	Network Management	To U21 Ethernet Port	Component
U08	NET 0	37	Storage server
U08	NET MGT	38	Storage server
U06*	NET 0	39	Storage server
U06*	NET MGT	40	Storage server
U04*	NET 0	41	Storage server
U04*	NET MGT	42	Storage server
U02*	NET 0	43	Storage server
U02*	NET MGT	44	Storage server
U22*	NET MGT 0	45	RDMA Network Fabric switch
U20*	NET MGT 0	46	RDMA Network Fabric switch
U01*	NET MGT 0	47	
	External	48	

¹ Un is the unit location in the rack, where *n* is the number.

* These marked components are required for all rack configurations, and together represent the minimum rack configuration.

Single-Phase Power Distribution Unit Cabling Connections RA21

Table 17-3 shows the cable connections from the power distribution units (PDUs) to the rack units in an Oracle Zero Data Loss Recovery Appliance RA21 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 17-3 Single-Phase PDU Cabling for the Recovery Appliance RA21 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U41	G4-6	G1-0	3 meters
U39	G4-4	G1-2	3 meters
U37	G5-1	G0-5	3 meters
U35	G1-5	G4-1	3 meters
U33	G2-6	G3-0	2 meters
U31	G3-3	G2-3	3 meters
U29	G3-6	G2-0	3 meters
U27	G3-2	G2-4	3 meters
U25	G4-1	G1-5	3 meters
U23	G5-0	G0-6	3meters

Table 17-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance RA21 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U22*	G1-6	G4-0	2 meters
U21*	G0-4	G5-2	2 meters
U20*	G0-3	G5-3	2 meters
U18	G4-2	G1-4	3 meters
U17*	G1-2	G4-4	3 meters
U16*	G1-4	G4-2	2 meters
U14	G2-1	G3-5	3 meters
U12	G2-0	G3-6	3 meters
U10	G2-3	G3-3	3 meters
U8	G2-4	G3-2	3 meters
U6*	G1-3	G4-3	3 meters
U4*	G0-2	G5-4	3 meters
U2*	G0-5	G5-1	3 meters
U1*	G2-5	G3-1	3 meters

¹ Un is the unit location in the rack, where *n* is the number.

* These items are required for all Recovery Appliances.

Three-Phase Power Distribution Unit Cabling Connections RA21

Table 17-4 shows the cable connections from the power distribution units (PDUs) to the rack units in an Oracle Zero Data Loss Recovery Appliance RA21 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 17-4 Three-Phase PDU Cabling for the Recovery Appliance RA21 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U41	G5-2	G2-4	3 meters
U39	G4-6	G1-0	3 meters
U37	G4-1	G1-5	3 meters
U35	G0-4	G3-2	3 meters
U33	G2-2	G5-4	3 meters
U31	G1-2	G4-4	3 meters
U29	G0-1	G3-5	3 meters
U27	G2-0	G5-6	3 meters
U25	G3-0	G0-6	3 meters
U23	G3-5	G0-1	3 meters

Table 17-4 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance RA21 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U22*	G1-6	G4-0	2 meters
U21*	G2-6	G5-0	2 meters
U20*	G1-5	G4-1	2 meters
U18	G1-4	G4-2	3 meters
U17*	G2-1	G5-5	3 meters
U16*	G1-3	G4-3	3 meters
U14	G0-5	G3-1	3 meters
U12	G2-5	G5-1	3 meters
U10	G0-3	G3-3	3 meters
U8	G2-3	G5-3	3 meters
U6*	G2-4	G5-2	3 meters
U4*	G0-6	G3-0	3 meters
U2*	G0-0	G3-6	3 meters
U1*	G1-0	G4-6	3 meters

¹ Un is the unit location in the rack, where *n* is the number.

* These items are required for all Recovery Appliances.

RDMA Network Fabric Switch-to-Switch Cabling Connections RA21

For RoCE Network Fabric, [Table 17-5](#) lists the ports and cable lengths for the connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) RDMA Network Fabric cables.

Table 17-5 RDMA Network Fabric Switch-to-Switch Cabling RA21

From RDMA Network Fabric Switch U20 Port ¹	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
04	04	1 meter
05	05	1 meter
06	06	1 meter
07	07	1 meter
30	30	1 meter
31	31	1 meter
32	32	1 meter
33	33	1 meter

¹ Un is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Server Cabling Connections RA21

The following tables list the location, ports, and cables for the connections from the RDMA Network Fabric switches to the computer and HC storage servers in a Oracle Zero Data Loss Recovery Appliance RA21 full rack. The port numbers and cable lengths are the same in racks with fewer HC Storage servers. The switches are located in rack units 20 and 22.

All cables are black quad data rate (QDR) cables.

Table 17-6 RDMA Network Fabric Switch to Compute Server Cabling for the Recovery Appliance RA21 Full Racks

From RDMA Network Fabric Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U20	21*	U17	1U compute server	PCIe 5, P1	3 meters
U20	22*	U16	1U compute server	PCIe 5, P1	3 meters
U22	21*	U17	1U compute server	PCIe 5, P2	3 meters
U22	22*	U16	1U compute server	PCIe 5, P2	3 meters

¹ Un is the unit location in the rack, where *n* is the number.

* All Recovery Appliances have these entries.

Table 17-7 RDMA Network Fabric Switch to HC Server Cabling for the Recovery Appliance RA21 Full Racks

From RDMA Network Fabric Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U20	8	U39	HC Storage server	PCIe 10, P1	3 meters
U20	9	U37	HC Storage server	PCIe 10, P1	3 meters
U20	10	U35	HC Storage server	PCIe 10, P1	3 meters
U20	11	U33	HC Storage server	PCIe 10, P1	3 meters
U20	12	U31	HC Storage server	PCIe 10, P1	2 meters
U20	13	U29	HC Storage server	PCIe 10, P1	2 meters
U20	14	U27	HC Storage server	PCIe 10, P1	2 meters
U20	15	U26	HC Storage server	PCIe 10, P1	2 meters
U20	16	U25	HC Storage server	PCIe 10, P1	2 meters
U20	17	U24	HC Storage server	PCIe 10, P1	2 meters
U20	18	U23	HC Storage server	PCIe 10, P1	2 meters
U20	19	U19	HC Storage server	PCIe 10, P1	2 meters
U20	20	U18	HC Storage server	PCIe 10, P1	2 meters
U20	23	U14	HC Storage server	PCIe 10, P1	2 meters

Table 17-7 (Cont.) RDMA Network Fabric Switch to HC Server Cabling for the Recovery Appliance RA21 Full Racks

From RDMA Network Fabric Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U20	24	U12	HC Storage server	PCIe 10, P1	2 meters
U20	25	U10	HC Storage server	PCIe 10, P1	2 meters
U20 and U22 should mirror one another					
U22	8	U39	HC Storage server	PCIe 10, P2	3 meters
U22	9	U37	HC Storage server	PCIe 10, P2	3 meters
U22	10	U35	HC Storage server	PCIe 10, P2	3 meters
U22	11	U33	HC Storage server	PCIe 10, P2	3 meters
U22	12	U31	HC Storage server	PCIe 10, P2	2 meters
U22	13	U29	HC Storage server	PCIe 10, P2	2 meters
U22	14	U27	HC Storage server	PCIe 10, P2	2 meters
U22	15	U26	HC Storage server	PCIe 10, P2	2 meters
U22	16	U25	HC Storage server	PCIe 10, P2	2 meters
U22	17	U24	HC Storage server	PCIe 10, P2	2 meters
U22	18	U23	HC Storage server	PCIe 10, P2	2 meters
U22	19	U19	HC Storage server	PCIe 10, P2	2 meters
U22	20	U18	HC Storage server	PCIe 10, P2	2 meters
U22	23	U14	HC Storage server	PCIe 10, P2	2 meters
U22	24	U12	HC Storage server	PCIe 10, P2	2 meters
U22	25	U10	HC Storage server	PCIe 10, P2	2 meters

¹ Un is the unit location in the rack, where *n* is the number.

* All Recovery Appliances have these entries.

In-Rack Cabling Tables for Recovery Appliance X8M

This appendix identifies the location of the various components in a Recovery Appliance X8M full rack and a minimum configuration. It also lists the cabling connections among components within the rack.

Cable Color Coding

The cables used in Recovery Appliance are color coded as follows:

- **Black:** RDMA Network Fabric 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

Recovery Appliance RA21 and X8M Rack Layout

Figure 18-1 shows the front and rear views of a Recovery Appliance RA21 and X8M 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 spine switch, which is used in multirack installations. Do not install any other unit in this space.

Figure 18-1 Rack Layout of the Recovery Appliance RA21 and X8M Full Rack




Figure 18-2 shows the front and rear views of a Recovery Appliance RA21 and X8M minimum configuration.

Figure 18-2 Rack Layout of the Recovery Appliance RA21 and X8M Minimum Configuration

Rear View			Front View		
U42	1U Solid Filler	U42	U42	1U Solid Filler	U42
U41	1U Solid Filler	U41	U41	1U Solid Filler	U41
U40	1U Solid Filler	U40	U40	1U Solid Filler	U40
U39	1U Solid Filler	U39	U39	1U Solid Filler	U39
U38	1U Solid Filler	U38	U38	1U Solid Filler	U38
U37	1U Solid Filler	U37	U37	1U Solid Filler	U37
U36	1U Solid Filler	U36	U36	1U Solid Filler	U36
U35	1U Solid Filler	U35	U35	1U Solid Filler	U35
U34	1U Solid Filler	U34	U34	1U Solid Filler	U34
U33	1U Solid Filler	U33	U33	1U Solid Filler	U33
U32	1U Solid Filler	U32	U32	1U Solid Filler	U32
U31	1U Solid Filler	U31	U31	1U Solid Filler	U31
U30	1U Solid Filler	U30	U30	1U Solid Filler	U30
U29	1U Solid Filler	U29	U29	1U Solid Filler	U29
U28	1U Solid Filler	U28	U28	1U Solid Filler	U28
U27	1U Solid Filler	U27	U27	1U Solid Filler	U27
U26	1U Solid Filler	U26	U26	1U Solid Filler	U26
U25	1U Solid Filler	U25	U25	1U Solid Filler	U25
U24	1U Solid Filler	U24	U24	1U Solid Filler	U24
U23	1U Solid Filler	U23	U23	1U Solid Filler	U23
U22	RDMA Network Fabric	U22	U22	1U Vented Filler	U22
U21	Ethernet Switch	U21	U21	1U Vented Filler	U21
U20	RDMA Network Fabric	U20	U20	1U Vented Filler	U20
U19	1U Solid Filler	U19	U19	1U Solid Filler	U19
U18	1U Solid Filler	U18	U18	1U Solid Filler	U18
U17	Compute Server	U17	U17	Compute Server	U17
U16	Compute Server	U16	U16	Compute Server	U16
U15	1U Solid Filler	U15	U15	1U Solid Filler	U15
U14	1U Solid Filler	U14	U14	1U Solid Filler	U14
U13	1U Solid Filler	U13	U13	1U Solid Filler	U13
U12	1U Solid Filler	U12	U12	1U Solid Filler	U12
U11	1U Solid Filler	U11	U11	1U Solid Filler	U11
U10	1U Solid Filler	U10	U10	1U Solid Filler	U10
U9	1U Solid Filler	U9	U9	1U Solid Filler	U9
U8	1U Solid Filler	U8	U8	1U Solid Filler	U8
U7	Storage Server	U7	U7	Storage Server	U7
U6	Storage Server	U6	U6	Storage Server	U6
U5	Storage Server	U5	U5	Storage Server	U5
U4	Storage Server	U4	U4	Storage Server	U4
U3	Storage Server	U3	U3	Storage Server	U3
U2	Storage Server	U2	U2	Storage Server	U2
U1	1U Solid Filler	U1	U1	1U Solid Filler	U1

Oracle Integrated Lights Out Manager Cabling Connections X8M

Table 19-1 shows the cable connections from the servers to the Oracle ILOM switches in a Recovery Appliance X8M 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 18-1 ILOM Cabling for the Recovery Appliance X8M 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 U_n is the unit location in the rack, where n is the number.

Administrative Ethernet Port Cabling Tables X8M

Table 18-2 shows the cable connections from the servers to the gigabit Ethernet switch in a Recovery Appliance X8M full rack. The port numbers are the same in racks with fewer storage servers.

The Management Network Switch is located in rack unit 21 on Recovery Appliance X8M. Ethernet cables are blue.

In the following tables, U_n denotes the unit height in the rack, where n is the number, such as U20. The cables used for the Management Network Switch are blue in color. The following tables are in this section:

Table 18-2 Ethernet Cabling for the Recovery Appliance X8M 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
U31	Storage server	9
U29	Storage server	11
U27	Storage server	13
U25	Storage server	17
U23	Storage server	21
U22	RDMA Network Fabric switch	45
U20	RDMA Network Fabric 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.

Single-Phase Power Distribution Unit Cabling Connections X8M

Table 19-3 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8M 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 18-3 Single-Phase PDU Cabling for the Recovery Appliance X8M Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U41	G4-2	G1-4	2 meters
U39	G4-4	G1-2	2 meters
U37	G5-3	G2-3	2 meters
U35	G3-2	G0-4	2 meters

Table 18-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X8M Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U33	G5-4	G2-2	2 meters
U31	G5-1	G2-5	2 meters
U29	G3-0	G0-6	2 meters
U27	G3-5	G0-1	2 meters
U25	G4-5	G1-1	2 meters
U23	G4-0	G1-6	2 meters
U22	G1-6	G4-0	2 meters
U21	G2-6	G5-0	2 meters
U20	G1-5	G5-1	2 meters
U18	G1-4	G4-2	2 meters
U17	G2-3	G4-3	2 meters
U16	G2-0	G5-6	2 meters
U14	G1-0	G4-6	2 meters
U12	G0-2	G3-4	2 meters
U10	G0-1	G3-5	2 meters
U8	G0-5	G3-1	2 meters
U6	G2-4	G5-2	2 meters
U4	G0-3	G3-3	2 meters
U2	G1-3	G4-3	2 meters
U1	G1-1	G4-5	2 meters

¹ Un is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections X8M

Table 19-3 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8M 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 18-4 Three-Phase PDU Cabling for the Recovery Appliance X8M Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U41	G4-4	G1-2	2 meters
U39	G5-1	G0-5	2 meters
U37	G5-5	G0-1	2 meters
U35	G4-1	G1-5	2 meters
U33	G5-3	G0-3	2 meters

Table 18-4 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X8M Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U31	G5-4	G0-2	2 meters
U29	G4-6	G1-0	2 meters
U27	G3-4	G2-2	2 meters
U25	G3-1	G2-5	2 meters
U23	G3-0	G2-6	2 meters
U22	G1-6	G4-0	2 meters
U21	G1-4	G4-2	2 meters
U20	G1-5	G4-1	2 meters
U18	G2-6	G3-0	2 meters
U17	G0-5	G5-1	2 meters
U16	G1-1	G4-5	2 meters
U14	G2-2	G3-4	2 meters
U12	G2-5	G3-1	2 meters
U10	G2-4	G3-2	2 meters
U8	G0-2	G5-4	2 meters
U6	G1-0	G4-6	2 meters
U4	G0-0	G5-6	2 meters
U2	G0-4	G5-2	2 meters
U1	G4-2	G1-4	2 meters

¹ Un is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Switch Cabling ConnectionsX8M

For RoCE Network Fabric, [Table 19-7](#) lists the ports and cable lengths for the connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) RDMA Network Fabric cables.

Table 18-5 RDMA Network Fabric Switch-to-Switch Cabling X8M

From RDMA Network Fabric Switch U20 Port ¹	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
04	04	1 meter
05	05	1 meter
06	06	1 meter
07	07	1 meter

Table 18-5 (Cont.) RDMA Network Fabric Switch-to-Switch Cabling X8M

From RDMA Network Fabric Switch U20 Port ¹	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
30	30	1 meter
31	31	1 meter
32	32	1 meter
33	33	3 meter

¹ Un is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Server Cabling Connections X8M

Table 19-8 lists the location, ports, and cables for the InfiniBand connections from the gateway leaf switches to the servers in a Recovery Appliance X8M 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) cables.

Table 18-6 RDMA Network Fabric Switch-to-Server Cabling for the Recovery Appliance X8M Full Racks

From RDMA Network Fabric Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U22	15	U41	Storage server	PCle 1, P2	3 meters
U22	08	U39	Storage server	PCle 1, P2	3 meters
U22	09	U37	Storage server	PCle 1, P2	3 meters
U22	10	U35	Storage server	PCle 1, P2	3 meters
U22	11	U33	Storage server	PCle 1, P2	3 meters
U22	12	U31	Storage server	PCle 1, P2	3 meters
U22	13	U29	Storage server	PCle 1, P2	3 meters
U22	14	U27	Storage server	PCle 1, P2	3 meters
U22	16	U25	Storage server	PCle 1, P2	2 meters
U22	18	U23	Storage server	PCle 1, P2	2 meters
U22	20	U18	Storage server	PCle 1, P2	2 meters
U22	21	U17	Compute server	PCle 1, P2	2 meters
U22	22	U16	Compute server	PCle 1, P2	2 meters
U22	23	U14	Storage server	PCle 1, P2	3 meters
U22	24	U12	Storage server	PCle 1, P2	3 meters
U22	25	U10	Storage server	PCle 1, P2	3 meters
U22	26	U8	Storage server	PCle 1, P2	3 meters

Table 18-6 (Cont.) RDMA Network Fabric Switch-to-Server Cabling for the Recovery Appliance X8M Full Racks

From RDMA Network Fabric Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U22	27	U6	Storage server	PCle 1, P2	3 meters
U22	28	U4	Storage server	PCle 1, P2	3 meters
U22	29	U2	Storage server	PCle 1, P2	3 meters
U20	15	U41	Storage server	PCle 1, P1	3 meters
U20	08	U39	Storage server	PCle 1, P1	3 meters
U20	09	U37	Storage server	PCle 1, P1	3 meters
U20	10	U35	Storage server	PCle 1, P1	3 meters
U20	11	U33	Storage server	PCle 1, P1	3 meters
U20	12	U31	Storage server	PCle 1, P1	3 meters
U20	13	U29	Storage server	PCle 1, P1	3 meters
U20	14	U27	Storage server	PCle 1, P1	3 meters
U20	16	U25	Storage server	PCle 1, P1	2 meters
U20	18	U23	Storage server	PCle 1, P1	2 meters
U20	20	U18	Storage server	PCle 1, P1	2 meters
U20	21	U17	Compute server	PCle 1, P1	2 meters
U20	22	U16	Compute server	PCle 1, P1	2 meters
U20	23	U14	Storage server	PCle 1, P1	2 meters
U20	24	U12	Storage server	PCle 1, P1	3 meters
U20	25	U10	Storage server	PCle 1, P1	3 meters
U20	26	U8	Storage server	PCle 1, P1	3 meters
U20	27	U6	Storage server	PCle 1, P1	3 meters
U20	28	U4	Storage server	PCle 1, P1	3 meters
U20	29	U2	Storage server	PCle 1, P1	3 meters

¹ Un is the unit location in the rack, where *n* is the number.

In-Rack Cabling Tables for Recovery Appliance X8, X7, X6, and X5

This appendix identifies the location of the various components in a Recovery Appliance X8, 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](#)
- [Recovery Appliance X8, X7, X6, and X5 Rack Layout](#)
- [Oracle Integrated Lights Out Manager Cabling Connections](#)
- [Administrative Gigabit Ethernet Cabling Connections](#)
- [Single-Phase Power Distribution Unit Cabling Connections X8](#)
- [Single-Phase Power Distribution Unit Cabling Connections X7 and Earlier](#)
- [Three-Phase Power Distribution Unit Cabling Connections X8](#)
- [Three-Phase Power Distribution Unit Cabling Connections](#)
- [RDMA Network Fabric Switch-to-Switch Cabling Connections](#)
- [RDMA Network Fabric Switch-to-Server Cabling Connections](#)

Cable Color Coding

The cables used in Recovery Appliance are color coded as follows:

- **Black:** RDMA Network Fabric 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

Recovery Appliance X8, X7, X6, and X5 Rack Layout

Figure 19-1 shows the front and rear views of a Recovery Appliance X8, 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 multirack installations. Do not install any other unit in this space.

Figure 19-1 Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Full Rack




Figure 19-2 shows the front and rear views of a Recovery Appliance X8, X7, X6, or X5 minimum configuration.

Figure 19-2 Rack Layout of the Recovery Appliance X8, X7, X6, or X5 Minimum Configuration

Rear View			Front View		
U42	1U Solid Filler	U42	U42	1U Solid Filler	U42
U41	1U Solid Filler	U41	U41	1U Solid Filler	U41
U40	1U Solid Filler	U40	U40	1U Solid Filler	U40
U39	1U Solid Filler	U39	U39	1U Solid Filler	U39
U38	1U Solid Filler	U38	U38	1U Solid Filler	U38
U37	1U Solid Filler	U37	U37	1U Solid Filler	U37
U36	1U Solid Filler	U36	U36	1U Solid Filler	U36
U35	1U Solid Filler	U35	U35	1U Solid Filler	U35
U34	1U Solid Filler	U34	U34	1U Solid Filler	U34
U33	1U Solid Filler	U33	U33	1U Solid Filler	U33
U32	1U Solid Filler	U32	U32	1U Solid Filler	U32
U31	1U Solid Filler	U31	U31	1U Solid Filler	U31
U30	1U Solid Filler	U30	U30	1U Solid Filler	U30
U29	1U Solid Filler	U29	U29	1U Solid Filler	U29
U28	1U Solid Filler	U28	U28	1U Solid Filler	U28
U27	1U Solid Filler	U27	U27	1U Solid Filler	U27
U26	1U Solid Filler	U26	U26	1U Solid Filler	U26
U25	1U Solid Filler	U25	U25	1U Solid Filler	U25
U24	1U Solid Filler	U24	U24	1U Solid Filler	U24
U23	1U Solid Filler	U23	U23	1U Solid Filler	U23
U22	RDMA Network Fabric	U22	U22	1U Vented Filler	U22
U21	Ethernet Switch	U21	U21	1U Vented Filler	U21
U20	RDMA Network Fabric	U20	U20	1U Vented Filler	U20
U19	1U Solid Filler	U19	U19	1U Solid Filler	U19
U18	1U Solid Filler	U18	U18	1U Solid Filler	U18
U17	Compute Server	U17	U17	Compute Server	U17
U16	Compute Server	U16	U16	Compute Server	U16
U15	1U Solid Filler	U15	U15	1U Solid Filler	U15
U14	1U Solid Filler	U14	U14	1U Solid Filler	U14
U13	1U Solid Filler	U13	U13	1U Solid Filler	U13
U12	1U Solid Filler	U12	U12	1U Solid Filler	U12
U11	1U Solid Filler	U11	U11	1U Solid Filler	U11
U10	1U Solid Filler	U10	U10	1U Solid Filler	U10
U9	1U Solid Filler	U9	U9	1U Solid Filler	U9
U8	1U Solid Filler	U8	U8	1U Solid Filler	U8
U7	Storage Server	U7	U7	Storage Server	U7
U6	Storage Server	U6	U6	Storage Server	U6
U5	Storage Server	U5	U5	Storage Server	U5
U4	Storage Server	U4	U4	Storage Server	U4
U3	Storage Server	U3	U3	Storage Server	U3
U2	Storage Server	U2	U2	Storage Server	U2
U1	1U Solid Filler	U1	U1	1U Solid Filler	U1

Oracle Integrated Lights Out Manager Cabling Connections

Table 19-1 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 19-1 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.

Administrative Gigabit Ethernet Cabling Connections

Table 19-2 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 19-2 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

Table 19-2 (Cont.) Gigabit Ethernet Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

From Rack Unit ¹	Component	To Ethernet Port
U33	Storage server	7
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.

Single-Phase Power Distribution Unit Cabling Connections X8

Table 19-3 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8 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 19-3 Single-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U41	G3-6	G2-0	2 meters
U39	G5-4	G0-2	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-2	G0-4	2 meters
U33	G5-0	G0-6	2 meters
U31	G4-2	G1-4	2 meters
U29	G4-1	G1-5	2 meters

Table 19-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U27	G4-0	G1-6	2 meters
U25	G3-4	G2-2	2 meters
U23	G2-5	G3-1	2 meters
U22	G0-5	G5-1	2 meters
U21	G1-2	G4-4	2 meters
U20	G0-4	G5-2	2 meters
U18	G2-4	G3-2	2 meters
U17	G1-1	G4-5	2 meters
U16	G0-3	G5-3	2 meters
U14	G2-3	G3-3	2 meters
U12	G2-2	G3-4	2 meters
U10	G2-1	G3-5	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-2	G5-4	2 meters
U4	G0-1	G5-5	2 meters
U2	G0-0	G5-6	2 meters
U1	G2-0	G3-6	2 meters

¹ Un is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cabling Connections X7 and Earlier

Table 19-4 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 19-4 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
U33	G5-0	G0-6	2 meters
U31	G4-4	G1-2	2 meters

Table 19-4 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
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.

Three-Phase Power Distribution Unit Cabling Connections X8

Table 19-3 shows the cable connections from the power distribution units (PDUs) to the rack units in a Recovery Appliance X8 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 19-5 Three-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U41	G5-5	G2-1	2 meters
U39	G3-0	G0-6	2 meters
U37	G5-6	G2-0	2 meters
U35	G4-5	G1-1	2 meters
U33	G4-4	G1-2	2 meters
U31	G3-5	G0-1	2 meters
U29	G4-1	G1-5	2 meters

Table 19-5 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X8 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
U27	G5-0	G2-6	2 meters
U25	G4-0	G1-6	2 meters
U23	G2-6	G5-0	2 meters
U22	G2-3	G5-3	2 meters
U21	G1-3	G4-3	2 meters
U20	G0-3	G3-3	2 meters
U18	G0-6	G3-0	2 meters
U17	G1-2	G4-4	2 meters
U16	G1-1	G4-5	2 meters
U14	G2-2	G5-4	2 meters
U12	G0-2	G3-4	2 meters
U10	G0-1	G3-5	2 meters
U8	G2-1	G5-5	2 meters
U6	G2-0	G5-6	2 meters
U4	G1-0	G4-6	2 meters
U2	G0-0	G3-6	2 meters
U1	G3-2	G0-4	2 meters

¹ Un is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cabling Connections

Table 19-6 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 19-6 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
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

Table 19-6 (Cont.) Three-Phase PDU Cabling for the Recovery Appliance X7, X6, or X5 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
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.

RDMA Network Fabric Switch-to-Switch Cabling Connections

For InfiniBand Network Fabric, [Table 19-7](#) lists the ports and cable lengths for the connections between the two gateway leaf switches, which are located in rack units 20 and 22. All cables are black quad data rate (QDR) RDMA Network Fabric cables.

Table 19-7 RDMA Network Fabric Switch-to-Switch Cabling

From RDMA Network Fabric Switch U20 Port ¹	To RDMA Network Fabric Switch U22 Port unless specified otherwise	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
U20 8B	U01 1B	3 meter

Table 19-7 (Cont.) RDMA Network Fabric Switch-to-Switch Cabling

From RDMA Network Fabric Switch U20 Port ¹	To RDMA Network Fabric Switch U22 Port unless specified otherwise	Cable Length
U01 0B	U22 8B	3 meter

¹ Un is the unit location in the rack, where *n* is the number.

RDMA Network Fabric Switch-to-Server Cabling Connections

For InfiniBand Network Fabric, [Table 19-8](#) lists the location, ports, and cables for the connections from the gateway leaf switches to the servers in a Recovery Appliance X8, X7, 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) cables.

Table 19-8 RDMA Network Fabric Switch-to-Server Cabling for the Recovery Appliance X8, X7, X6 or X5 Full Racks

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U22	6B	U41	Storage server	PCle 7, P1	3 meters
U22	0A	U39	Storage server	PCle 7, P1	3 meters
U22	0B	U37	Storage server	PCle 7, P1	3 meters
U22	1A	U35	Storage server	PCle 7, P1	3 meters
U22	1B	U33	Storage server	PCle 7, P1	3 meters
U22	2A	U31	Storage server	PCle 7, P1	3 meters
U22	2B	U29	Storage server	PCle 7, P1	3 meters
U22	3B	U27	Storage server	PCle 7, P1	3 meters
U22	4A	U25	Storage server	PCle 7, P1	2 meters
U22	5A	U23	Storage server	PCle 7, P1	2 meters
U22	13A	U18	Storage server	PCle 7, P2	2 meters
U22	13B	U17	Compute server	PCle 7, P2	2 meters
U22	14B	U16	Compute server	PCle 7, P2	2 meters
U22	14A	U14	Storage server	PCle 7, P2	3 meters
U22	15A	U12	Storage server	PCle 7, P2	3 meters
U22	15B	U10	Storage server	PCle 7, P2	3 meters
U22	16A	U8	Storage server	PCle 7, P2	3 meters
U22	16B	U6	Storage server	PCle 7, P2	3 meters
U22	17A	U4	Storage server	PCle 7, P2	3 meters
U22	17B	U2	Storage server	PCle 7, P2	3 meters
U20	6B	U41	Storage server	PCle 7, P2	3 meters
U20	0A	U39	Storage server	PCle 7, P2	3 meters

Table 19-8 (Cont.) RDMA Network Fabric Switch-to-Server Cabling for the Recovery Appliance X8, X7, X6 or X5 Full Racks

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit	Device	Port	Cable Length
U20	0B	U37	Storage server	PCle 7, P2	3 meters
U20	1A	U35	Storage server	PCle 7, P2	3 meters
U20	1B	U33	Storage server	PCle 7, P2	3 meters
U20	2A	U31	Storage server	PCle 7, P2	3 meters
U20	2B	U29	Storage server	PCle 7, P2	3 meters
U20	3B	U27	Storage server	PCle 7, P2	3 meters
U20	4A	U25	Storage server	PCle 7, P2	2 meters
U20	5A	U23	Storage server	PCle 7, P2	2 meters
U20	13A	U18	Storage server	PCle 7, P1	2 meters
U20	13B	U17	Compute server	PCle 2, P1	2 meters
U20	14B	U16	Compute server	PCle 2, P1	2 meters
U20	14A	U14	Storage server	PCle 7, P1	2 meters
U20	15A	U12	Storage server	PCle 7, P1	3 meters
U20	15B	U10	Storage server	PCle 7, P1	3 meters
U20	16A	U8	Storage server	PCle 7, P1	3 meters
U20	16B	U6	Storage server	PCle 7, P1	3 meters
U20	17A	U4	Storage server	PCle 7, P1	3 meters
U20	17B	U2	Storage server	PCle 7, P1	3 meters

¹ Un is the unit location in the rack, where n is the number.

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](#)
- [Oracle Integrated Lights Out Manager Cabling Connections](#)
- [Administrative Gigabit Ethernet Cabling Connections](#)
- [Single-Phase Power Distribution Unit Cabling Connections](#)
- [Three-Phase Power Distribution Unit Cabling Connections](#)
- [InfiniBand Switch-to-Switch Cabling Connections](#)
- [InfiniBand Switch-to-Server Cabling Connections](#)

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

Recovery Appliance X4 Rack Layout

[Figure 20-1](#) 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 20-1 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		41		1U Solid Filler	
40		Storage Server		40		Storage Server	
39		Storage Server		39		Storage Server	
38		Storage Server		38		Storage Server	
37		Storage Server		37		Storage Server	
36		Storage Server		36		Storage Server	
35		Storage Server		35		Storage Server	
34		Storage Server		34		Storage Server	
33		Storage Server		33		Storage Server	
32		Storage Server		32		Storage Server	
31		Storage Server		31		Storage Server	
30	INP1	Storage Server	INP1	30	INP1	Storage Server	INP1
29		Storage Server		29		Storage Server	
28		Storage Server		28		Storage Server	
27		1U Solid Filler		27		1U Solid Filler	
26		1U Solid Filler		26		1U Solid Filler	
25		1U Solid Filler		25		1U Solid Filler	
24		1U Solid Filler		24		1U Solid Filler	
23		1U Solid Filler		23		1U Solid Filler	
22		Infiniband Gateway Switch		22		1U Vented Filler	
21		Ethernet Switch		21		1U Vented Filler	
20		Infiniband Gateway Switch		20		1U Vented Filler	
19		1U Solid Filler		19		1U Solid Filler	
18	INP0	1U Solid Filler	INP2	18	INP2	1U Solid Filler	INP0
17		Compute Server		17		Compute Server	
16		Compute Server		16		Compute Server	
15		Storage Server		15		Storage Server	
14		Storage Server		14		Storage Server	
13		Storage Server		13		Storage Server	
12		Storage Server		12		Storage Server	
11		Storage Server		11		Storage Server	
10		Storage Server		10		Storage Server	
9		Storage Server		9		Storage Server	
8		Storage Server		8		Storage Server	
7		Storage Server		7		Storage Server	
6		Storage Server		6		Storage Server	
5		Storage Server		5		Storage Server	
4		Storage Server		4		Storage Server	
3		Storage Server		3		Storage Server	
2		1U Solid Filler		2		1U Solid Filler	
1		1U Solid Filler		1		1U Solid Filler	

Figure 20-2 shows the front and rear views of a Recovery Appliance X4 minimum configuration.

Figure 20-2 Rack Layout of the Recovery Appliance X4 Minimum Configuration

Rear View			Front View		
U42	1U Solid Filler	U42	U42	1U Solid Filler	U42
U41	1U Solid Filler	U41	U41	1U Solid Filler	U41
U40	1U Solid Filler	U40	U40	1U Solid Filler	U40
U39	1U Solid Filler	U39	U39	1U Solid Filler	U39
U38	1U Solid Filler	U38	U38	1U Solid Filler	U38
U37	1U Solid Filler	U37	U37	1U Solid Filler	U37
U36	1U Solid Filler	U36	U36	1U Solid Filler	U36
U35	1U Solid Filler	U35	U35	1U Solid Filler	U35
U34	1U Solid Filler	U34	U34	1U Solid Filler	U34
U33	1U Solid Filler	U33	U33	1U Solid Filler	U33
U32	1U Solid Filler	U32	U32	1U Solid Filler	U32
U31	1U Solid Filler	U31	U31	1U Solid Filler	U31
U30	1U Solid Filler	U30	U30	1U Solid Filler	U30
U29	1U Solid Filler	U29	U29	1U Solid Filler	U29
U28	1U Solid Filler	U28	U28	1U Solid Filler	U28
U27	1U Solid Filler	U27	U27	1U Solid Filler	U27
U26	1U Solid Filler	U26	U26	1U Solid Filler	U26
U25	1U Solid Filler	U25	U25	1U Solid Filler	U25
U24	1U Solid Filler	U24	U24	1U Solid Filler	U24
U23	1U Solid Filler	U23	U23	1U Solid Filler	U23
U22	RDMA Network Fabric	U22	U22	1U Vented Filler	U22
U21	Ethernet Switch	U21	U21	1U Vented Filler	U21
U20	RDMA Network Fabric	U20	U20	1U Vented Filler	U20
U19	1U Solid Filler	U19	U19	1U Solid Filler	U19
U18	1U Solid Filler	U18	U18	1U Solid Filler	U18
U17	Compute Server	U17	U17	Compute Server	U17
U16	Compute Server	U16	U16	Compute Server	U16
U15	1U Solid Filler	U15	U15	1U Solid Filler	U15
U14	1U Solid Filler	U14	U14	1U Solid Filler	U14
U13	1U Solid Filler	U13	U13	1U Solid Filler	U13
U12	1U Solid Filler	U12	U12	1U Solid Filler	U12
U11	1U Solid Filler	U11	U11	1U Solid Filler	U11
U10	1U Solid Filler	U10	U10	1U Solid Filler	U10
U9	1U Solid Filler	U9	U9	1U Solid Filler	U9
U8	1U Solid Filler	U8	U8	1U Solid Filler	U8
U7	Storage Server	U7	U7	Storage Server	U7
U6	Storage Server	U6	U6	Storage Server	U6
U5	Storage Server	U5	U5	Storage Server	U5
U4	Storage Server	U4	U4	Storage Server	U4
U3	Storage Server	U3	U3	Storage Server	U3
U2	Storage Server	U2	U2	Storage Server	U2
U1	1U Solid Filler	U1	U1	1U Solid Filler	U1

Oracle Integrated Lights Out Manager Cabling Connections

Table 20-1 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 20-1 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.

Administrative Gigabit Ethernet Cabling Connections

Table 20-2 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 20-2 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

Table 20-2 (Cont.) Gigabit Ethernet Cabling for the Recovery Appliance X4 Full Racks

From Rack Unit ¹	Component	To Ethernet Port
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.

Single-Phase Power Distribution Unit Cabling Connections

Table 20-3 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 20-3 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
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters

Table 20-3 (Cont.) Single-Phase PDU Cabling for the Recovery Appliance X4 Full Racks

Rack Unit ¹	PDU-A	PDU-B	Cable Length
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.

Three-Phase Power Distribution Unit Cabling Connections

Table 20-4 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 20-4 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.

InfiniBand Switch-to-Switch Cabling Connections

Table 20-5 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 20-5 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.

InfiniBand Switch-to-Server Cabling Connections

Table 20-6 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 20-6 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	PCle 3, P1	3 meters
U22	0B	U37	Storage server	PCle 3, P1	3 meters
U22	1A	U35	Storage server	PCle 3, P1	3 meters
U22	1B	U33	Storage server	PCle 3, P1	3 meters
U22	2A	U31	Storage server	PCle 3, P1	3 meters
U22	2B	U29	Storage server	PCle 3, P1	3 meters
U22	3B	U27	Storage server	PCle 3, P1	3 meters
U22	13B	U17	Compute server	PCle 2, P2	2 meters
U22	14B	U16	Compute server	PCle 2, P2	2 meters
U22	14A	U14	Storage server	PCle 3, P2	3 meters
U22	15A	U12	Storage server	PCle 3, P2	3 meters
U22	15B	U10	Storage server	PCle 3, P2	3 meters

Table 20-6 (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	16A	U8	Storage server	PCle 3, P2	3 meters
U22	16B	U6	Storage server	PCle 3, P2	3 meters
U22	17A	U4	Storage server	PCle 3, P2	3 meters
U22	17B	U2	Storage server	PCle 3, P2	3 meters
U20	0A	U39	Storage server	PCle 3, P2	3 meters
U20	0B	U37	Storage server	PCle 3, P2	3 meters
U20	1A	U35	Storage server	PCle 3, P2	3 meters
U20	1B	U33	Storage server	PCle 3, P2	3 meters
U20	2A	U31	Storage server	PCle 3, P2	3 meters
U20	2B	U29	Storage server	PCle 3, P2	3 meters
U20	3B	U27	Storage server	PCle 3, P2	3 meters
U20	13B	U17	Compute server	PCle 2, P1	2 meters
U20	14B	U16	Compute server	PCle 2, P1	2 meters
U20	14A	U14	Storage server	PCle 3, P1	2 meters
U20	15A	U12	Storage server	PCle 3, P1	3 meters
U20	15B	U10	Storage server	PCle 3, P1	3 meters
U20	16A	U8	Storage server	PCle 3, P1	3 meters
U20	16B	U6	Storage server	PCle 3, P1	3 meters
U20	17A	U4	Storage server	PCle 3, P1	3 meters
U20	17B	U2	Storage server	PCle 3, P1	3 meters

¹ Un is the unit location in the rack, where *n* is the number.

Part VI

Extending Oracle Zero Data Loss Recovery Appliance

You can extend Oracle Zero Data Loss Recovery Appliance either by adding servers to the current configuration or by cabling together multiple racks.

Here are considerations when extending Recovery Appliance racks:

- You can extend Recovery Appliance from a fixed or custom configuration to another configuration by adding storage servers up to the allowed maximum.
- You can cable together multiple Recovery Appliance racks subject to the following:
 - Oracle Zero Data Loss Recovery Appliance RA23 can be multi-racked with RA21 or X8M.

Oracle Zero Data Loss Recovery Appliance X8-2 can be multi-racked with X7-2 and earlier ZDLRA Racks.
 - All racks in the multi-rack configuration must use the same RDMA Network Fabric.

For X8M and later, all racks must use RoCE Network Fabric.

For X8 and earlier, all racks must use InfiniBand Network Fabric.
 - The racks can have different storage disk drives:
 - * RA23 have 22 TB.
 - * RA21 have 18 TB
 - * X8M have 14 TB
- Prior to extending a system across multiple racks, you must acquire the appropriate RDMA Network Fabric switches and transceivers.

Multiple Recovery Appliances can run as separate environments while sharing the RDMA Network Fabric. If you are planning to utilize multiple Recovery Appliances in this manner, then note the following:

Extending a Rack by Adding Another Rack

You can extend your ZDLRA Rack by adding another rack and configuring the racks together.

Overview of Adding Another Rack to an Existing System

Review the following notes before cabling racks together.

- The procedures for extending racks with RoCE Network Fabric (X8M and later) are different than the procedures for racks with InfiniBand Network Fabric (X8 and earlier.)
- Racks with InfiniBand Network Fabric can be cabled together with no downtime. Depending on the procedure being used, racks with RoCE Network Fabric might require downtime when cabling racks together.
- Cabling within a live network must be done carefully in order to avoid potentially serious disruptions.
- There can be performance degradation while cabling the racks together. This degradation results from data retransmission due to packet loss and reduced network bandwidth when a cable is unplugged.
- Redundancy with the RDMA Network Fabric can be compromised while cabling the racks together. This occurs whenever the RDMA Network Fabric ports or switches are taken offline and all traffic must use the remaining switches.
- Only the existing racks are operational when adding racks. It is assumed that the servers on any new racks are initially powered down.
- The software running on the systems cannot have problems related to RDMA Network Fabric restarts. To verify the configuration, run `infinicheck` separately on each rack before connecting multiple racks together.
- It is assumed that each ZDLRA Rack has three RDMA Network Fabric switches already installed.
- The new racks have been configured with the appropriate IP addresses to be migrated into the expanded system prior to any cabling, and there are no duplicate IP addresses.
- Racks with RoCE Network Fabric use one loopback IP interface on each spine switch and two loopback IP interfaces on each leaf switch. The IP addressing scheme uses IANA 'Shared Address Space' 100.64.0.0/10. This ensures that there is no overlap with IPv4 addresses in the network using other schemes.
 - Leaf loopback0 IPs are assigned as 100.64.0.101, 100.64.0.102, 100.64.0.103, and so on.
 - Leaf loopback1 IPs are assigned as 100.64.1.101, 100.64.1.102, 100.64.1.103, and so on.
 - Spine loopback0 IPs are assigned as 100.64.0.201, 100.64.0.202, up to 100.64.0.208.

Cabling Two Racks Together

Choose from the available methods based on your system specifications and operational requirements.

Cabling Two RoCE Network Fabric Racks Together with No Down Time

If your operational requirements cannot tolerate any scheduled down time, then choose from the following procedures to extend your existing RoCE Network Fabric rack by adding another rack.

Extending an RA21 or Later Model Rack with No Down Time by Adding Another RA21 or Later Model Rack

WARNING:

Take time to read and understand this procedure before implementation. Pay careful attention to the instructions that surround the command examples. A system outage may occur if the procedure is not applied correctly.

Note:

For additional background information, see [Understanding Multi-Rack Cabling for RA21 and Later Model Racks](#).

Use this procedure to extend a typical RA21 or later model rack by cabling it together with a second RA21 or later model rack. The primary rack (designated R1) and all of the systems it supports remain online throughout the procedure. At the beginning of the procedure, the additional rack (designated R2) is shut down.

The following is an outline of the procedure:

- Preparation (steps [1](#) and [2](#))
In this phase, you prepare the racks, switches, and cables. Also, you install and cable the spine switches in both racks.
- Configuration and Physical Cabling
In this phase, you reconfigure the leaf switches and finalize the cabling to the spine switches. These tasks are carefully orchestrated to avoid downtime on the primary system, as follows:
 - Partially configure the lower leaf switches (step [3](#))
In this step, you reconfigure the switch ports on the lower leaf switches. There is no physical cabling performed in this step.
 - Partially configure the upper leaf switches (step [4](#))

In this step, you reconfigure the switch ports on the upper leaf switches, remove the inter-switch cables that connect the leaf switches in both racks and connect the cables between the upper leaf switches and the spine switches.

- Finalize the lower leaf switches (step 5)

In this step, you finalize the switch port configuration on the lower leaf switches. You also complete the physical cabling by connecting the cables between the lower leaf switches and the spine switches.

- Finalize the upper leaf switches (step 6)

In this step, you finalize the switch port configuration on the upper leaf switches.

- Validation and Testing (steps 7 and 8)

In this phase, you validate and test the RoCE Network Fabric across both of the interconnect racks.

After completing the procedure, both racks share the RoCE Network Fabric, and the combined system is ready for further configuration. For example, you can extend existing disk groups and Oracle RAC databases to consume resources across both racks.

 Note:

- This procedure applies only to typical rack configurations that initially have leaf switches with the following specifications:
 - The inter-switch ports are ports 4 to 7, and ports 30 to 33.
 - The storage server ports are ports 8 to 14, and ports 23 to 29.
 - The database server ports are ports 15 to 22.

For other rack configurations (for example, X9M-8 systems with three database servers and 11 storage servers) a different procedure and different RoCE Network Fabric switch configuration files are required. Contact Oracle for further guidance.

- The procedure uses the following naming abbreviations and conventions:
 - The abbreviation for the existing rack is R1, and the new rack is R2.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies a spine switch.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the existing rack (R1).
- Most operations must be performed in multiple locations. For example, step [1.h](#) instructs you to update the firmware on all the RoCE Network Fabric leaf switches (R1LL, R1UL, R2LL, and R2UL). Pay attention to the instructions and keep track of your actions.

 Tip:

When a step must be performed on multiple switches, the instruction contains a list of the applicable switches. For example, (R1LL, R1UL, R2LL, and R2UL). You can use this list as a checklist to keep track of your actions.

- Perform operations sequentially, and complete every operation before proceeding. For example, run the entire command sequence at [3.a.i](#) as one operation and complete it before proceeding.
- All of commands that are run on a RoCE Network Fabric switch must be run while connected to the switch management interface as the switch administrator.

1. Prepare the systems.

- a. Position the new rack (R2) so that it is physically near the existing rack (R1).

The RDMA Network Fabric cables must be able to reach the switches in each rack.

For the required cross-rack cabling information, see [Two-Rack Cabling for RA21 and Later Model Racks](#).

- b. Power on all of the servers and network switches in the new rack (R2).

This includes the database servers, storage servers, RoCE Network Fabric leaf switches, and the Management Network Switch.

- c. Prepare the RoCE Network Fabric cables that you will use to interconnect the racks.
Label both ends of every cable.
For the required cross-rack cabling information, see [Two-Rack Cabling for RA21 and Later Model Racks](#).
- d. Connect the new rack (R2) to your existing management network.
Ensure that there are no IP address conflicts across the racks and that you can access the management interfaces on the RoCE Network Fabric switches.
- e. Ensure that you have a backup of the current switch configuration for each RoCE Network Fabric switch (R1LL, R1UL, R2LL, and R2UL).
See Backing Up Settings on the RoCE Network Fabric Switch in *Oracle Exadata Database Machine Maintenance Guide*.
- f. Download the required RoCE Network Fabric switch configuration files.
This procedure requires specific RoCE Network Fabric switch configuration files, which you must download from [My Oracle Support document 2704997.1](#).

 WARNING:

You must use different switch configuration files depending on whether your system uses Exadata Secure RDMA Fabric Isolation. Ensure that you download the correct archive that matches your system configuration.

For system configurations without Secure Fabric, download `online_multi-rack_14uplinks.zip`. For system configurations with Secure Fabric, download `online_SF_enabled_multi-rack_14uplinks.zip`.

Download and extract the archive containing the required RoCE Network Fabric switch configuration files. Place the files on a server with access to the management interfaces on the RoCE Network Fabric switches.

- g. Copy the required RoCE Network Fabric switch configuration files to the leaf switches on both racks.

You can use the following commands to copy the required configuration files to all of the RoCE Network Fabric switches on a system without Secure Fabric enabled:

- i. `# scp roce_multi_14uplinks_online_step3_R1_LL.cfg admin@R1LL_IP:/`
- ii. `# scp roce_multi_14uplinks_online_step3_R2_LL.cfg admin@R2LL_IP:/`
- iii. `# scp roce_multi_14uplinks_online_step4_R1_UL.cfg admin@R1UL_IP:/`
- iv. `# scp roce_multi_14uplinks_online_step4_R2_UL.cfg admin@R2UL_IP:/`
- v. `# scp roce_multi_14uplinks_online_step5.cfg admin@R1LL_IP:/`
- vi. `# scp roce_multi_14uplinks_online_step5.cfg admin@R2LL_IP:/`

On a system with Secure Fabric enabled, you can use the following commands:

- i. # scp roce_SF_multi_14uplinks_online_step3_R1_LL.cfg admin@R1LL_IP:/
- ii. # scp roce_SF_multi_14uplinks_online_step3_R2_LL.cfg admin@R2LL_IP:/
- iii. # scp roce_SF_multi_14uplinks_online_step4_R1_UL.cfg admin@R1UL_IP:/
- iv. # scp roce_SF_multi_14uplinks_online_step4_R2_UL.cfg admin@R2UL_IP:/
- v. # scp roce_SF_multi_14uplinks_online_step5.cfg admin@R1LL_IP:/
- vi. # scp roce_SF_multi_14uplinks_online_step5.cfg admin@R2LL_IP:/

In the above commands, substitute the appropriate IP address or host name where applicable. For example, in place of *R1LL_IP*, substitute the management IP address or host name for the lower leaf switch (LL) on the existing rack (R1).

 Note:

The command examples in the rest of this procedure use the configuration files for a system configuration without Secure Fabric enabled. If required, adjust the commands to use the Secure Fabric-enabled switch configuration files.

- h. Update the firmware to the latest available release on all of the RoCE Network Fabric leaf switches (R1LL, R1UL, R2LL, and R2UL).

See Updating RoCE Network Fabric Switch Firmware in *Oracle Exadata Database Machine Maintenance Guide*.

- i. Examine the RoCE Network Fabric leaf switches (R1LL, R1UL, R2LL, and R2UL) and confirm the port categories for the cabled ports.

Run the `show interface status` command on every RoCE Network Fabric leaf switch:

- i. R1LL# **show interface status**
- ii. R1UL# **show interface status**
- iii. R2LL# **show interface status**
- iv. R2UL# **show interface status**

Examine the output and confirm the port categories as follows:

- Confirm that the inter-switch ports are ports 4 to 7, and ports 30 to 33.
- Confirm that the storage server ports are ports 8 to 14, and ports 23 to 29.
- Confirm that the database server ports are ports 15 to 22.

For example:

R1LL# **show interface status**

```

-----
-----
Port          Name                Status      Vlan      Duplex  Speed  Type
-----
-----
mgmt0         --                  connected   routed    full    1000   --
-----
-----
Port          Name                Status      Vlan      Duplex  Speed  Type
-----
-----
Eth1/1        --                  xcvrAbsen  1          auto    auto   --
Eth1/2        --                  xcvrAbsen  1          auto    auto   --
Eth1/3        --                  xcvrAbsen  1          auto    auto   --
Eth1/4        ISL1                connected   trunk      full    100G
QSFP-100G-CR4
Eth1/5        ISL2                connected   trunk      full    100G
QSFP-100G-CR4
Eth1/6        ISL3                connected   trunk      full    100G
QSFP-100G-CR4
Eth1/7        ISL4                connected   trunk      full    100G
QSFP-100G-CR4
Eth1/8        celadm14            connected   3888       full    100G
QSFP-100G-CR4
Eth1/9        celadm13            connected   3888       full    100G
QSFP-100G-CR4
Eth1/10       celadm12            connected   3888       full    100G
QSFP-100G-CR4
Eth1/11       celadm11            connected   3888       full    100G
QSFP-100G-CR4
Eth1/12       celadm10            connected   3888       full    100G
QSFP-100G-CR4
Eth1/13       celadm09            connected   3888       full    100G
QSFP-100G-CR4
Eth1/14       celadm08            connected   3888       full    100G
QSFP-100G-CR4
Eth1/15       adm08               connected   3888       full    100G
QSFP-100G-CR4
Eth1/16       adm07               connected   3888       full    100G
QSFP-100G-CR4
Eth1/17       adm06               connected   3888       full    100G
QSFP-100G-CR4
Eth1/18       adm05               connected   3888       full    100G
QSFP-100G-CR4
Eth1/19       adm04               connected   3888       full    100G
QSFP-100G-CR4
Eth1/20       adm03               connected   3888       full    100G
QSFP-100G-CR4
Eth1/21       adm02               connected   3888       full    100G
QSFP-100G-CR4
Eth1/22       adm01               connected   3888       full    100G

```

QSFP-100G-CR4						
Eth1/23	celadm07	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

- j. For each rack (R1 and R2), confirm the RoCE Network Fabric cabling by running the `verify_roce_cables.py` script.

The `verify_roce_cables.py` script uses two input files; one for database servers and storage servers (`nodes.rackN`), and another for switches (`switches.rackN`). In each file, every server or switch must be listed on separate lines. Use fully qualified domain names or IP addresses for each server and switch.

See My Oracle Support document 2587717.1 for download and detailed usage instructions.

Run the `verify_roce_cables.py` script against both of the racks:

- i.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
```
- ii.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2
```

Check that output in the `CABLE OK?` columns contains the `OK` status.

The following example shows the expected command results:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
```

```

SWITCH PORT (EXPECTED PEER)                LOWER LEAF (rack1sw-rocea0) :
CABLE OK?                UPPER LEAF (rack1sw-roceb0) : CABLE OK?
-----
Eth1/4 (ISL peer switch) : rack1sw-rocea0 Ethernet1/4 :
OK rack1sw-roceb0 Ethernet1/4 : OK
Eth1/5 (ISL peer switch) : rack1sw-rocea0 Ethernet1/5 :
OK rack1sw-roceb0 Ethernet1/5 : OK
Eth1/6 (ISL peer switch) : rack1sw-rocea0 Ethernet1/6 :
OK rack1sw-roceb0 Ethernet1/6 : OK
Eth1/7 (ISL peer switch) : rack1sw-rocea0 Ethernet1/7 :
OK rack1sw-roceb0 Ethernet1/7 : OK
Eth1/8 (RU39) : rack1celadm14 port-1 :
OK rack1celadm14 port-2 : OK
Eth1/9 (RU37) : rack1celadm13 port-1 :
OK rack1celadm13 port-2 : OK
Eth1/10 (RU35) : rack1celadm12 port-1 :
OK rack1celadm12 port-2 : OK
Eth1/11 (RU33) : rack1celadm11 port-1 :
OK rack1celadm11 port-2 : OK
Eth1/12 (RU31) : rack1celadm10 port-1 :
OK rack1celadm10 port-2 : OK
Eth1/13 (RU29) : rack1celadm09 port-1 :
OK rack1celadm09 port-2 : OK
Eth1/14 (RU27) : rack1celadm08 port-1 :
OK rack1celadm08 port-2 : OK
Eth1/15 (RU26) : rack1adm08 port-1 :
OK rack1adm08 port-2 : OK
Eth1/16 (RU25) : rack1adm07 port-1 :
OK rack1adm07 port-2 : OK
Eth1/17 (RU24) : rack1adm06 port-1 :
OK rack1adm06 port-2 : OK
Eth1/18 (RU23) : rack1adm05 port-1 :
OK rack1adm05 port-2 : OK
Eth1/19 (RU19) : rack1adm04 port-1 :
OK rack1adm04 port-2 : OK
Eth1/20 (RU18) : rack1adm03 port-1 :
OK rack1adm03 port-2 : OK
Eth1/21 (RU17) : rack1adm02 port-1 :
OK rack1adm02 port-2 : OK
Eth1/22 (RU16) : rack1adm01 port-1 :
OK rack1adm01 port-2 : OK
Eth1/23 (RU14) : rack1celadm07 port-1 :
OK rack1celadm07 port-2 : OK
Eth1/24 (RU12) : rack1celadm06 port-1 :
OK rack1celadm06 port-2 : OK
Eth1/25 (RU10) : rack1celadm05 port-1 :
OK rack1celadm05 port-2 : OK
Eth1/26 (RU08) : rack1celadm04 port-1 :
OK rack1celadm04 port-2 : OK
Eth1/27 (RU06) : rack1celadm03 port-1 :
OK rack1celadm03 port-2 : OK
Eth1/28 (RU04) : rack1celadm02 port-1 :
OK rack1celadm02 port-2 : OK
Eth1/29 (RU02) : rack1celadm01 port-1 :
OK rack1celadm01 port-2 : OK

```

```

Eth1/30 (ISL peer switch) : rack1sw-rocea0 Ethernet1/30 :
OK rack1sw-roceb0 Ethernet1/30 : OK
Eth1/31 (ISL peer switch) : rack1sw-rocea0 Ethernet1/31 :
OK rack1sw-roceb0 Ethernet1/31 : OK
Eth1/32 (ISL peer switch) : rack1sw-rocea0 Ethernet1/32 :
OK rack1sw-roceb0 Ethernet1/32 : OK
Eth1/33 (ISL peer switch) : rack1sw-rocea0 Ethernet1/33 :
OK rack1sw-roceb0 Ethernet1/33 : OK

```

- k. For each rack (R1 and R2), verify the RoCE Network Fabric operation by using the `infinicheck` command.

- Use `infinicheck` with the `-z` option to clear the files that were created during the last run of the `infinicheck` command.
- Use `infinicheck` with the `-s` option to set up user equivalence for password-less SSH across the RoCE Network Fabric.
- Finally, verify the RoCE Network Fabric operation by using `infinicheck` with the `-b` option, which is recommended on newly imaged machines where it is acceptable to suppress the `cellip.ora` and `cellinit.ora` configuration checks.

In each command, the hosts input file (`hosts.rack1` and `hosts.rack2`) contains a list of database server RoCE Network Fabric IP addresses (2 RoCE Network Fabric IP addresses for each database server), and the cells input file (`cells.rack1` and `cells.rack2`) contains a list of RoCE Network Fabric IP addresses for the storage servers (2 RoCE Network Fabric IP addresses for each storage server).

- i. Use the following recommended command sequence on the existing rack (R1):

```

i. # cd /opt/oracle.SupportTools/ibdiagtools
   # ./infinicheck -g hosts.rack1 -c cells.rack1 -z

ii. # cd /opt/oracle.SupportTools/ibdiagtools
    # ./infinicheck -g hosts.rack1 -c cells.rack1 -s

iii. # cd /opt/oracle.SupportTools/ibdiagtools
     # ./infinicheck -g hosts.rack1 -c cells.rack1 -b

```

- ii. Use the following recommended command sequence on the new rack (R2):

```

i. # cd /opt/oracle.SupportTools/ibdiagtools
   # ./infinicheck -g hosts.rack2 -c cells.rack2 -z

ii. # cd /opt/oracle.SupportTools/ibdiagtools
    # ./infinicheck -g hosts.rack2 -c cells.rack2 -s

iii. # cd /opt/oracle.SupportTools/ibdiagtools
     # ./infinicheck -g hosts.rack2 -c cells.rack2 -b

```

The following example shows the expected command results for the final command in the sequence:

```

# cd /opt/oracle.SupportTools/ibdiagtools
# ./infinicheck -g hosts.rackN -c cells.rackN -b

```

```

INFINICHECK
[Network Connectivity, Configuration and Performance]

#### FABRIC TYPE TESTS ####
System type identified: RoCE
Verifying User Equivalence of user=root from all DBs to all CELLS.
#### RoCE CONFIGURATION TESTS ####
Checking for presence of RoCE devices on all DBs and CELLS
[SUCCESS].... RoCE devices on all DBs and CELLS look good
Checking for RoCE Policy Routing settings on all DBs and CELLS
[SUCCESS].... RoCE Policy Routing settings look good
Checking for RoCE DSCP ToS mapping on all DBs and CELLS
[SUCCESS].... RoCE DSCP ToS settings look good
Checking for RoCE PFC settings and DSCP mapping on all DBs and
CELLs
[SUCCESS].... RoCE PFC and DSCP settings look good
Checking for RoCE interface MTU settings. Expected value : 2300
[SUCCESS].... RoCE interface MTU settings look good
Verifying switch advertised DSCP on all DBs and CELLS ports ( )
[SUCCESS].... Advertised DSCP settings from RoCE switch looks good
#### CONNECTIVITY TESTS ####
[COMPUTE NODES -> STORAGE CELLS]
(60 seconds approx.)
(Will walk through QoS values: 0-6) [SUCCESS].....Results OK
[SUCCESS]..... All can talk to all storage cells
[COMPUTE NODES -> COMPUTE NODES]
...

```

2. Install the spine switches (R1SS and R2SS).

- a. Physically install and power up the spine switches in the existing rack (R1SS) and the new rack (R2SS).
 - i. Physically install each spine switch in RU1.
 - ii. For each spine switch, ensure that the management Ethernet interface is connected to the management network and then supply power.
 - iii. On each spine switch, perform the initial configuration steps outlined in Configuring the Cisco Nexus C9336C-FX2 Switch. Skip the step for applying the golden configuration settings as you will do this later.
 - iv. For each spine switch, perform a ping test to the management Ethernet interface to ensure that the switch is online and accessible.
- b. Apply the golden configuration settings to the new spine switches.

See Applying Golden Configuration Settings on RoCE Network Fabric Switches in *Oracle Exadata Database Machine Maintenance Guide*.

You can use the instance of `patchmgr` that you previously used to update the firmware on the leaf switches (in step 1.h).

Use a switch list file (`spines.lst`) to apply the golden configuration settings to both spine switches using one `patchmgr` command:

```

# cat spines.lst
R1SS_IP:mspine.201
R2SS_IP:mspine.202

```

```
# ./patchmgr --roceswitches spines.lst --apply-config -log_dir /tmp/spinelogs
```

 Note:

In the switch list file, *R1SS_IP* is the management IP address or host name for the spine switch on the existing rack (R1SS) and *R2SS_IP* is the management IP address or host name for the spine switch on the new rack (R2SS).

- c. Upgrade the firmware on the spine switches.

See Updating RoCE Network Fabric Switch Firmware in *Oracle Exadata Database Machine Maintenance Guide*.

You can use the instance of `patchmgr` that you used in the previous step.

Use a switch list file (`spines.lst`) to perform the firmware upgrade on both spine switches using one `patchmgr` command:

```
# cat spines.lst
R1SS_IP:mspine.201
R2SS_IP:mspine.202

# ./patchmgr --roceswitches spines.lst --upgrade -log_dir /tmp/spinelogs
```

 Note:

In the switch list file, *R1SS_IP* is the management IP address or host name for the spine switch on the existing rack (R1SS) and *R2SS_IP* is the management IP address or host name for the spine switch on the new rack (R2SS).

- d. Connect the RoCE Network Fabric cables to the spine switches (R1SS and R2SS).

 WARNING:

At this stage, only connect the cables to the spine switches.

To avoid later complications, ensure that each cable connects to the correct switch and port.

DO NOT CONNECT ANY OF THE CABLES TO THE LEAF SWITCHES.

Use the cables that you prepared earlier (in step 1.c).

For the required cross-rack cabling information, see [Two-Rack Cabling for RA21 and Later Model Racks](#).

3. Perform the first round of configuration on the lower leaf switches (R1LL and R2LL).

Perform this step on the lower leaf switches (R1LL and R2LL) only.

 Note:

During this step, the lower leaf switch ports are shut down. While the R1LL ports are down, R1UL exclusively supports the RoCE Network Fabric. During this time, there is no redundancy in the RoCE Network Fabric, and availability cannot be maintained if R1UL goes down.

- a. Shut down the switch ports on the lower leaf switches (R1LL and R2LL).

- i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/1-36
R1LL(config-if-range)# shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

- ii. Repeat the command sequence on R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/1-36
R2LL(config-if-range)# shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

- b. Reconfigure the lower leaf switch ports (R1LL and R2LL) .

For each switch, you must use the correct corresponding switch configuration file, which you earlier copied to the switch (in step 1.g).

- i. On R1LL, the switch configuration file name must end with `step3_R1_LL.cfg`:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# run-script bootflash:///
roce_multi_14uplinks_online_step3_R1_LL.cfg | grep 'none'
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```


- ii. On R2LL, the switch configuration file name must end with `step3_R2_LL.cfg`:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# run-script bootflash:///
roce_multi_14uplinks_online_step3_R2_LL.cfg | grep 'none'
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

 Note:

This step can take approximately 5 to 8 minutes on each switch.

- c. Start the inter-switch ports on the lower leaf switches (R1LL and R2LL) .

- i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/1-7, ethernet 1/30-36
R1LL(config-if-range)# no shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

- ii. Repeat the command sequence on R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/1-7, ethernet 1/30-36
R2LL(config-if-range)# no shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

- d. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- e. Verify the status of the inter-switch ports on the lower leaf switches (R1LL and R2LL) .
Run the `show interface status` command on each lower leaf switch:

i. R1LL# **show interface status**

ii. R2LL# **show interface status**

Examine the output to ensure that the inter-switch ports are connected.

For example:

R1LL# **show interface status**

```

-----
-----
Port              Name              Status      Vlan      Duplex  Speed  Type
-----
mgmt0             --              connected   routed    full    1000   --
-----
-----
Port              Name              Status      Vlan      Duplex  Speed  Type
-----
Eth1/1            --              xcvrAbsen  1          auto    auto   --
Eth1/2            --              xcvrAbsen  1          auto    auto   --
Eth1/3            --              xcvrAbsen  1          auto    auto   --
Eth1/4            ISL1            connected   trunk      full    100G
QSFP-100G-CR4
Eth1/5            ISL2            connected   trunk      full    100G
QSFP-100G-CR4
Eth1/6            ISL3            connected   trunk      full    100G
QSFP-100G-CR4
Eth1/7            ISL4            connected   trunk      full    100G
QSFP-100G-CR4
Eth1/8            celadm14        disabled    3888       full    100G
QSFP-100G-CR4
Eth1/9            celadm13        disabled    3888       full    100G
QSFP-100G-CR4
Eth1/10           celadm12        disabled    3888       full    100G
QSFP-100G-CR4
Eth1/11           celadm11        disabled    3888       full    100G
QSFP-100G-CR4
Eth1/12           celadm10        disabled    3888       full    100G
QSFP-100G-CR4
Eth1/13           celadm09        disabled    3888       full    100G
QSFP-100G-CR4
Eth1/14           celadm08        disabled    3888       full    100G
QSFP-100G-CR4
Eth1/15           adm08           disabled    3888       full    100G
QSFP-100G-CR4
Eth1/16           adm07           disabled    3888       full    100G
QSFP-100G-CR4
Eth1/17           adm06           disabled    3888       full    100G
QSFP-100G-CR4
Eth1/18           adm05           disabled    3888       full    100G
QSFP-100G-CR4

```

Eth1/19	adm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

- f. Start the storage server ports on the lower leaf switches (R1LL and R2LL) .

i. On R1LL:

```

R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R1LL(config-if-range)# no shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#

```

ii. Repeat the command sequence on R2LL:

```

R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R2LL(config-if-range)# no shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#

```

- g. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- h. Verify the status of the storage server ports on the lower leaf switches (R1LL and R2LL).

Run the `show interface status` command on each lower leaf switch:

i. R1LL# `show interface status`

ii. R2LL# `show interface status`

Examine the output to ensure that the storage server ports are connected.

For example:

R1LL# `show interface status`

```

-----
-----
Port          Name          Status    Vlan      Duplex  Speed  Type
-----
-----
mgmt0         --            connected routed     full    1000   --
-----
-----
Port          Name          Status    Vlan      Duplex  Speed  Type
-----
-----
Eth1/1        --            xcvrAbsen 1      auto    auto   --
Eth1/2        --            xcvrAbsen 1      auto    auto   --
Eth1/3        --            xcvrAbsen 1      auto    auto   --
Eth1/4        ISL1          connected trunk     full    100G
QSFP-100G-CR4
Eth1/5        ISL2          connected trunk     full    100G
QSFP-100G-CR4
Eth1/6        ISL3          connected trunk     full    100G
QSFP-100G-CR4
Eth1/7        ISL4          connected trunk     full    100G
QSFP-100G-CR4
Eth1/8        celadm14      connected 3888     full    100G
QSFP-100G-CR4

```

Eth1/9	celadm13	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/10	celadm12	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/11	celadm11	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/12	celadm10	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/13	celadm09	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/14	celadm08	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/15	adm08	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/16	adm07	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/17	adm06	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/18	adm05	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/19	adm04	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/20	adm03	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/21	adm02	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/22	adm01	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/23	celadm07	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/24	celadm06	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/25	celadm05	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/26	celadm04	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/27	celadm03	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/28	celadm02	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/29	celadm01	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/30	ISL5	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/31	ISL6	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/32	ISL7	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/33	ISL8	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/34	--	xcvrAbsen 1	auto	auto	--
Eth1/35	--	xcvrAbsen 1	auto	auto	--
Eth1/36	--	xcvrAbsen 1	auto	auto	--
Po100	--	connected trunk	full	100G	--
Lo0	Routing loopback i	connected routed	auto	auto	--
Lo1	VTEP loopback inte	connected routed	auto	auto	--

```
Vlan1      --          down      routed      auto      auto      --
nve1       --          connected --          auto      auto      --
```

- i. Start the database server ports on the lower leaf switches (R1LL and R2LL).

- i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/15-22
R1LL(config-if-range)# no shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

- ii. Repeat the command sequence on R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/15-22
R2LL(config-if-range)# no shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

- j. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
 - k. Verify the status of the database server ports on the lower leaf switches (R1LL and R2LL).

Run the `show interface status` command on each lower leaf switch:

- i. R1LL# `show interface status`

- ii. R2LL# `show interface status`

Examine the output to ensure that the database server ports are `connected`.

For example:

```
R1LL# show interface status

-----
Port          Name          Status      Vlan      Duplex Speed  Type
-----
mgmt0         --            connected   routed    full   1000   --
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Eth1/1	--	xcvrAbsen	1	auto	auto	--
Eth1/2	--	xcvrAbsen	1	auto	auto	--
Eth1/3	--	xcvrAbsen	1	auto	auto	--
Eth1/4	ISL1	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/5	ISL2	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/6	ISL3	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/7	ISL4	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/8	celadm14	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/9	celadm13	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/10	celadm12	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/11	celadm11	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/12	celadm10	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/13	celadm09	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/14	celadm08	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/15	adm08	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/16	adm07	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/17	adm06	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/18	adm05	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/19	adm04	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	connected	3888	full	100G	

QSFP-100G-CR4						
Eth1/28	celadm02	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

 Note:

Before proceeding, ensure that you have completed all of the actions in step 3 on both lower leaf switches (R1LL and R2LL). If not, then ensure that you go back and perform the missing actions.

4. Perform the first round of configuration on the upper leaf switches (R1UL and R2UL).
Perform this step on the upper leaf switches (R1UL and R2UL) only.

 Note:

At the start of this step, the upper leaf switch ports are shut down. While the R1UL ports are down, R1LL exclusively supports the RoCE Network Fabric on the existing rack. During this time, there is no redundancy in the RoCE Network Fabric, and availability cannot be maintained if R1LL goes down.

- a. Shut down the upper leaf switch ports (R1UL and R2UL).
 - i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/1-36
R1UL(config-if-range)# shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
```



```
R1UL(config)# <Ctrl-Z>
R1UL#
```

- ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/1-36
R2UL(config-if-range)# shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- b. On both racks, remove the inter-switch links between the leaf switches (R1LL to R1UL, and R2LL to R2UL).

On every leaf switch, remove the cables for the inter-switch links:

- i. On R1LL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
 - ii. On R1UL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
 - iii. On R2LL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
 - iv. On R2UL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
- c. On both racks, cable the upper leaf switch to both of the spine switches (R1UL and R2UL to R1SS and R2SS).

Connect the cables from the spine switches that you prepared earlier (in step 2.d).

Cable the switches as described in [Two-Rack Cabling for RA21 and Later Model Racks](#):

- i. On R1UL, cable ports 01, 02, 03, 04, 05, 06, 07, 30, 31, 32, 33, 34, 35, and 36 to R1SS and R2SS.
- ii. On R2UL, cable ports 01, 02, 03, 04, 05, 06, 07, 30, 31, 32, 33, 34, 35, and 36 to R1SS and R2SS.

Note:

Ensure that each cable connects to the correct switch and port at both ends. In addition to physically checking each connection, you can run the `show lldp neighbors` command on each network switch and examine the output to confirm correct connections. You can individually check each cable connection to catch and correct errors quickly.

- d. Reconfigure the upper leaf switch ports (R1UL and R2UL).

For each switch, you must use the correct corresponding switch configuration file, which you earlier copied to the switch (in step 1.g):

- i. On R1UL, the switch configuration file name must end with `step4_R1_UL.cfg`:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# run-script bootflash:///
roce_multi_14uplinks_online_step4_R1_UL.cfg | grep 'none'
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

- ii. On R2UL, the switch configuration file name must end with `step4_R2_UL.cfg`:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# run-script bootflash:///
roce_multi_14uplinks_online_step4_R2_UL.cfg | grep 'none'
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

Note:

This step can take approximately 5 to 8 minutes on each switch.

- e. Check the status of the RoCE Network Fabric ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

- i. R1UL# `show interface status`

- ii. R2UL# `show interface status`

Examine the output to ensure that all of the cabled ports are disabled.

For example:

```
R1UL# show interface status
```

```
-----
-----
Port          Name                Status    Vlan    Duplex Speed
Type
-----
-----
```

```

mgmt0      --      connected routed    full 1000    --
-----
-----
Port      Name      Status  Vlan      Duplex Speed
Type
-----
-----
Eth1/1      RouterPort1    disabled routed    full 100G
QSFP-100G-CR4
Eth1/2      RouterPort2    disabled routed    full 100G
QSFP-100G-CR4
Eth1/3      RouterPort3    disabled routed    full 100G
QSFP-100G-CR4
Eth1/4      RouterPort4    disabled routed    full 100G
QSFP-100G-CR4
Eth1/5      RouterPort5    disabled routed    full 100G
QSFP-100G-CR4
Eth1/6      RouterPort6    disabled routed    full 100G
QSFP-100G-CR4
Eth1/7      RouterPort7    disabled routed    full 100G
QSFP-100G-CR4
Eth1/8      celadm14      disabled 3888      full 100G
QSFP-100G-CR4
Eth1/9      celadm13      disabled 3888      full 100G
QSFP-100G-CR4
Eth1/10     celadm12      disabled 3888      full 100G
QSFP-100G-CR4
Eth1/11     celadm11      disabled 3888      full 100G
QSFP-100G-CR4
Eth1/12     celadm10      disabled 3888      full 100G
QSFP-100G-CR4
Eth1/13     celadm09      disabled 3888      full 100G
QSFP-100G-CR4
Eth1/14     celadm08      disabled 3888      full 100G
QSFP-100G-CR4
Eth1/15     adm08         disabled 3888      full 100G
QSFP-100G-CR4
Eth1/16     adm07         disabled 3888      full 100G
QSFP-100G-CR4
Eth1/17     adm06         disabled 3888      full 100G
QSFP-100G-CR4
Eth1/18     adm05         disabled 3888      full 100G
QSFP-100G-CR4
Eth1/19     adm04         disabled 3888      full 100G
QSFP-100G-CR4
Eth1/20     adm03         disabled 3888      full 100G
QSFP-100G-CR4
Eth1/21     adm02         disabled 3888      full 100G
QSFP-100G-CR4
Eth1/22     adm01         disabled 3888      full 100G
QSFP-100G-CR4
Eth1/23     celadm07      disabled 3888      full 100G
QSFP-100G-CR4
Eth1/24     celadm06      disabled 3888      full 100G
QSFP-100G-CR4

```

Eth1/25	celadm05	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	RouterPort8	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/31	RouterPort9	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/32	RouterPort10	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/33	RouterPort11	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/34	RouterPort12	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/35	RouterPort13	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/36	RouterPort14	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

 Note:

Before proceeding, ensure that you have completed all of the actions to this point in step 4 on both upper leaf switches (R1UL and R2UL). If not, then ensure that you go back and perform the missing actions.

- f. Verify the configuration of the upper leaf switches.

You can use the instance of `patchmgr` that you previously used to update the switch firmware (in step 1.h).

Use a switch list file (`ul.lst`) to check both upper leaf switches using one `patchmgr` command:

```
# cat ul.lst
R1UL_IP:mleaf_u14.102
R2UL_IP:mleaf_u14.104
```

On a system with Secure Fabric enabled, use the `msfleaf_u14` tag in the switch list file:

```
# cat ul.lst
R1UL_IP:msfleaf_u14.102
R2UL_IP:msfleaf_u14.104
```

The following shows the recommended command and an example of the expected results:

```
# ./patchmgr --roceswitches ul.lst --verify-config -log_dir /tmp/log
2020-08-10 13:40:09 -0700      :Working: Initiating config
verification... Expect up to 6 minutes for each switch
Mon Aug 10 13:40:13 PDT 2020 1 of 4 :Verifying config on switch ...
...
Mon Aug 10 13:40:32 PDT 2020:      [INFO      ] Config matches
template: ...
Mon Aug 10 13:40:32 PDT 2020:      [SUCCESS   ] Config validation
successful!
2020-08-10 13:40:32 -0700
Config check on RoCE switch(es)
2020-08-10 13:40:32 -0700
Completed run of command: ./patchmgr --roceswitches ul.lst --verify-
config -log_dir /tmp/log
2020-08-10 13:40:32 -0700      :INFO      : config attempted on nodes in
file ul.lst: [R1UL_IP R2UL_IP]
2020-08-10 13:40:32 -0700      :INFO      : For details, check the
following files in /tmp/log:
2020-08-10 13:40:32 -0700      :INFO      : - updateRoceSwitch.log
2020-08-10 13:40:32 -0700      :INFO      : - updateRoceSwitch.trc
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.stdout
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.stderr
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.log
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.trc
2020-08-10 13:40:32 -0700      :INFO      : Exit status:0
2020-08-10 13:40:32 -0700      :INFO      : Exiting.
```

In the command output, verify that the switch configuration is good for both upper leaf switches. You can ignore messages about the ports that are down.

5. Finalize the configuration of the lower leaf switches (R1LL and R2LL).

Perform this step on the lower leaf switches (R1LL and R2LL) only.

a. Reconfigure the lower leaf switch ports (R1LL and R2LL).

Run the following command sequence on both of the lower leaf switches (R1LL and R2LL).

You must use the correct switch configuration file, which you earlier copied to the switch (in step 1.g). In this step, the configuration file name must end with `step5.cfg`.

i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# run-script bootflash:///
roce_multi_14uplinks_online_step5.cfg | grep 'none'
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

ii. Repeat the command sequence on R2LL:

```

R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# run-script bootflash:///
roce_multi_14uplinks_online_step5.cfg | grep 'none'
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#

```

 Note:

This step can take approximately 5 to 8 minutes on each switch.

b. On both racks, cable the lower leaf switch to both of the spine switches (R1LL and R2LL to R1SS and R2SS).

Connect the cables from the spine switches that you prepared earlier (in step 2.d).

Cable the switches as described in [Two-Rack Cabling for RA21 and Later Model Racks](#):

- i. On R1LL, cable ports 01, 02, 03, 04, 05, 06, 07, 30, 31, 32, 33, 34, 35, and 36 to R1SS and R2SS.
- ii. On R2LL, cable ports 01, 02, 03, 04, 05, 06, 07, 30, 31, 32, 33, 34, 35, and 36 to R1SS and R2SS.

 Note:

Ensure that each cable connects to the correct switch and port at both ends. In addition to physically checking each connection, you can run the `show lldp neighbors` command on each network switch and examine the output to confirm correct connections. You can individually check each cable connection to catch and correct errors quickly.

c. On the lower leaf switches, verify that all of the cabled RoCE Network Fabric ports are connected (R1LL and R2LL).

Run the `show interface status` command on each lower leaf switch:

i. R1LL# **show interface status**

ii. R2LL# **show interface status**

Examine the output to ensure that all of the cabled ports are `connected`.

For example:

```
R1LL# show interface status
```

Port Type	Name	Status	Vlan	Duplex	Speed
mgmt0	--	connected	routed	full 1000	--
Port Type	Name	Status	Vlan	Duplex	Speed
Eth1/1	RouterPort1	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/2	RouterPort2	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/3	RouterPort3	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/4	RouterPort4	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/5	RouterPort5	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/6	RouterPort6	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/7	RouterPort7	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/8	celadm14	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/9	celadm13	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/10	celadm12	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/11	celadm11	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/12	celadm10	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/13	celadm09	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/14	celadm08	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/15	adm08	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/16	adm07	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/17	adm06	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/18	adm05	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/19	adm04	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/20	adm03	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/21	adm02	connected	3888	full	100G
QSFP-100G-CR4					

Eth1/22	adm01	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/23	celadm07	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/24	celadm06	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/25	celadm05	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/26	celadm04	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/27	celadm03	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/28	celadm02	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/29	celadm01	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/30	RouterPort8	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/31	RouterPort9	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/32	RouterPort10	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/33	RouterPort11	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/34	RouterPort12	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/35	RouterPort13	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/36	RouterPort14	connected routed	full 100G	
QSFP-100G-CR4				
Lo0	Routing loopback i	connected routed	auto auto	--
Lo1	VTEP loopback inte	connected routed	auto auto	--
Vlan1	--	down routed	auto auto	--
nve1	--	connected --	auto auto	--

 Note:

Before proceeding, ensure that you have completed all of the actions to this point in step 5 on both lower leaf switches (R1LL and R2LL). If not, then ensure that you go back and perform the missing actions.

- d. Verify the configuration of the lower leaf switches.

You can use the instance of `patchmgr` that you previously used to update the switch firmware (in step 1.h).

Use a switch list file (`ll.lst`) to check both lower leaf switches using one `patchmgr` command:

```
# cat ll.lst
R1LL_IP:mleaf_u14.101
R2LL_IP:mleaf_u14.103
```


On a system with Secure Fabric enabled, use the `msfleaf_u14` tag in the switch list file:

```
# cat ll.lst
R1LL_IP:msfleaf_u14.101
R2LL_IP:msfleaf_u14.103
```

The following shows the recommended command and an example of the expected results:

```
# ./patchmgr --roceswitches ll.lst --verify-config -log_dir /tmp/log
2020-08-10 13:45:09 -0700      :Working: Initiating config
verification... Expect up to 6 minutes for each switch
Mon Aug 10 13:45:13 PDT 2020 1 of 4 :Verifying config on switch ...
...
Mon Aug 10 13:45:32 PDT 2020:      [INFO      ] Config matches
template: ...
Mon Aug 10 13:45:32 PDT 2020:      [SUCCESS   ] Config validation
successful!
2020-08-10 13:45:32 -0700
  Config check on RoCE switch(es)
2020-08-10 13:45:32 -0700
  Completed run of command: ./patchmgr --roceswitches ll.lst --verify-
config -log_dir /tmp/log
2020-08-10 13:45:32 -0700      :INFO      : config attempted on nodes in
file ll.lst: [R1LL_IP R2LL_IP]
2020-08-10 13:45:32 -0700      :INFO      : For details, check the
following files in /tmp/log:
2020-08-10 13:45:32 -0700      :INFO      : - updateRoceSwitch.log
2020-08-10 13:45:32 -0700      :INFO      : - updateRoceSwitch.trc
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.stdout
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.stderr
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.log
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.trc
2020-08-10 13:45:32 -0700      :INFO      : Exit status:0
2020-08-10 13:45:32 -0700      :INFO      : Exiting.
```

In the command output, verify that the switch configuration is good for both lower leaf switches.

- e. Verify that nve is up on the lower leaf switches (R1LL and R2LL).

Run the following command on each lower leaf switch and examine the output:

- i. R1LL# **show nve peers**

- ii. R2LL# **show nve peers**

At this point, you should see one nve peer with State=Up.

For example:

```
R1LL# show nve peers
Interface      Peer-IP State LearnType  Uptime      Router-Mac
```

```
-----
nve1          100.64.1.103      Up          CP 00:04:29          n/a
```

- f. Verify that BGP is up on the lower leaf switches (R1LL and R2LL).

Run the following command on each lower leaf switch and examine the output:

- i. R1LL# **show logging log | grep BGP**

- ii. R2LL# **show logging log | grep BGP**

Look for two entries with `Up` in the rightmost column that are associated with different IP addresses.

For example:

```
R1LL# show logging log | grep BGP
2020 Aug 10 13:47:13 R1LL %BGP-5-ADJCHANGE: bgp- [29342] (default)
neighbor 100.64.0.201 Up
2020 Aug 10 13:47:24 R1LL %BGP-5-ADJCHANGE: bgp- [29342] (default)
neighbor 100.64.0.202 Up
```

6. Finalize the configuration of the upper leaf switches (R1UL and R2UL).

Perform this step on the upper leaf switches (R1UL and R2UL) only.

- a. Start the inter-switch ports on the upper leaf switches (R1UL and R2UL).

- i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/1-7, ethernet 1/30-36
R1UL(config-if-range)# no shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

- ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/1-7, ethernet 1/30-36
R2UL(config-if-range)# no shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- b. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- c. Verify the status of the inter-switch ports on the upper leaf switches (R1UL and R2UL).
Run the `show interface status` command on each upper leaf switch:

i. R1UL# **show interface status**

ii. R2UL# **show interface status**

Examine the output to ensure that the inter-switch ports are connected.

For example:

R1UL# **show interface status**

Port Type	Name	Status	Vlan	Duplex	Speed
mgmt0	--	connected	routed	full	1000
Eth1/1	RouterPort1	connected	routed	full	100G
Eth1/2	RouterPort2	connected	routed	full	100G
Eth1/3	RouterPort3	connected	routed	full	100G
Eth1/4	RouterPort4	connected	routed	full	100G
Eth1/5	RouterPort5	connected	routed	full	100G
Eth1/6	RouterPort6	connected	routed	full	100G
Eth1/7	RouterPort7	connected	routed	full	100G
Eth1/8	celadm14	disabled	3888	full	100G
Eth1/9	celadm13	disabled	3888	full	100G
Eth1/10	celadm12	disabled	3888	full	100G
Eth1/11	celadm11	disabled	3888	full	100G
Eth1/12	celadm10	disabled	3888	full	100G
Eth1/13	celadm09	disabled	3888	full	100G

Eth1/14	celadm08	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/15	adm08	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/16	adm07	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/17	adm06	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/18	adm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/19	adm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	RouterPort8	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/31	RouterPort9	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/32	RouterPort10	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/33	RouterPort11	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/34	RouterPort12	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/35	RouterPort13	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/36	RouterPort14	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

- d. Start the storage server ports on the upper leaf switches (R1UL and R2UL).

i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R1UL(config-if-range)# no shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R2UL(config-if-range)# no shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- e. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- f. Verify the status of the storage server ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

i. R1UL# `show interface status`

ii. R2UL# `show interface status`

Examine the output to ensure that the storage server ports are connected.

For example:

```
R1UL# show interface status
```

```
-----
-----
Port          Name          Status    Vlan      Duplex Speed
Type
-----
-----
mgmt0         --            connected routed    full 1000  --
-----
-----
```

Port Type	Name	Status	Vlan	Duplex	Speed

Eth1/1	RouterPort1	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/2	RouterPort2	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/3	RouterPort3	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/4	RouterPort4	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/5	RouterPort5	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/6	RouterPort6	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/7	RouterPort7	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/8	celadm14	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/9	celadm13	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/10	celadm12	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/11	celadm11	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/12	celadm10	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/13	celadm09	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/14	celadm08	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/15	adm08	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/16	adm07	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/17	adm06	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/18	adm05	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/19	adm04	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/20	adm03	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/21	adm02	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/22	adm01	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/23	celadm07	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/24	celadm06	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/25	celadm05	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/26	celadm04	<i>connected</i>	3888	full	100G
QSFP-100G-CR4					

Eth1/27	celadm03	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/28	celadm02	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/29	celadm01	connected 3888	full 100G	
QSFP-100G-CR4				
Eth1/30	RouterPort8	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/31	RouterPort9	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/32	RouterPort10	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/33	RouterPort11	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/34	RouterPort12	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/35	RouterPort13	connected routed	full 100G	
QSFP-100G-CR4				
Eth1/36	RouterPort14	connected routed	full 100G	
QSFP-100G-CR4				
Lo0	Routing loopback i	connected routed	auto auto	--
Lo1	VTEP loopback inte	connected routed	auto auto	--
Vlan1	--	down routed	auto auto	--
nve1	--	connected --	auto auto	--

g. Start the database server ports on the upper leaf switches (R1UL and R2UL).

i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/15-22
R1UL(config-if-range)# no shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/15-22
R2UL(config-if-range)# no shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- h. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- i. Verify the status of the database server ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

i. R1UL# `show interface status`

ii. R2UL# `show interface status`

Examine the output to ensure that the database server ports are connected.

For example:

R1UL# `show interface status`

```
-----
-----
Port          Name          Status    Vlan    Duplex Speed
Type
-----
-----
mgmt0         --           connected routed    full 1000  --
-----
-----
Port          Name          Status    Vlan    Duplex Speed
Type
-----
-----
Eth1/1        RouterPort1   connected routed    full 100G
QSFP-100G-CR4
Eth1/2        RouterPort2   connected routed    full 100G
QSFP-100G-CR4
Eth1/3        RouterPort3   connected routed    full 100G
QSFP-100G-CR4
Eth1/4        RouterPort4   connected routed    full 100G
QSFP-100G-CR4
Eth1/5        RouterPort5   connected routed    full 100G
QSFP-100G-CR4
Eth1/6        RouterPort6   connected routed    full 100G
QSFP-100G-CR4
Eth1/7        RouterPort7   connected routed    full 100G
QSFP-100G-CR4
Eth1/8        celadm14      connected 3888     full 100G
QSFP-100G-CR4
Eth1/9        celadm13      connected 3888     full 100G
QSFP-100G-CR4
Eth1/10       celadm12      connected 3888     full 100G
QSFP-100G-CR4
Eth1/11       celadm11      connected 3888     full 100G
QSFP-100G-CR4
Eth1/12       celadm10      connected 3888     full 100G
QSFP-100G-CR4
Eth1/13       celadm09      connected 3888     full 100G
```



```

QSFP-100G-CR4
Eth1/14      celadm08      connected 3888      full 100G
QSFP-100G-CR4
Eth1/15      adm08      connected 3888      full 100G
QSFP-100G-CR4
Eth1/16      adm07      connected 3888      full 100G
QSFP-100G-CR4
Eth1/17      adm06      connected 3888      full 100G
QSFP-100G-CR4
Eth1/18      adm05      connected 3888      full 100G
QSFP-100G-CR4
Eth1/19      adm04      connected 3888      full 100G
QSFP-100G-CR4
Eth1/20      adm03      connected 3888      full 100G
QSFP-100G-CR4
Eth1/21      adm02      connected 3888      full 100G
QSFP-100G-CR4
Eth1/22      adm01      connected 3888      full 100G
QSFP-100G-CR4
Eth1/23      celadm07      connected 3888      full 100G
QSFP-100G-CR4
Eth1/24      celadm06      connected 3888      full 100G
QSFP-100G-CR4
Eth1/25      celadm05      connected 3888      full 100G
QSFP-100G-CR4
Eth1/26      celadm04      connected 3888      full 100G
QSFP-100G-CR4
Eth1/27      celadm03      connected 3888      full 100G
QSFP-100G-CR4
Eth1/28      celadm02      connected 3888      full 100G
QSFP-100G-CR4
Eth1/29      celadm01      connected 3888      full 100G
QSFP-100G-CR4
Eth1/30      RouterPort8      connected routed      full 100G
QSFP-100G-CR4
Eth1/31      RouterPort9      connected routed      full 100G
QSFP-100G-CR4
Eth1/32      RouterPort10      connected routed      full 100G
QSFP-100G-CR4
Eth1/33      RouterPort11      connected routed      full 100G
QSFP-100G-CR4
Eth1/34      RouterPort12      connected routed      full 100G
QSFP-100G-CR4
Eth1/35      RouterPort13      connected routed      full 100G
QSFP-100G-CR4
Eth1/36      RouterPort14      connected routed      full 100G
QSFP-100G-CR4
Lo0          Routing loopback i connected routed      auto auto      --
Lo1          VTEP loopback inte connected routed      auto auto      --
Vlan1        --              down      routed      auto auto      --
nve1         --              connected --      auto auto      --

```

- j. Verify that nve is up on the leaf switches (R1LL, R1UL, R2LL, and R2UL).
Run the following command on each leaf switch and examine the output:

i. R1LL# **show nve peers**

ii. R1UL# **show nve peers**

iii. R2LL# **show nve peers**

iv. R2UL# **show nve peers**

In the output, you should see three nve peers with State=Up.

For example:

```
R1UL# show nve peers
Interface          Peer-IP State LearnType  Uptime      Router-Mac
-----
nve1               100.64.1.101 Up        CP 00:04:29 n/a
nve1               100.64.1.103 Up        CP 00:07:48 n/a
nve1               100.64.1.104 Up        CP 00:04:10 n/a
```

k. Verify that BGP is up on the upper leaf switches (R1UL and R2UL).

Run the following command on each upper leaf switch and examine the output:

i. R1UL# **show logging log | grep BGP**

ii. R2UL# **show logging log | grep BGP**

In the output, look for two entries with Up in the rightmost column that are associated with different IP addresses.

For example:

```
R1UL# show logging log | grep BGP
2020 Aug 10 13:57:13 R1UL %BGP-5-ADJCHANGE: bgp- [32782] (default)
neighbor 100.64.0.201 Up
2020 Aug 10 13:57:24 R1UL %BGP-5-ADJCHANGE: bgp- [32782] (default)
neighbor 100.64.0.202 Up
```

7. For each rack (R1 and R2), confirm the multi-rack cabling by running the `verify_roce_cables.py` script.

The `verify_roce_cables.py` script uses two input files; one for database servers and storage servers (`nodes.rackN`), and another for switches (`switches.rackN`). In each file, every server or switch must be listed on separate lines. Use fully qualified domain names or IP addresses for each server and switch.

See My Oracle Support document 2587717.1 for download and detailed usage instructions.

Run the `verify_roce_cables.py` script against both of the racks:

```
a. # cd /opt/oracle.SupportTools/ibdiagtools
   # ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
```

- b. `# cd /opt/oracle.SupportTools/ibdiagtools`
`# ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2`

Check the output of the `verify_roce_cables.py` script against the tables in [Two-Rack Cabling for RA21 and Later Model Racks](#). Also, check that output in the `CABLE OK?` columns contains the `OK` status.

The following examples show extracts of the expected command results:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
SWITCH PORT (EXPECTED PEER)  LOWER LEAF (rack1sw-rocea0) : CABLE OK?  UPPER LEAF (rack1sw-roceb0) : CABLE OK?
-----
----- : -----
...

```

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2
SWITCH PORT (EXPECTED PEER)  LOWER LEAF (rack2sw-rocea0) : CABLE OK?  UPPER LEAF (rack2sw-roceb0) : CABLE OK?
-----
----- : -----
...

```

8. Verify the RoCE Network Fabric operation across both interconnected racks by using the `infinicheck` command.

Use the following recommended command sequence to verify the RoCE Network Fabric operation across both racks.

In each command, `hosts.all` contains a list of database server RoCE Network Fabric IP addresses from both racks (2 RoCE Network Fabric IP addresses for each database server), and `cells.all` contains a list of RoCE Network Fabric IP addresses for the storage servers from both racks (2 RoCE Network Fabric IP addresses for each storage server).

- a. `# cd /opt/oracle.SupportTools/ibdiagtools`
`# ./infinicheck -g hosts.all -c cells.all -z`
- b. `# cd /opt/oracle.SupportTools/ibdiagtools`
`# ./infinicheck -g hosts.all -c cells.all -s`
- c. `# cd /opt/oracle.SupportTools/ibdiagtools`
`# ./infinicheck -g hosts.all -c cells.all -b`

See [step 1.k](#) for most information about each `infinicheck` command.

The following example shows the expected command results for the final command in the sequence:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./infinicheck -g hosts.all -c cells.all -b
```

```
INFINICHECK
```

[Network Connectivity, Configuration and Performance]

```

##### FABRIC TYPE TESTS #####
System type identified: RoCE
Verifying User Equivalence of user=root from all DBs to all CELLS.
##### RoCE CONFIGURATION TESTS #####
Checking for presence of RoCE devices on all DBs and CELLS
[SUCCESS].... RoCE devices on all DBs and CELLS look good
Checking for RoCE Policy Routing settings on all DBs and CELLS
[SUCCESS].... RoCE Policy Routing settings look good
Checking for RoCE DSCP ToS mapping on all DBs and CELLS
[SUCCESS].... RoCE DSCP ToS settings look good
Checking for RoCE PFC settings and DSCP mapping on all DBs and CELLS
[SUCCESS].... RoCE PFC and DSCP settings look good
Checking for RoCE interface MTU settings. Expected value : 2300
[SUCCESS].... RoCE interface MTU settings look good
Verifying switch advertised DSCP on all DBs and CELLS ports ( )
[SUCCESS].... Advertised DSCP settings from RoCE switch looks good
##### CONNECTIVITY TESTS #####
[COMPUTE NODES -> STORAGE CELLS]
(60 seconds approx.)
(Will walk through QoS values: 0-6) [SUCCESS].....Results OK
[SUCCESS]..... All can talk to all storage cells
[COMPUTE NODES -> COMPUTE NODES]
...

```

At this point, both racks share the RoCE Network Fabric, and the combined system is ready for further configuration.

Extending an X8M Rack with No Down Time by Adding an RA21 or Later Model Rack



WARNING:

Take time to read and understand this procedure before implementation. Pay careful attention to the instructions that surround the command examples. A system outage may occur if the procedure is not applied correctly.



Note:

This procedure assumes that the RoCE Network Fabric switches on the existing X8M rack contain the golden configuration settings from Oracle Exadata System Software 20.1.0 or later. Otherwise, before using this procedure, you must update the Oracle Exadata System Software and update the golden configuration settings on the RoCE Network Fabric switches. Downtime is required to update the golden configuration settings on the RoCE Network Fabric switches.

**Note:**

For additional background information, see [Understanding Multi-Rack Cabling for X8M Racks](#) and [Understanding Multi-Rack Cabling for RA21 and Later Model Racks](#).

Use this procedure to extend a typical X8M rack without down-time by cabling it together with an RA21 or later model rack. The primary rack (designated R1) and all of the systems it supports remain online throughout the procedure. At the beginning of the procedure, the additional rack (designated R2) is shut down.

The following is an outline of the procedure:

- Preparation (steps [1](#) and [2](#))
In this phase, you prepare the racks, switches, and cables. Also, you install and cable the spine switches in both racks.
- Configuration and Physical Cabling
In this phase, you reconfigure the leaf switches and finalize the cabling to the spine switches. These tasks are carefully orchestrated to avoid downtime on the primary system, as follows:
 - Partially configure the lower leaf switches (step [3](#))
In this step, you reconfigure the switch ports on the lower leaf switches. There is no physical cabling performed in this step.
 - Partially configure the upper leaf switches (step [4](#))
In this step, you reconfigure the switch ports on the upper leaf switches, remove the inter-switch cables that connect the leaf switches in both racks and connect the cables between the upper leaf switches and the spine switches.
 - Finalize the lower leaf switches (step [5](#))
In this step, you finalize the switch port configuration on the lower leaf switches. You also complete the physical cabling by connecting the cables between the lower leaf switches and the spine switches.
 - Finalize the upper leaf switches (step [6](#))
In this step, you finalize the switch port configuration on the upper leaf switches.
- Validation and Testing (steps [7](#) and [8](#))
In this phase, you validate and test the RoCE Network Fabric across both of the interconnect racks.

After completing the procedure, both racks share the RoCE Network Fabric, and the combined system is ready for further configuration. For example, you can extend existing disk groups and Oracle RAC databases to consume resources across both racks.

 Note:

- This procedure applies only to typical rack configurations that initially have leaf switches with the following specifications:
 - The inter-switch ports are ports 4 to 7, and ports 30 to 33.
 - The storage server ports are ports 8 to 14, and ports 23 to 29.
 - The database server ports are ports 15 to 22.

For other rack configurations (for example, 8-socket systems with three database servers and 11 storage servers) a different procedure and different RoCE Network Fabric switch configuration files are required. Contact Oracle for further guidance.

- The procedure uses the following naming abbreviations and conventions:
 - The abbreviation for the existing X8M rack is R1, and the new RA21 or later model rack is R2.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies a spine switch.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the existing rack (R1).
- Most operations must be performed in multiple locations. For example, step [1.h](#) instructs you to update the firmware on all the RoCE Network Fabric leaf switches (R1LL, R1UL, R2LL, and R2UL). Pay attention to the instructions and keep track of your actions.

 Tip:

When a step must be performed on multiple switches, the instruction contains a list of the applicable switches. For example, (R1LL, R1UL, R2LL, and R2UL). You can use this list as a checklist to keep track of your actions.

- Perform operations sequentially, and complete every operation before proceeding. For example, run the entire command sequence at [3.a.i](#) as one operation and complete it before proceeding.
- All of commands that are run on a RoCE Network Fabric switch must be run while connected to the switch management interface as the switch administrator.

1. Prepare the systems.

- a. Position the new RA21 or later model rack (R2) so that it is physically near the existing X8M rack (R1).

The RDMA Network Fabric cables must be able to reach the switches in each rack.

For the required cross-rack cabling information, see [Two-Rack Cabling for a System Combining an X8M Rack and a Later Model Rack](#).

- b. Power on all of the servers and network switches in the new rack (R2).

This includes the database servers, storage servers, RoCE Network Fabric leaf switches, and the Management Network Switch.

- c. Prepare the RoCE Network Fabric cables that you will use to interconnect the racks.
Label both ends of every cable.

For the required cross-rack cabling information, see [Two-Rack Cabling for a System Combining an X8M Rack and a Later Model Rack](#).

- d. Connect the new rack (R2) to your existing management network.

Ensure that there are no IP address conflicts across the racks and that you can access the management interfaces on the RoCE Network Fabric switches.

- e. Ensure that you have a backup of the current switch configuration for each RoCE Network Fabric switch (R1LL, R1UL, R2LL, and R2UL).

See *Backing Up Settings on the RoCE Network Fabric Switch* in *Oracle Exadata Database Machine Maintenance Guide*.

- f. Download the required RoCE Network Fabric switch configuration files.

This procedure requires specific RoCE Network Fabric switch configuration files, which you must download from [My Oracle Support document 2704997.1](#).

 WARNING:

You must use different switch configuration files depending on whether your system uses Exadata Secure RDMA Fabric Isolation. Ensure that you download the correct archive that matches your system configuration.

For system configurations without Secure Fabric, download `online_multi-rack_8and14uplinks.zip`. For system configurations with Secure Fabric, download `online_SF_enabled_multi-rack_8and14uplinks.zip`.

Download and extract the archive containing the required RoCE Network Fabric switch configuration files. Place the files on a server with access to the management interfaces on the RoCE Network Fabric switches.

- g. Copy the required RoCE Network Fabric switch configuration files to the leaf switches on both racks.

You can use the following commands to copy the required configuration files to all of the RoCE Network Fabric switches on a system without Secure Fabric enabled:

- i. `# scp roce_multi_online_step3_R1_LL.cfg admin@R1LL_IP:/`
- ii. `# scp roce_multi_14uplinks_online_step3_R2_LL.cfg admin@R2LL_IP:/`
- iii. `# scp roce_multi_online_step4_R1_UL.cfg admin@R1UL_IP:/`
- iv. `# scp roce_multi_14uplinks_online_step4_R2_UL.cfg admin@R2UL_IP:/`
- v. `# scp roce_multi_online_step5.cfg admin@R1LL_IP:/`

```
vi. # scp roce_multi_14uplinks_online_step5.cfg admin@R2LL_IP:/
```

On a system with Secure Fabric enabled, you can use the following commands:

```
i. # scp roce_SF_multi_online_step3_R1_LL.cfg admin@R1LL_IP:/
```

```
ii. # scp roce_SF_multi_14uplinks_online_step3_R2_LL.cfg admin@R2LL_IP:/
```

```
iii. # scp roce_SF_multi_online_step4_R1_UL.cfg admin@R1UL_IP:/
```

```
iv. # scp roce_SF_multi_14uplinks_online_step4_R2_UL.cfg admin@R2UL_IP:/
```

```
v. # scp roce_SF_multi_online_step5.cfg admin@R1LL_IP:/
```

```
vi. # scp roce_SF_multi_14uplinks_online_step5.cfg admin@R2LL_IP:/
```

In the above commands, substitute the appropriate IP address or host name where applicable. For example, in place of *R1LL_IP*, substitute the management IP address or host name for the lower leaf switch (LL) on the existing rack (R1).

Note:

The command examples in the rest of this procedure use the configuration files for a system configuration without Secure Fabric enabled. If required, adjust the commands to use the Secure Fabric-enabled switch configuration files.

- h. Update the firmware to the latest available release on all of the RoCE Network Fabric leaf switches (R1LL, R1UL, R2LL, and R2UL).

See Updating RoCE Network Fabric Switch Firmware in *Oracle Exadata Database Machine Maintenance Guide*.

- i. Examine the RoCE Network Fabric leaf switches (R1LL, R1UL, R2LL, and R2UL) and confirm the port categories for the cabled ports.

Run the `show interface status` command on every RoCE Network Fabric leaf switch:

```
i. R1LL# show interface status
```

```
ii. R1UL# show interface status
```

```
iii. R2LL# show interface status
```

```
iv. R2UL# show interface status
```

Examine the output and confirm the port categories as follows:

- Confirm that the inter-switch ports are ports 4 to 7, and ports 30 to 33.
- Confirm that the storage server ports are ports 8 to 14, and ports 23 to 29.

- Confirm that the database server ports are ports 15 to 22.

For example:

```
R1LL# show interface status
```

```
-----  
-----  
Port          Name          Status    Vlan      Duplex  Speed  Type  
-----  
-----  
mgmt0         --           connected routed     full    1000   --  
-----  
-----
```

```
Port          Name          Status    Vlan      Duplex  Speed  Type  
-----  
-----  
Eth1/1        --           xcvrAbsen 1      auto    auto   --  
Eth1/2        --           xcvrAbsen 1      auto    auto   --  
Eth1/3        --           xcvrAbsen 1      auto    auto   --  
Eth1/4        ISL1         connected trunk    full    100G  
QSFP-100G-CR4  
Eth1/5        ISL2         connected trunk    full    100G  
QSFP-100G-CR4  
Eth1/6        ISL3         connected trunk    full    100G  
QSFP-100G-CR4  
Eth1/7        ISL4         connected trunk    full    100G  
QSFP-100G-CR4  
Eth1/8        celadm14     connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/9        celadm13     connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/10       celadm12     connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/11       celadm11     connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/12       celadm10     connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/13       celadm09     connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/14       celadm08     connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/15       adm08        connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/16       adm07        connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/17       adm06        connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/18       adm05        connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/19       adm04        connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/20       adm03        connected 3888     full    100G  
QSFP-100G-CR4  
Eth1/21       adm02        connected 3888     full    100G
```

QSFP-100G-CR4						
Eth1/22	adm01	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

- j. For each rack (R1 and R2), confirm the RoCE Network Fabric cabling by running the `verify_roce_cables.py` script.

The `verify_roce_cables.py` script uses two input files; one for database servers and storage servers (`nodes.rackN`), and another for switches (`switches.rackN`). In each file, every server or switch must be listed on separate lines. Use fully qualified domain names or IP addresses for each server and switch.

See My Oracle Support document 2587717.1 for download and detailed usage instructions.

Run the `verify_roce_cables.py` script against both of the racks:

- i.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
```
- ii.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2
```

Check that output in the `CABLE OK?` columns contains the `OK` status.

The following example shows the expected command results:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
SWITCH PORT (EXPECTED PEER) LOWER LEAF (rack1sw-rocea0) :
CABLE OK? UPPER LEAF (rack1sw-roceb0) : CABLE OK?
-----
Eth1/4 (ISL peer switch) : rack1sw-rocea0 Ethernet1/4 :
OK rack1sw-roceb0 Ethernet1/4 : OK
Eth1/5 (ISL peer switch) : rack1sw-rocea0 Ethernet1/5 :
OK rack1sw-roceb0 Ethernet1/5 : OK
Eth1/6 (ISL peer switch) : rack1sw-rocea0 Ethernet1/6 :
OK rack1sw-roceb0 Ethernet1/6 : OK
Eth1/7 (ISL peer switch) : rack1sw-rocea0 Ethernet1/7 :
OK rack1sw-roceb0 Ethernet1/7 : OK
Eth1/8 (RU39) : rack1celadm14 port-1 :
OK rack1celadm14 port-2 : OK
Eth1/9 (RU37) : rack1celadm13 port-1 :
OK rack1celadm13 port-2 : OK
Eth1/10 (RU35) : rack1celadm12 port-1 :
OK rack1celadm12 port-2 : OK
Eth1/11 (RU33) : rack1celadm11 port-1 :
OK rack1celadm11 port-2 : OK
Eth1/12 (RU31) : rack1celadm10 port-1 :
OK rack1celadm10 port-2 : OK
Eth1/13 (RU29) : rack1celadm09 port-1 :
OK rack1celadm09 port-2 : OK
Eth1/14 (RU27) : rack1celadm08 port-1 :
OK rack1celadm08 port-2 : OK
Eth1/15 (RU26) : rack1adm08 port-1 :
OK rack1adm08 port-2 : OK
Eth1/16 (RU25) : rack1adm07 port-1 :
OK rack1adm07 port-2 : OK
Eth1/17 (RU24) : rack1adm06 port-1 :
OK rack1adm06 port-2 : OK
Eth1/18 (RU23) : rack1adm05 port-1 :
OK rack1adm05 port-2 : OK
Eth1/19 (RU19) : rack1adm04 port-1 :
OK rack1adm04 port-2 : OK
Eth1/20 (RU18) : rack1adm03 port-1 :
OK rack1adm03 port-2 : OK
Eth1/21 (RU17) : rack1adm02 port-1 :
OK rack1adm02 port-2 : OK
Eth1/22 (RU16) : rack1adm01 port-1 :
OK rack1adm01 port-2 : OK
Eth1/23 (RU14) : rack1celadm07 port-1 :
OK rack1celadm07 port-2 : OK
Eth1/24 (RU12) : rack1celadm06 port-1 :
OK rack1celadm06 port-2 : OK
Eth1/25 (RU10) : rack1celadm05 port-1 :
OK rack1celadm05 port-2 : OK
Eth1/26 (RU08) : rack1celadm04 port-1 :
OK rack1celadm04 port-2 : OK
Eth1/27 (RU06) : rack1celadm03 port-1 :
OK rack1celadm03 port-2 : OK
```

```

Eth1/28 (RU04)          :          rack1celadm02 port-1 :
OK                      rack1celadm02 port-2 : OK
Eth1/29 (RU02)          :          rack1celadm01 port-1 :
OK                      rack1celadm01 port-2 : OK
Eth1/30 (ISL peer switch) :          rack1sw-rocea0 Ethernet1/30 :
OK                      rack1sw-roceb0 Ethernet1/30 : OK
Eth1/31 (ISL peer switch) :          rack1sw-rocea0 Ethernet1/31 :
OK                      rack1sw-roceb0 Ethernet1/31 : OK
Eth1/32 (ISL peer switch) :          rack1sw-rocea0 Ethernet1/32 :
OK                      rack1sw-roceb0 Ethernet1/32 : OK
Eth1/33 (ISL peer switch) :          rack1sw-rocea0 Ethernet1/33 :
OK                      rack1sw-roceb0 Ethernet1/33 : OK

```

- k. For each rack (R1 and R2), verify the RoCE Network Fabric operation by using the `infinicheck` command.
- Use `infinicheck` with the `-z` option to clear the files that were created during the last run of the `infinicheck` command.
 - Use `infinicheck` with the `-s` option to set up user equivalence for password-less SSH across the RoCE Network Fabric.
 - Finally, verify the RoCE Network Fabric operation by using `infinicheck` with the `-b` option, which is recommended on newly imaged machines where it is acceptable to suppress the `cellip.ora` and `cellinit.ora` configuration checks.

In each command, the hosts input file (`hosts.rack1` and `hosts.rack2`) contains a list of database server RoCE Network Fabric IP addresses (2 RoCE Network Fabric IP addresses for each database server), and the cells input file (`cells.rack1` and `cells.rack2`) contains a list of RoCE Network Fabric IP addresses for the storage servers (2 RoCE Network Fabric IP addresses for each storage server).

- i. Use the following recommended command sequence on the existing rack (R1):

```

i. # cd /opt/oracle.SupportTools/ibdiagtools
   # ./infinicheck -g hosts.rack1 -c cells.rack1 -z

ii. # cd /opt/oracle.SupportTools/ibdiagtools
    # ./infinicheck -g hosts.rack1 -c cells.rack1 -s

iii. # cd /opt/oracle.SupportTools/ibdiagtools
      # ./infinicheck -g hosts.rack1 -c cells.rack1 -b

```

- ii. Use the following recommended command sequence on the new rack (R2):

```

i. # cd /opt/oracle.SupportTools/ibdiagtools
   # ./infinicheck -g hosts.rack2 -c cells.rack2 -z

ii. # cd /opt/oracle.SupportTools/ibdiagtools
    # ./infinicheck -g hosts.rack2 -c cells.rack2 -s

iii. # cd /opt/oracle.SupportTools/ibdiagtools
      # ./infinicheck -g hosts.rack2 -c cells.rack2 -b

```

The following example shows the expected command results for the final command in the sequence:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./infinicheck -g hosts.rackN -c cells.rackN -b

INFINICHECK
    [Network Connectivity, Configuration and Performance]

        ##### FABRIC TYPE TESTS #####
System type identified: RoCE
Verifying User Equivalence of user=root from all DBs to all CELLS.
        ##### RoCE CONFIGURATION TESTS #####
            Checking for presence of RoCE devices on all DBs and CELLS
[SUCCESS].... RoCE devices on all DBs and CELLS look good
            Checking for RoCE Policy Routing settings on all DBs and CELLS
[SUCCESS].... RoCE Policy Routing settings look good
            Checking for RoCE DSCP ToS mapping on all DBs and CELLS
[SUCCESS].... RoCE DSCP ToS settings look good
            Checking for RoCE PFC settings and DSCP mapping on all DBs and
CELLs
[SUCCESS].... RoCE PFC and DSCP settings look good
            Checking for RoCE interface MTU settings. Expected value : 2300
[SUCCESS].... RoCE interface MTU settings look good
            Verifying switch advertised DSCP on all DBs and CELLS ports ( )
[SUCCESS].... Advertised DSCP settings from RoCE switch looks good
        ##### CONNECTIVITY TESTS #####
            [COMPUTE NODES -> STORAGE CELLS]
                (60 seconds approx.)
            (Will walk through QoS values: 0-6) [SUCCESS].....Results OK
[SUCCESS]..... All can talk to all storage cells
            [COMPUTE NODES -> COMPUTE NODES]
...

```

2. Install the spine switches (R1SS and R2SS).

- a. Physically install and power up the spine switches in the existing rack (R1SS) and the new rack (R2SS).
 - i. Physically install each spine switch in RU1.
 - ii. For each spine switch, ensure that the management Ethernet interface is connected to the management network and then supply power.
 - iii. On each spine switch, perform the initial configuration steps outlined in Configuring the Cisco Nexus C9336C-FX2 Switch. Skip the step for applying the golden configuration settings as you will do this later.
 - iv. For each spine switch, perform a ping test to the management Ethernet interface to ensure that the switch is online and accessible.
- b. Apply the golden configuration settings to the new spine switches.

See Applying Golden Configuration Settings on RoCE Network Fabric Switches in *Oracle Exadata Database Machine Maintenance Guide*.

You can use the instance of `patchmgr` that you previously used to update the firmware on the leaf switches (in step 1.h).

Use a switch list file (`spines.lst`) to apply the golden configuration settings to both spine switches using one `patchmgr` command:

```
# cat spines.lst
R1SS_IP:mspine.201
R2SS_IP:mspine.202

# ./patchmgr --roceswitches spines.lst --apply-config -log_dir /tmp/
spinelogs
```

 Note:

In the switch list file, `R1SS_IP` is the management IP address or host name for the spine switch on the existing rack (R1SS) and `R2SS_IP` is the management IP address or host name for the spine switch on the new rack (R2SS).

- c. Upgrade the firmware on the spine switches.

See Updating RoCE Network Fabric Switch Firmware in *Oracle Exadata Database Machine Maintenance Guide*.

You can use the instance of `patchmgr` that you used in the previous step.

Use a switch list file (`spines.lst`) to perform the firmware upgrade on both spine switches using one `patchmgr` command:

```
# cat spines.lst
R1SS_IP:mspine.201
R2SS_IP:mspine.202

# ./patchmgr --roceswitches spines.lst --upgrade -log_dir /tmp/spinelogs
```

 Note:

In the switch list file, `R1SS_IP` is the management IP address or host name for the spine switch on the existing rack (R1SS) and `R2SS_IP` is the management IP address or host name for the spine switch on the new rack (R2SS).

- d. Connect the RoCE Network Fabric cables to the spine switches (R1SS and R2SS).

 WARNING:

At this stage, only connect the cables to the spine switches.

To avoid later complications, ensure that each cable connects to the correct switch and port.

DO NOT CONNECT ANY OF THE CABLES TO THE LEAF SWITCHES.

Use the cables that you prepared earlier (in step 1.c).

For the required cross-rack cabling information, see [Two-Rack Cabling for a System Combining an X8M Rack and a Later Model Rack](#).

3. Perform the first round of configuration on the lower leaf switches (R1LL and R2LL).

Perform this step on the lower leaf switches (R1LL and R2LL) only.

 Note:

During this step, the lower leaf switch ports are shut down. While the R1LL ports are down, R1UL exclusively supports the RoCE Network Fabric. During this time, there is no redundancy in the RoCE Network Fabric, and availability cannot be maintained if R1UL goes down.

- a. Shut down the switch ports on the lower leaf switches (R1LL and R2LL).

- i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/4-33
R1LL(config-if-range)# shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

- ii. On R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/1-36
R2LL(config-if-range)# shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

- b. Reconfigure the lower leaf switch ports (R1LL and R2LL) .

For each switch, you must use the correct corresponding switch configuration file, which you earlier copied to the switch (in step 1.g).

- i. On R1LL, the switch configuration file name must end with `step3_R1_LL.cfg`:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# run-script bootflash:///
```

```

roce_multi_online_step3_R1_LL.cfg | grep 'none'
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#

```

- ii. On R2LL, the switch configuration file name must end with `step3_R2_LL.cfg`:

```

R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# run-script bootflash:///
roce_multi_14uplinks_online_step3_R2_LL.cfg | grep 'none'
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#

```

 Note:

This step can take approximately 5 to 8 minutes on each switch.

- c. Start the inter-switch ports on the lower leaf switches (R1LL and R2LL) .

- i. On R1LL:

```

R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/4-7, ethernet 1/30-33
R1LL(config-if-range)# no shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#

```

- ii. On R2LL:

```

R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/1-7, ethernet 1/30-36
R2LL(config-if-range)# no shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete

```



```
R2LL(config)# <Ctrl-Z>
R2LL#
```

- d. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- e. Verify the status of the inter-switch ports on the lower leaf switches (R1LL and R2LL) .
Run the `show interface status` command on each lower leaf switch:

i. R1LL# **show interface status**

ii. R2LL# **show interface status**

Examine the output to ensure that the inter-switch ports are connected.

For example:

```
R1LL# show interface status
```

```
-----
-----
Port          Name          Status    Vlan    Duplex Speed  Type
-----
-----
mgmt0         --            connected routed   full   1000   --
-----
-----
Port          Name          Status    Vlan    Duplex Speed  Type
-----
-----
Eth1/1        --            xcvrAbsen 1      auto   auto   --
Eth1/2        --            xcvrAbsen 1      auto   auto   --
Eth1/3        --            xcvrAbsen 1      auto   auto   --
Eth1/4        ISL1          connected trunk   full   100G
QSFP-100G-CR4
Eth1/5        ISL2          connected trunk   full   100G
QSFP-100G-CR4
Eth1/6        ISL3          connected trunk   full   100G
QSFP-100G-CR4
Eth1/7        ISL4          connected trunk   full   100G
QSFP-100G-CR4
Eth1/8        celadm14      disabled  3888    full   100G
QSFP-100G-CR4
Eth1/9        celadm13      disabled  3888    full   100G
QSFP-100G-CR4
Eth1/10       celadm12      disabled  3888    full   100G
QSFP-100G-CR4
Eth1/11       celadm11      disabled  3888    full   100G
QSFP-100G-CR4
Eth1/12       celadm10      disabled  3888    full   100G
QSFP-100G-CR4
Eth1/13       celadm09      disabled  3888    full   100G
QSFP-100G-CR4
Eth1/14       celadm08      disabled  3888    full   100G
QSFP-100G-CR4
```

Eth1/15	adm08	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/16	adm07	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/17	adm06	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/18	adm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/19	adm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	<i>connected</i>	trunk	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	<i>connected</i>	trunk	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	<i>connected</i>	trunk	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	<i>connected</i>	trunk	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

f. Start the storage server ports on the lower leaf switches (R1LL and R2LL) .

i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R1LL(config-if-range)# no shut
R1LL(config-if-range)# exit
```

```
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

- ii. Repeat the command sequence on R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R2LL(config-if-range)# no shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

- g. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- h. Verify the status of the storage server ports on the lower leaf switches (R1LL and R2LL).

Run the `show interface status` command on each lower leaf switch:

- i. R1LL# `show interface status`

- ii. R2LL# `show interface status`

Examine the output to ensure that the storage server ports are connected.

For example:

```
R1LL# show interface status
```

```
-----
-----
Port          Name          Status    Vlan    Duplex Speed  Type
-----
-----
mgmt0         --            connected routed   full   1000   --
-----
-----
Port          Name          Status    Vlan    Duplex Speed  Type
-----
-----
Eth1/1        --            xcvrAbsen 1      auto   auto   --
Eth1/2        --            xcvrAbsen 1      auto   auto   --
Eth1/3        --            xcvrAbsen 1      auto   auto   --
Eth1/4        ISL1          connected trunk   full   100G
QSFP-100G-CR4
Eth1/5        ISL2          connected trunk   full   100G
```

QSFP-100G-CR4					
Eth1/6	ISL3	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/7	ISL4	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/8	celadm14	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/9	celadm13	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/10	celadm12	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/11	celadm11	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/12	celadm10	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/13	celadm09	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/14	celadm08	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/15	adm08	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/16	adm07	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/17	adm06	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/18	adm05	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/19	adm04	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/20	adm03	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/21	adm02	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/22	adm01	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/23	celadm07	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/24	celadm06	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/25	celadm05	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/26	celadm04	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/27	celadm03	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/28	celadm02	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/29	celadm01	<i>connected</i> 3888	full	100G	
QSFP-100G-CR4					
Eth1/30	ISL5	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/31	ISL6	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/32	ISL7	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/33	ISL8	connected trunk	full	100G	

```

QSFP-100G-CR4
Eth1/34      --          xcvrAbsen 1          auto  auto  --
Eth1/35      --          xcvrAbsen 1          auto  auto  --
Eth1/36      --          xcvrAbsen 1          auto  auto  --
Po100        --          connected trunk      full  100G  --
Lo0          Routing loopback i connected routed  auto  auto  --
Lo1          VTEP loopback inte connected routed  auto  auto  --
Vlan1        --          down      routed      auto  auto  --
nve1         --          connected --          auto  auto  --

```

- i. Start the database server ports on the lower leaf switches (R1LL and R2LL).

- i. On R1LL:

```

R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/15-22
R1LL(config-if-range)# no shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#

```

- ii. Repeat the command sequence on R2LL:

```

R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/15-22
R2LL(config-if-range)# no shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#

```

- j. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- k. Verify the status of the database server ports on the lower leaf switches (R1LL and R2LL).

Run the `show interface status` command on each lower leaf switch:

- i. R1LL# `show interface status`

- ii. R2LL# `show interface status`

Examine the output to ensure that the database server ports are `connected`.

For example:

R1LL# **show interface status**

```

-----
-----
Port          Name                Status    Vlan      Duplex  Speed  Type
-----
mgmt0         --                  connected routed    full    1000   --
-----
-----
Port          Name                Status    Vlan      Duplex  Speed  Type
-----
Eth1/1        --                  xcvrAbsen 1      auto    auto   --
Eth1/2        --                  xcvrAbsen 1      auto    auto   --
Eth1/3        --                  xcvrAbsen 1      auto    auto   --
Eth1/4        ISL1                connected trunk    full    100G
QSFP-100G-CR4
Eth1/5        ISL2                connected trunk    full    100G
QSFP-100G-CR4
Eth1/6        ISL3                connected trunk    full    100G
QSFP-100G-CR4
Eth1/7        ISL4                connected trunk    full    100G
QSFP-100G-CR4
Eth1/8        celadm14            connected 3888      full    100G
QSFP-100G-CR4
Eth1/9        celadm13            connected 3888      full    100G
QSFP-100G-CR4
Eth1/10       celadm12            connected 3888      full    100G
QSFP-100G-CR4
Eth1/11       celadm11            connected 3888      full    100G
QSFP-100G-CR4
Eth1/12       celadm10            connected 3888      full    100G
QSFP-100G-CR4
Eth1/13       celadm09            connected 3888      full    100G
QSFP-100G-CR4
Eth1/14       celadm08            connected 3888      full    100G
QSFP-100G-CR4
Eth1/15       adm08               connected 3888      full    100G
QSFP-100G-CR4
Eth1/16       adm07               connected 3888      full    100G
QSFP-100G-CR4
Eth1/17       adm06               connected 3888      full    100G
QSFP-100G-CR4
Eth1/18       adm05               connected 3888      full    100G
QSFP-100G-CR4
Eth1/19       adm04               connected 3888      full    100G
QSFP-100G-CR4
Eth1/20       adm03               connected 3888      full    100G
QSFP-100G-CR4
Eth1/21       adm02               connected 3888      full    100G
QSFP-100G-CR4
Eth1/22       adm01               connected 3888      full    100G

```

QSFP-100G-CR4						
Eth1/23	celadm07	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

 Note:

Before proceeding, ensure that you have completed all of the actions in step 3 on both lower leaf switches (R1LL and R2LL). If not, then ensure that you go back and perform the missing actions.

4. Perform the first round of configuration on the upper leaf switches (R1UL and R2UL).
Perform this step on the upper leaf switches (R1UL and R2UL) only.

 Note:

At the start of this step, the upper leaf switch ports are shut down. While the R1UL ports are down, R1LL exclusively supports the RoCE Network Fabric on the existing rack. During this time, there is no redundancy in the RoCE Network Fabric, and availability cannot be maintained if R1LL goes down.

- a. Shut down the upper leaf switch ports (R1UL and R2UL).

i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/4-33
R1UL(config-if-range)# shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

ii. On R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/1-36
R2UL(config-if-range)# shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

b. On both racks, remove the inter-switch links between the leaf switches (R1LL to R1UL, and R2LL to R2UL).

On every leaf switch, remove the cables for the inter-switch links:

- i. On R1LL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
- ii. On R1UL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
- iii. On R2LL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
- iv. On R2UL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.

c. On both racks, cable the upper leaf switch to both of the spine switches (R1UL and R2UL to R1SS and R2SS).

Connect the cables from the spine switches that you prepared earlier (in step 2.d).

Cable the switches as described in [Two-Rack Cabling for a System Combining an X8M Rack and a Later Model Rack](#):

- i. On R1UL, cable ports 04, 05, 06, 07, 30, 31, 32, and 33 to R1SS and R2SS.
- ii. On R2UL, cable ports 01, 02, 03, 04, 05, 06, 07, 30, 31, 32, 33, 34, 35, and 36 to R1SS and R2SS.

 Note:

Ensure that each cable connects to the correct switch and port at both ends. In addition to physically checking each connection, you can run the `show lldp neighbors` command on each network switch and examine the output to confirm correct connections. You can individually check each cable connection to catch and correct errors quickly.

d. Reconfigure the upper leaf switch ports (R1UL and R2UL).

For each switch, you must use the correct corresponding switch configuration file, which you earlier copied to the switch (in step 1.g):

i. On R1UL, the switch configuration file name must end with `step4_R1_UL.cfg`:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# run-script bootflash:///
roce_multi_online_step4_R1_UL.cfg | grep 'none'
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

ii. On R2UL, the switch configuration file name must end with `step4_R2_UL.cfg`:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# run-script bootflash:///
roce_multi_14uplinks_online_step4_R2_UL.cfg | grep 'none'
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

 Note:

This step can take approximately 5 to 8 minutes on each switch.

e. Check the status of the RoCE Network Fabric ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

i. R1UL# `show interface status`

ii. R2UL# `show interface status`

Examine the output to ensure that all of the cabled ports are disabled.

The following example shows the expected output on the RA21 or later model rack (R2UL). On the X8M rack (R1UL), ports 01, 02, 03, 34, 35, and 36 are not physically connected.

R2UL# **show interface status**

Port Type	Name	Status	Vlan	Duplex	Speed
mgmt0	--	connected	routed	full	1000 --
Port Type	Name	Status	Vlan	Duplex	Speed
Eth1/1 QSFP-100G-CR4	RouterPort1	disabled	routed	full	100G
Eth1/2 QSFP-100G-CR4	RouterPort2	disabled	routed	full	100G
Eth1/3 QSFP-100G-CR4	RouterPort3	disabled	routed	full	100G
Eth1/4 QSFP-100G-CR4	RouterPort4	disabled	routed	full	100G
Eth1/5 QSFP-100G-CR4	RouterPort5	disabled	routed	full	100G
Eth1/6 QSFP-100G-CR4	RouterPort6	disabled	routed	full	100G
Eth1/7 QSFP-100G-CR4	RouterPort7	disabled	routed	full	100G
Eth1/8 QSFP-100G-CR4	celadm14	disabled	3888	full	100G
Eth1/9 QSFP-100G-CR4	celadm13	disabled	3888	full	100G
Eth1/10 QSFP-100G-CR4	celadm12	disabled	3888	full	100G
Eth1/11 QSFP-100G-CR4	celadm11	disabled	3888	full	100G
Eth1/12 QSFP-100G-CR4	celadm10	disabled	3888	full	100G
Eth1/13 QSFP-100G-CR4	celadm09	disabled	3888	full	100G
Eth1/14 QSFP-100G-CR4	celadm08	disabled	3888	full	100G
Eth1/15 QSFP-100G-CR4	adm08	disabled	3888	full	100G
Eth1/16 QSFP-100G-CR4	adm07	disabled	3888	full	100G
Eth1/17 QSFP-100G-CR4	adm06	disabled	3888	full	100G

Eth1/18	adm05	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/19	adm04	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	<i>disabled</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	RouterPort8	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/31	RouterPort9	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/32	RouterPort10	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/33	RouterPort11	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/34	RouterPort12	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/35	RouterPort13	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/36	RouterPort14	<i>disabled</i>	routed	full	100G	
QSFP-100G-CR4						
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

 Note:

Before proceeding, ensure that you have completed all of the actions to this point in step 4 on both upper leaf switches (R1UL and R2UL). If not, then ensure that you go back and perform the missing actions.

- f. Verify the configuration of the upper leaf switches.

You can use the instance of `patchmgr` that you previously used to update the switch firmware (in step 1.h).

Use a switch list file (`ul.lst`) to check both upper leaf switches using one `patchmgr` command:

```
# cat ul.lst
R1UL_IP:mleaf.102
R2UL_IP:mleaf_u14.104
```

On a system with Secure Fabric enabled, use the `msfleaf` and `msfleaf_u14` tags in the switch list file:

```
# cat ul.lst
R1UL_IP:msfleaf.102
R2UL_IP:msfleaf_u14.104
```

The following shows the recommended command and an example of the expected results:

```
# ./patchmgr --roceswitches ul.lst --verify-config -log_dir /tmp/log
2020-08-10 13:40:09 -0700      :Working: Initiating config
verification... Expect up to 6 minutes for each switch
Mon Aug 10 13:40:13 PDT 2020 1 of 4 :Verifying config on switch ...
...
Mon Aug 10 13:40:32 PDT 2020:      [INFO      ] Config matches
template: ...
Mon Aug 10 13:40:32 PDT 2020:      [SUCCESS   ] Config validation
successful!
2020-08-10 13:40:32 -0700
  Config check on RoCE switch(es)
2020-08-10 13:40:32 -0700
  Completed run of command: ./patchmgr --roceswitches ul.lst --verify-
config -log_dir /tmp/log
2020-08-10 13:40:32 -0700      :INFO      : config attempted on nodes in
file ul.lst: [R1UL_IP R2UL_IP]
2020-08-10 13:40:32 -0700      :INFO      : For details, check the
following files in /tmp/log:
2020-08-10 13:40:32 -0700      :INFO      : - updateRoceSwitch.log
2020-08-10 13:40:32 -0700      :INFO      : - updateRoceSwitch.trc
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.stdout
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.stderr
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.log
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.trc
2020-08-10 13:40:32 -0700      :INFO      : Exit status:0
2020-08-10 13:40:32 -0700      :INFO      : Exiting.
```

In the command output, verify that the switch configuration is good for both upper leaf switches. You can ignore messages about the ports that are down.

5. Finalize the configuration of the lower leaf switches (R1LL and R2LL).

Perform this step on the lower leaf switches (R1LL and R2LL) only.

a. Reconfigure the lower leaf switch ports (R1LL and R2LL).

Run the following command sequence on both of the lower leaf switches (R1LL and R2LL).

You must use the correct switch configuration file, which you earlier copied to the switch (in step 1.g). In this step, the configuration file name must end with `step5.cfg`.

i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# run-script bootflash:///roce_multi_online_step5.cfg |
grep 'none'
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

ii. On R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# run-script bootflash:///
roce_multi_14uplinks_online_step5.cfg | grep 'none'
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

 Note:

This step can take approximately 5 to 8 minutes on each switch.

- b. On both racks, cable the lower leaf switch to both of the spine switches (R1LL and R2LL to R1SS and R2SS).

Connect the cables from the spine switches that you prepared earlier (in step 2.d).

Cable the switches as described in [Two-Rack Cabling for a System Combining an X8M Rack and a Later Model Rack](#):

- i. On R1LL, cable ports 04, 05, 06, 07, 30, 31, 32, and 33 to R1SS and R2SS.
- ii. On R2LL, cable ports 01, 02, 03, 04, 05, 06, 07, 30, 31, 32, 33, 34, 35, and 36 to R1SS and R2SS.

 Note:

Ensure that each cable connects to the correct switch and port at both ends. In addition to physically checking each connection, you can run the `show lldp neighbors` command on each network switch and examine the output to confirm correct connections. You can individually check each cable connection to catch and correct errors quickly.

- c. On the lower leaf switches, verify that all of the cabled RoCE Network Fabric ports are connected (R1LL and R2LL).

Run the `show interface status` command on each lower leaf switch:

i. R1LL# `show interface status`

ii. R2LL# `show interface status`

Examine the output to ensure that all of the cabled ports are connected.

The following example shows the expected output on the RA21 or later model rack (R2LL). On the X8M rack (R1LL), ports 01, 02, 03, 34, 35, and 36 are not physically connected.

R2LL# `show interface status`

```

-----
-----
Port          Name          Status   Vlan    Duplex Speed
Type
-----
-----
mgmt0         --              connected routed    full 1000  --
-----
-----
Port          Name          Status   Vlan    Duplex Speed
Type
-----
-----
Eth1/1        RouterPort1    connected routed    full 100G
QSFP-100G-CR4
Eth1/2        RouterPort2    connected routed    full 100G
QSFP-100G-CR4
Eth1/3        RouterPort3    connected routed    full 100G
QSFP-100G-CR4
Eth1/4        RouterPort4    connected routed    full 100G
QSFP-100G-CR4
Eth1/5        RouterPort5    connected routed    full 100G
QSFP-100G-CR4
Eth1/6        RouterPort6    connected routed    full 100G
QSFP-100G-CR4
Eth1/7        RouterPort7    connected routed    full 100G
QSFP-100G-CR4
Eth1/8        celadm14       connected 3888     full 100G
QSFP-100G-CR4
Eth1/9        celadm13       connected 3888     full 100G
QSFP-100G-CR4
Eth1/10       celadm12       connected 3888     full 100G
QSFP-100G-CR4
Eth1/11       celadm11       connected 3888     full 100G
QSFP-100G-CR4
Eth1/12       celadm10       connected 3888     full 100G
QSFP-100G-CR4
Eth1/13       celadm09       connected 3888     full 100G

```

QSFP-100G-CR4					
Eth1/14	celadm08	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/15	adm08	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/16	adm07	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/17	adm06	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/18	adm05	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/19	adm04	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/20	adm03	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/21	adm02	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/22	adm01	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/23	celadm07	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/24	celadm06	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/25	celadm05	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/26	celadm04	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/27	celadm03	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/28	celadm02	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/29	celadm01	connected	3888	full	100G
QSFP-100G-CR4					
Eth1/30	RouterPort8	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/31	RouterPort9	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/32	RouterPort10	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/33	RouterPort11	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/34	RouterPort12	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/35	RouterPort13	connected	routed	full	100G
QSFP-100G-CR4					
Eth1/36	RouterPort14	connected	routed	full	100G
QSFP-100G-CR4					
Lo0	Routing loopback i	connected	routed	auto	auto --
Lo1	VTEP loopback inte	connected	routed	auto	auto --
Vlan1	--	down	routed	auto	auto --
nve1	--	connected	--	auto	auto --

 Note:

Before proceeding, ensure that you have completed all of the actions to this point in step 5 on both lower leaf switches (R1LL and R2LL). If not, then ensure that you go back and perform the missing actions.

d. Verify the configuration of the lower leaf switches.

You can use the instance of `patchmgr` that you previously used to update the switch firmware (in step 1.h).

Use a switch list file (`ll.lst`) to check both lower leaf switches using one `patchmgr` command:

```
# cat ll.lst
R1LL_IP:mleaf.101
R2LL_IP:mleaf_u14.103
```

On a system with Secure Fabric enabled, use the `msfleaf` and `msfleaf_u14` tags in the switch list file:

```
# cat ll.lst
R1LL_IP:msfleaf.101
R2LL_IP:msfleaf_u14.103
```

The following shows the recommended command and an example of the expected results:

```
# ./patchmgr --roceswitches ll.lst --verify-config -log_dir /tmp/log
2020-08-10 13:45:09 -0700      :Working: Initiating config
verification... Expect up to 6 minutes for each switch
Mon Aug 10 13:45:13 PDT 2020 1 of 4 :Verifying config on switch ...
...
Mon Aug 10 13:45:32 PDT 2020:      [INFO      ] Config matches
template: ...
Mon Aug 10 13:45:32 PDT 2020:      [SUCCESS   ] Config validation
successful!
2020-08-10 13:45:32 -0700
  Config check on RoCE switch(es)
2020-08-10 13:45:32 -0700
  Completed run of command: ./patchmgr --roceswitches ll.lst --verify-
config -log_dir /tmp/log
2020-08-10 13:45:32 -0700      :INFO      : config attempted on nodes in
file ll.lst: [R1LL_IP R2LL_IP]
2020-08-10 13:45:32 -0700      :INFO      : For details, check the
following files in /tmp/log:
2020-08-10 13:45:32 -0700      :INFO      : - updateRoceSwitch.log
2020-08-10 13:45:32 -0700      :INFO      : - updateRoceSwitch.trc
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.stdout
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.stderr
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.log
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.trc
```



```
2020-08-10 13:45:32 -0700      :INFO    : Exit status:0
2020-08-10 13:45:32 -0700      :INFO    : Exiting.
```

In the command output, verify that the switch configuration is good for both lower leaf switches.

- e. Verify that nve is up on the lower leaf switches (R1LL and R2LL).

Run the following command on each lower leaf switch and examine the output:

- i. R1LL# **show nve peers**

- ii. R2LL# **show nve peers**

At this point, you should see one nve peer with State=Up.

For example:

```
R1LL# show nve peers
Interface          Peer-IP State LearnType  Uptime          Router-Mac
-----
nve1               100.64.1.103 Up        CP 00:04:29          n/a
```

- f. Verify that BGP is up on the lower leaf switches (R1LL and R2LL).

Run the following command on each lower leaf switch and examine the output:

- i. R1LL# **show logging log | grep BGP**

- ii. R2LL# **show logging log | grep BGP**

Look for two entries with Up in the rightmost column that are associated with different IP addresses.

For example:

```
R1LL# show logging log | grep BGP
2020 Aug 10 13:47:13 R1LL %BGP-5-ADJCHANGE: bgp- [29342] (default)
neighbor 100.64.0.201 Up
2020 Aug 10 13:47:24 R1LL %BGP-5-ADJCHANGE: bgp- [29342] (default)
neighbor 100.64.0.202 Up
```

6. Finalize the configuration of the upper leaf switches (R1UL and R2UL).

Perform this step on the upper leaf switches (R1UL and R2UL) only.

- a. Start the inter-switch ports on the upper leaf switches (R1UL and R2UL).

- i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/4-7, ethernet 1/30-33
R1UL(config-if-range)# no shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
```

```
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

ii. On R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/1-7, ethernet 1/30-36
R2UL(config-if-range) # no shut
R2UL(config-if-range) # exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- b. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.

- c. Verify the status of the inter-switch ports on the upper leaf switches (R1UL and R2UL).
Run the `show interface status` command on each upper leaf switch:

i. R1UL# `show interface status`

ii. R2UL# `show interface status`

Examine the output to ensure that the inter-switch ports are connected.

The following example shows the expected output on the RA21 or later model rack (R2UL). On the X8M rack (R1UL), ports 01, 02, 03, 34, 35, and 36 are not physically connected.

R2UL# `show interface status`

```
-----
-----
Port          Name          Status    Vlan      Duplex Speed
Type
-----
-----
mgmt0         --            connected routed    full 1000  --
-----
-----
Port          Name          Status    Vlan      Duplex Speed
Type
-----
-----
Eth1/1        RouterPort1   connected routed    full 100G
QSFP-100G-CR4
Eth1/2        RouterPort2   connected routed    full 100G
QSFP-100G-CR4
Eth1/3        RouterPort3   connected routed    full 100G
QSFP-100G-CR4
```

Eth1/4	RouterPort4	<i>connected</i>	routed	full	100G
QSFP-100G-CR4					
Eth1/5	RouterPort5	<i>connected</i>	routed	full	100G
QSFP-100G-CR4					
Eth1/6	RouterPort6	<i>connected</i>	routed	full	100G
QSFP-100G-CR4					
Eth1/7	RouterPort7	<i>connected</i>	routed	full	100G
QSFP-100G-CR4					
Eth1/8	celadm14	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/9	celadm13	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/10	celadm12	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/11	celadm11	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/12	celadm10	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/13	celadm09	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/14	celadm08	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/15	adm08	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/16	adm07	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/17	adm06	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/18	adm05	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/19	adm04	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/20	adm03	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/21	adm02	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/22	adm01	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/23	celadm07	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/24	celadm06	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/25	celadm05	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/26	celadm04	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/27	celadm03	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/28	celadm02	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/29	celadm01	disabled	3888	full	100G
QSFP-100G-CR4					
Eth1/30	RouterPort8	<i>connected</i>	routed	full	100G
QSFP-100G-CR4					
Eth1/31	RouterPort9	<i>connected</i>	routed	full	100G
QSFP-100G-CR4					

Eth1/32	RouterPort10	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/33	RouterPort11	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/34	RouterPort12	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/35	RouterPort13	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/36	RouterPort14	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Lo0	Routing loopback i	<i>connected</i>	routed	auto	auto	--
Lo1	VTEP loopback inte	<i>connected</i>	routed	auto	auto	--
Vlan1	--	<i>down</i>	routed	auto	auto	--
nve1	--	<i>connected</i>	--	auto	auto	--

- d. Start the storage server ports on the upper leaf switches (R1UL and R2UL).

- i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R1UL(config-if-range)# no shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

- ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R2UL(config-if-range)# no shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- e. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- f. Verify the status of the storage server ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

- i. R1UL# `show interface status`

- ii. R2UL# `show interface status`

Examine the output to ensure that the storage server ports are connected.

The following example shows the expected output on the RA21 or later model rack (R2UL). On the X8M rack (R1UL), ports 01, 02, 03, 34, 35, and 36 are not physically connected.

R2UL# **show interface status**

```

-----
Port          Name          Status  Vlan    Duplex Speed
Type
-----
mgmt0         --           connected routed   full 1000  --
-----

Port          Name          Status  Vlan    Duplex Speed
Type
-----
Eth1/1        RouterPort1    connected routed   full 100G
QSFP-100G-CR4
Eth1/2        RouterPort2    connected routed   full 100G
QSFP-100G-CR4
Eth1/3        RouterPort3    connected routed   full 100G
QSFP-100G-CR4
Eth1/4        RouterPort4    connected routed   full 100G
QSFP-100G-CR4
Eth1/5        RouterPort5    connected routed   full 100G
QSFP-100G-CR4
Eth1/6        RouterPort6    connected routed   full 100G
QSFP-100G-CR4
Eth1/7        RouterPort7    connected routed   full 100G
QSFP-100G-CR4
Eth1/8        celadm14       connected 3888    full 100G
QSFP-100G-CR4
Eth1/9        celadm13       connected 3888    full 100G
QSFP-100G-CR4
Eth1/10       celadm12       connected 3888    full 100G
QSFP-100G-CR4
Eth1/11       celadm11       connected 3888    full 100G
QSFP-100G-CR4
Eth1/12       celadm10       connected 3888    full 100G
QSFP-100G-CR4
Eth1/13       celadm09       connected 3888    full 100G
QSFP-100G-CR4
Eth1/14       celadm08       connected 3888    full 100G
QSFP-100G-CR4
Eth1/15       adm08          disabled 3888    full 100G
QSFP-100G-CR4
Eth1/16       adm07          disabled 3888    full 100G
QSFP-100G-CR4
Eth1/17       adm06          disabled 3888    full 100G
QSFP-100G-CR4

```

Eth1/18	adm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/19	adm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	RouterPort8	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/31	RouterPort9	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/32	RouterPort10	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/33	RouterPort11	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/34	RouterPort12	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/35	RouterPort13	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/36	RouterPort14	connected	routed	full	100G	
QSFP-100G-CR4						
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

- g. Start the database server ports on the upper leaf switches (R1UL and R2UL).

- i. On R1UL:

```

R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/15-22
R1UL(config-if-range)# no shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete

```

```
R1UL(config)# <Ctrl-Z>
R1UL#
```

- ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/15-22
R2UL(config-if-range)# no shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- h. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- i. Verify the status of the database server ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

- i. R1UL# `show interface status`

- ii. R2UL# `show interface status`

Examine the output to ensure that the database server ports are connected.

The following example shows the expected output on the RA21 or later model rack (R2UL). On the X8M rack (R1UL), ports 01, 02, 03, 34, 35, and 36 are not physically connected.

```
R2UL# show interface status
```

```
-----
-----
Port          Name          Status    Vlan      Duplex Speed
Type
-----
-----
mgmt0         --            connected routed    full 1000  --
-----
-----
Port          Name          Status    Vlan      Duplex Speed
Type
-----
-----
Eth1/1        RouterPort1    connected routed    full 100G
QSFP-100G-CR4
Eth1/2        RouterPort2    connected routed    full 100G
QSFP-100G-CR4
Eth1/3        RouterPort3    connected routed    full 100G
QSFP-100G-CR4
```

Eth1/4	RouterPort4	connected routed	full 100G
QSFP-100G-CR4			
Eth1/5	RouterPort5	connected routed	full 100G
QSFP-100G-CR4			
Eth1/6	RouterPort6	connected routed	full 100G
QSFP-100G-CR4			
Eth1/7	RouterPort7	connected routed	full 100G
QSFP-100G-CR4			
Eth1/8	celadm14	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/9	celadm13	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/10	celadm12	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/11	celadm11	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/12	celadm10	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/13	celadm09	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/14	celadm08	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/15	adm08	<i>connected</i> 3888	full 100G
QSFP-100G-CR4			
Eth1/16	adm07	<i>connected</i> 3888	full 100G
QSFP-100G-CR4			
Eth1/17	adm06	<i>connected</i> 3888	full 100G
QSFP-100G-CR4			
Eth1/18	adm05	<i>connected</i> 3888	full 100G
QSFP-100G-CR4			
Eth1/19	adm04	<i>connected</i> 3888	full 100G
QSFP-100G-CR4			
Eth1/20	adm03	<i>connected</i> 3888	full 100G
QSFP-100G-CR4			
Eth1/21	adm02	<i>connected</i> 3888	full 100G
QSFP-100G-CR4			
Eth1/22	adm01	<i>connected</i> 3888	full 100G
QSFP-100G-CR4			
Eth1/23	celadm07	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/24	celadm06	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/25	celadm05	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/26	celadm04	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/27	celadm03	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/28	celadm02	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/29	celadm01	connected 3888	full 100G
QSFP-100G-CR4			
Eth1/30	RouterPort8	connected routed	full 100G
QSFP-100G-CR4			
Eth1/31	RouterPort9	connected routed	full 100G
QSFP-100G-CR4			

Eth1/32	RouterPort10	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/33	RouterPort11	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/34	RouterPort12	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/35	RouterPort13	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/36	RouterPort14	connected	routed	full	100G	
QSFP-100G-CR4						
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

- j. Verify that nve is up on the leaf switches (R1LL, R1UL, R2LL, and R2UL).
Run the following command on each leaf switch and examine the output:

i. R1LL# **show nve peers**

ii. R1UL# **show nve peers**

iii. R2LL# **show nve peers**

iv. R2UL# **show nve peers**

In the output, you should see three nve peers with State=Up.

For example:

```
R1UL# show nve peers
Interface          Peer-IP State LearnType  Uptime      Router-Mac
-----
nve1               100.64.1.101 Up       CP 00:04:29 n/a
nve1               100.64.1.103 Up       CP 00:07:48 n/a
nve1               100.64.1.104 Up       CP 00:04:10 n/a
```

- k. Verify that BGP is up on the upper leaf switches (R1UL and R2UL).
Run the following command on each upper leaf switch and examine the output:

i. R1UL# **show logging log | grep BGP**

ii. R2UL# **show logging log | grep BGP**

In the output, look for two entries with Up in the rightmost column that are associated with different IP addresses.

For example:

```
R1UL# show logging log | grep BGP
2020 Aug 10 13:57:13 R1UL %BGP-5-ADJCHANGE: bgp- [32782] (default)
neighbor 100.64.0.201 Up
```

```
2020 Aug 10 13:57:24 R1UL %BGP-5-ADJCHANGE: bgp- [32782] (default)
neighbor 100.64.0.202 Up
```

7. For each rack (R1 and R2), confirm the multi-rack cabling by running the `verify_roce_cables.py` script.

The `verify_roce_cables.py` script uses two input files; one for database servers and storage servers (`nodes.rackN`), and another for switches (`switches.rackN`). In each file, every server or switch must be listed on separate lines. Use fully qualified domain names or IP addresses for each server and switch.

See My Oracle Support document 2587717.1 for download and detailed usage instructions.

Run the `verify_roce_cables.py` script against both of the racks:

- a.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
```
- b.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2
```

Check the output of the `verify_roce_cables.py` script against the tables in [Two-Rack Cabling for a System Combining an X8M Rack and a Later Model Rack](#). Also, check that output in the `CABLE OK?` columns contains the `OK` status.

The following examples show extracts of the expected command results:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
SWITCH PORT (EXPECTED PEER)  LOWER LEAF (rack1sw-rocea0) : CABLE OK?  UPPER LEAF (rack1sw-
roceb0) : CABLE OK?
-----
...

# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2
SWITCH PORT (EXPECTED PEER)  LOWER LEAF (rack2sw-rocea0) : CABLE OK?  UPPER LEAF (rack2sw-
roceb0) : CABLE OK?
-----
...

```

8. Verify the RoCE Network Fabric operation across both interconnected racks by using the `infinicheck` command.

Use the following recommended command sequence to verify the RoCE Network Fabric operation across both racks.

In each command, `hosts.all` contains a list of database server RoCE Network Fabric IP addresses from both racks (2 RoCE Network Fabric IP addresses for each database server), and `cells.all` contains a list of RoCE Network Fabric IP addresses for the storage servers from both racks (2 RoCE Network Fabric IP addresses for each storage server).

- a. `# cd /opt/oracle.SupportTools/ibdiagtools`
`# ./infinicheck -g hosts.all -c cells.all -z`
- b. `# cd /opt/oracle.SupportTools/ibdiagtools`
`# ./infinicheck -g hosts.all -c cells.all -s`
- c. `# cd /opt/oracle.SupportTools/ibdiagtools`
`# ./infinicheck -g hosts.all -c cells.all -b`

See step 1.k for most information about each `infinicheck` command.

The following example shows the expected command results for the final command in the sequence:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./infinicheck -g hosts.all -c cells.all -b

INFINICHECK
    [Network Connectivity, Configuration and Performance]

        ##### FABRIC TYPE TESTS #####
System type identified: RoCE
Verifying User Equivalence of user=root from all DBs to all CELLS.
        ##### RoCE CONFIGURATION TESTS #####
        Checking for presence of RoCE devices on all DBs and CELLS
[SUCCESS].... RoCE devices on all DBs and CELLS look good
        Checking for RoCE Policy Routing settings on all DBs and CELLS
[SUCCESS].... RoCE Policy Routing settings look good
        Checking for RoCE DSCP ToS mapping on all DBs and CELLS
[SUCCESS].... RoCE DSCP ToS settings look good
        Checking for RoCE PFC settings and DSCP mapping on all DBs and CELLS
[SUCCESS].... RoCE PFC and DSCP settings look good
        Checking for RoCE interface MTU settings. Expected value : 2300
[SUCCESS].... RoCE interface MTU settings look good
        Verifying switch advertised DSCP on all DBs and CELLS ports ( )
[SUCCESS].... Advertised DSCP settings from RoCE switch looks good
        ##### CONNECTIVITY TESTS #####
        [COMPUTE NODES -> STORAGE CELLS]
        (60 seconds approx.)
        (Will walk through QoS values: 0-6) [SUCCESS].....Results OK
[SUCCESS]..... All can talk to all storage cells
        [COMPUTE NODES -> COMPUTE NODES]
...

```

At this point, both racks share the RoCE Network Fabric, and the combined system is ready for further configuration.

Two-Rack Cabling for a System Combining an X8M Rack and a Later Model Rack

This section provides the cabling details to connect an X8M rack with an RA21 or later model rack, both of which use RoCE Network Fabric.

 Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first (X8M) rack is R1, and the second (RA21 or later) rack is R2.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X8M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X8M-8 or X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.

The following illustration shows the cable connections for the spine switches when cabling a two-rack hybrid system with an X8M rack and an RA21 or later model rack:

Rack 1 Spine (R1SS)																	
1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35
		R1 UL P5	R1 UL P7	R1 UL P4	R1 UL P6	R1 LL P5	R1 LL P7	R1 LL P4	R1 LL P6								
2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36
		R2 UL P1	R2 UL P2	R2 UL P3	R2 UL P5	R2 UL P5	R2 UL P6	R2 UL P7	R2 LL P1	R2 LL P2	R2 LL P3	R2 LL P4	R2 LL P5	R2 LL P6	R2 LL P7		

Rack 2 Spine (R2SS)																	
1	3	5	7	9	11	13	15	17	19	21	23	25	27	29	31	33	35
		R1 UL P31	R1 UL P33	R1 UL P30	R1 UL P32	R1 LL P31	R1 LL P33	R1 LL P30	R1 LL P32								
2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36
		R2 UL P30	R2 UL P31	R2 UL P32	R2 UL P33	R2 UL P34	R2 UL P35	R2 UL P36	R2 LL P30	R2 LL P31	R2 LL P32	R2 LL P33	R2 LL P34	R2 LL P35	R2 LL P36		

The following tables contain details for all of the RoCE Network Fabric cabling connections in a two-rack hybrid system with an X8M rack and a later model rack.

Table 21-1 Leaf Switch Connections for the X8M Rack (R1)

Leaf Switch	Connection	Cable Length
R1UL to R1SS	R1UL-P5 to R1SS-P5 R1UL-P7 to R1SS-P7 R1UL-P4 to R1SS-P9 R1UL-P6 to R1SS-P11	3 meters
R1UL to R2SS	R1UL-P31 to R2SS-P5 R1UL-P33 to R2SS-P7 R1UL-P30 to R2SS-P9 R1UL-P32 to R2SS-P11	5 meters
R1LL to R1SS	R1LL-P5 to R1SS-P13 R1LL-P7 to R1SS-P15 R1LL-P4 to R1SS-P17 R1LL-P6 to R1SS-P19	3 meters
R1LL to R2SS	R1LL-P31 to R2SS-P13 R1LL-P33 to R2SS-P15 R1LL-P30 to R2SS-P17 R1LL-P32 to R2SS-P19	5 meters

Table 21-2 Leaf Switch Connections for the RA21 or Later Model Rack (R2)

Switch Pair	Connections	Cable Length
R2UL to R1SS	R2UL-P1 to R1SS-P6 R2UL-P2 to R1SS-P8 R2UL-P3 to R1SS-P10 R2UL-P4 to R1SS-P12 R2UL-P5 to R1SS-P14 R2UL-P6 to R1SS-P16 R2UL-P7 to R1SS-P18	5 meters
R2UL to R2SS	R2UL-P30 to R2SS-P6 R2UL-P31 to R2SS-P8 R2UL-P32 to R2SS-P10 R2UL-P33 to R2SS-P12 R2UL-P34 to R2SS-P14 R2UL-P35 to R2SS-P16 R2UL-P36 to R2SS-P18	3 meters
R2LL to R1SS	R2LL-P1 to R1SS-P20 R2LL-P2 to R1SS-P22 R2LL-P3 to R1SS-P24 R2LL-P4 to R1SS-P26 R2LL-P5 to R1SS-P28 R2LL-P6 to R1SS-P30 R2LL-P7 to R1SS-P32	5 meters

Table 21-2 (Cont.) Leaf Switch Connections for the RA21 or Later Model Rack (R2)

Switch Pair	Connections	Cable Length
R2LL to R2SS	R2LL-P30 to R2SS-P20 R2LL-P31 to R2SS-P22 R2LL-P32 to R2SS-P24 R2LL-P33 to R2SS-P26 R2LL-P34 to R2SS-P28 R2LL-P35 to R2SS-P30 R2LL-P36 to R2SS-P32	3 meters

Extending an X8M Rack with No Down Time by Adding Another X8M Rack

WARNING:

Take time to read and understand this procedure before implementation. Pay careful attention to the instructions that surround the command examples. A system outage may occur if the procedure is not applied correctly.

Note:

This procedure assumes that the RoCE Network Fabric switches on the X8M racks contain the golden configuration settings from Oracle Exadata System Software 20.1.0 or later. Otherwise, before using this procedure, you must update the Oracle Exadata System Software and update the golden configuration settings on the RoCE Network Fabric switches. Downtime is required to update the golden configuration settings on the RoCE Network Fabric switches.

Note:

For additional background information, see [Understanding Multi-Rack Cabling for X8M Racks](#).

Use this procedure to extend a typical X8M rack without down-time by cabling it together with a second X8M rack. The primary rack (designated R1) and all of the systems it supports remain online throughout the procedure. At the beginning of the procedure, the additional rack (designated R2) is shut down.

The following is an outline of the procedure:

- Preparation (steps [1](#) and [2](#))
In this phase, you prepare the racks, switches, and cables. Also, you install and cable the spine switches in both racks.
- Configuration and Physical Cabling

In this phase, you reconfigure the leaf switches and finalize the cabling to the spine switches. These tasks are carefully orchestrated to avoid downtime on the primary system, as follows:

- Partially configure the lower leaf switches (step 3)
In this step, you reconfigure the switch ports on the lower leaf switches. There is no physical cabling performed in this step.
- Partially configure the upper leaf switches (step 4)
In this step, you reconfigure the switch ports on the upper leaf switches, remove the inter-switch cables that connect the leaf switches in both racks and connect the cables between the upper leaf switches and the spine switches.
- Finalize the lower leaf switches (step 5)
In this step, you finalize the switch port configuration on the lower leaf switches. You also complete the physical cabling by connecting the cables between the lower leaf switches and the spine switches.
- Finalize the upper leaf switches (step 6)
In this step, you finalize the switch port configuration on the upper leaf switches.
- Validation and Testing (steps 7 and 8)
In this phase, you validate and test the RoCE Network Fabric across both of the interconnect racks.

After completing the procedure, both racks share the RoCE Network Fabric, and the combined system is ready for further configuration. For example, you can extend existing disk groups and Oracle RAC databases to consume resources across both racks.

 Note:

- This procedure applies only to typical rack configurations that initially have leaf switches with the following specifications:

- The inter-switch ports are ports 4 to 7, and ports 30 to 33.
- The storage server ports are ports 8 to 14, and ports 23 to 29.
- The database server ports are ports 15 to 22.

For other rack configurations (for example, X8M-8 systems with three database servers and 11 storage servers) a different procedure and different RoCE Network Fabric switch configuration files are required. Contact Oracle for further guidance.

- The procedure uses the following naming abbreviations and conventions:
 - The abbreviation for the existing rack is R1, and the new rack is R2.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies a spine switch.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the existing rack (R1).
- Most operations must be performed in multiple locations. For example, step [1.h](#) instructs you to update the firmware on all the RoCE Network Fabric leaf switches (R1LL, R1UL, R2LL, and R2UL). Pay attention to the instructions and keep track of your actions.

 Tip:

When a step must be performed on multiple switches, the instruction contains a list of the applicable switches. For example, (R1LL, R1UL, R2LL, and R2UL). You can use this list as a checklist to keep track of your actions.

- Perform operations sequentially, and complete every operation before proceeding. For example, run the entire command sequence at [3.a.i](#) as one operation and complete it before proceeding.
- All of commands that are run on a RoCE Network Fabric switch must be run while connected to the switch management interface as the switch administrator.

1. Prepare the systems.

- a. Position the new rack (R2) so that it is physically near the existing rack (R1).

The RDMA Network Fabric cables must be able to reach the switches in each rack.

For the required cross-rack cabling information, see [Two-Rack Cabling for X8M Racks](#).

- b. Power on all of the servers and network switches in the new rack (R2).

This includes the database servers, storage servers, RoCE Network Fabric leaf switches, and the Management Network Switch.

- c. Prepare the RoCE Network Fabric cables that you will use to interconnect the racks.

Label both ends of every cable.

For the required cross-rack cabling information, see [Two-Rack Cabling for X8M Racks](#).

- d. Connect the new rack (R2) to your existing management network.

Ensure that there are no IP address conflicts across the racks and that you can access the management interfaces on the RoCE Network Fabric switches.

- e. Ensure that you have a backup of the current switch configuration for each RoCE Network Fabric switch (R1LL, R1UL, R2LL, and R2UL).

See *Backing Up Settings on the RoCE Network Fabric Switch* in *Oracle Exadata Database Machine Maintenance Guide*.

- f. Download the required RoCE Network Fabric switch configuration files.

This procedure requires specific RoCE Network Fabric switch configuration files, which you must download from [My Oracle Support document 2704997.1](#).

 WARNING:

You must use different switch configuration files depending on whether your system uses Exadata Secure RDMA Fabric Isolation. Ensure that you download the correct archive that matches your system configuration.

For system configurations without Secure Fabric, download `online_multi-rack.zip`. For system configurations with Secure Fabric, download `online_SF_enabled_multi-rack.zip`.

Download and extract the archive containing the required RoCE Network Fabric switch configuration files. Place the files on a server with access to the management interfaces on the RoCE Network Fabric switches.

- g. Copy the required RoCE Network Fabric switch configuration files to the leaf switches on both racks.

You can use the following commands to copy the required configuration files to all of the RoCE Network Fabric switches on a system without Secure Fabric enabled:

```
i. # scp roce_multi_online_step3_R1_LL.cfg admin@R1LL_IP:/
```

```
ii. # scp roce_multi_online_step3_R2_LL.cfg admin@R2LL_IP:/
```

```
iii. # scp roce_multi_online_step4_R1_UL.cfg admin@R1UL_IP:/
```

```
iv. # scp roce_multi_online_step4_R2_UL.cfg admin@R2UL_IP:/
```

```
v. # scp roce_multi_online_step5.cfg admin@R1LL_IP:/
```

```
vi. # scp roce_multi_online_step5.cfg admin@R2LL_IP:/
```

On a system with Secure Fabric enabled, you can use the following commands:

```
i. # scp roce_SF_multi_online_step3_R1_LL.cfg admin@R1LL_IP:/
```

- ii. # scp roce_SF_multi_online_step3_R2_LL.cfg admin@R2LL_IP:/
- iii. # scp roce_SF_multi_online_step4_R1_UL.cfg admin@R1UL_IP:/
- iv. # scp roce_SF_multi_online_step4_R2_UL.cfg admin@R2UL_IP:/
- v. # scp roce_SF_multi_online_step5.cfg admin@R1LL_IP:/
- vi. # scp roce_SF_multi_online_step5.cfg admin@R2LL_IP:/

In the above commands, substitute the appropriate IP address or host name where applicable. For example, in place of *R1LL_IP*, substitute the management IP address or host name for the lower leaf switch (LL) on the existing rack (R1).

 Note:

The command examples in the rest of this procedure use the configuration files for a system configuration without Secure Fabric enabled. If required, adjust the commands to use the Secure Fabric-enabled switch configuration files.

- h. Update the firmware to the latest available release on all of the RoCE Network Fabric leaf switches (R1LL, R1UL, R2LL, and R2UL).
See Updating RoCE Network Fabric Switch Firmware in *Oracle Exadata Database Machine Maintenance Guide*.
- i. Examine the RoCE Network Fabric leaf switches (R1LL, R1UL, R2LL, and R2UL) and confirm the port categories for the cabled ports.

Run the `show interface status` command on every RoCE Network Fabric leaf switch:

- i. R1LL# `show interface status`
- ii. R1UL# `show interface status`
- iii. R2LL# `show interface status`
- iv. R2UL# `show interface status`

Examine the output and confirm the port categories as follows:

- Confirm that the inter-switch ports are ports 4 to 7, and ports 30 to 33.
- Confirm that the storage server ports are ports 8 to 14, and ports 23 to 29.
- Confirm that the database server ports are ports 15 to 22.

For example:

```
R1LL# show interface status
```

```
-----
Port          Name          Status    Vlan      Duplex Speed  Type
-----
mgmt0         --              connected routed    full   1000   --
-----

Port          Name          Status    Vlan      Duplex Speed  Type
-----
Eth1/1        --              xcvrAbsen 1      auto   auto   --
Eth1/2        --              xcvrAbsen 1      auto   auto   --
Eth1/3        --              xcvrAbsen 1      auto   auto   --
Eth1/4        ISL1            connected trunk    full   100G
QSFP-100G-CR4
Eth1/5        ISL2            connected trunk    full   100G
QSFP-100G-CR4
Eth1/6        ISL3            connected trunk    full   100G
QSFP-100G-CR4
Eth1/7        ISL4            connected trunk    full   100G
QSFP-100G-CR4
Eth1/8        celadm14        connected 3888      full   100G
QSFP-100G-CR4
Eth1/9        celadm13        connected 3888      full   100G
QSFP-100G-CR4
Eth1/10       celadm12        connected 3888      full   100G
QSFP-100G-CR4
Eth1/11       celadm11        connected 3888      full   100G
QSFP-100G-CR4
Eth1/12       celadm10        connected 3888      full   100G
QSFP-100G-CR4
Eth1/13       celadm09        connected 3888      full   100G
QSFP-100G-CR4
Eth1/14       celadm08        connected 3888      full   100G
QSFP-100G-CR4
Eth1/15       adm08           connected 3888      full   100G
QSFP-100G-CR4
Eth1/16       adm07           connected 3888      full   100G
QSFP-100G-CR4
Eth1/17       adm06           connected 3888      full   100G
QSFP-100G-CR4
Eth1/18       adm05           connected 3888      full   100G
QSFP-100G-CR4
Eth1/19       adm04           connected 3888      full   100G
QSFP-100G-CR4
Eth1/20       adm03           connected 3888      full   100G
QSFP-100G-CR4
Eth1/21       adm02           connected 3888      full   100G
QSFP-100G-CR4
Eth1/22       adm01           connected 3888      full   100G
QSFP-100G-CR4
Eth1/23       celadm07        connected 3888      full   100G
QSFP-100G-CR4
Eth1/24       celadm06        connected 3888      full   100G
QSFP-100G-CR4
```

Eth1/25	celadm05	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

- j. For each rack (R1 and R2), confirm the RoCE Network Fabric cabling by running the `verify_roce_cables.py` script.

The `verify_roce_cables.py` script uses two input files; one for database servers and storage servers (`nodes.rackN`), and another for switches (`switches.rackN`). In each file, every server or switch must be listed on separate lines. Use fully qualified domain names or IP addresses for each server and switch.

See My Oracle Support document 2587717.1 for download and detailed usage instructions.

Run the `verify_roce_cables.py` script against both of the racks:

- i.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
```
- ii.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2
```

Check that output in the `CABLE OK?` columns contains the `OK` status.

The following example shows the expected command results:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
SWITCH PORT (EXPECTED PEER)          LOWER LEAF (rack1sw-rocea0) :
CABLE OK?                UPPER LEAF (rack1sw-roceb0) : CABLE OK?
-----
Eth1/4 (ISL peer switch) : rack1sw-rocea0 Ethernet1/4 :
```

```
OK rack1sw-roceb0 Ethernet1/4 : OK
Eth1/5 (ISL peer switch) : rack1sw-rocea0 Ethernet1/5 :
OK rack1sw-roceb0 Ethernet1/5 : OK
Eth1/6 (ISL peer switch) : rack1sw-rocea0 Ethernet1/6 :
OK rack1sw-roceb0 Ethernet1/6 : OK
Eth1/7 (ISL peer switch) : rack1sw-rocea0 Ethernet1/7 :
OK rack1sw-roceb0 Ethernet1/7 : OK
Eth1/8 (RU39) : rack1celadm14 port-1 :
OK rack1celadm14 port-2 : OK
Eth1/9 (RU37) : rack1celadm13 port-1 :
OK rack1celadm13 port-2 : OK
Eth1/10 (RU35) : rack1celadm12 port-1 :
OK rack1celadm12 port-2 : OK
Eth1/11 (RU33) : rack1celadm11 port-1 :
OK rack1celadm11 port-2 : OK
Eth1/12 (RU31) : rack1celadm10 port-1 :
OK rack1celadm10 port-2 : OK
Eth1/13 (RU29) : rack1celadm09 port-1 :
OK rack1celadm09 port-2 : OK
Eth1/14 (RU27) : rack1celadm08 port-1 :
OK rack1celadm08 port-2 : OK
Eth1/15 (RU26) : rack1adm08 port-1 :
OK rack1adm08 port-2 : OK
Eth1/16 (RU25) : rack1adm07 port-1 :
OK rack1adm07 port-2 : OK
Eth1/17 (RU24) : rack1adm06 port-1 :
OK rack1adm06 port-2 : OK
Eth1/18 (RU23) : rack1adm05 port-1 :
OK rack1adm05 port-2 : OK
Eth1/19 (RU19) : rack1adm04 port-1 :
OK rack1adm04 port-2 : OK
Eth1/20 (RU18) : rack1adm03 port-1 :
OK rack1adm03 port-2 : OK
Eth1/21 (RU17) : rack1adm02 port-1 :
OK rack1adm02 port-2 : OK
Eth1/22 (RU16) : rack1adm01 port-1 :
OK rack1adm01 port-2 : OK
Eth1/23 (RU14) : rack1celadm07 port-1 :
OK rack1celadm07 port-2 : OK
Eth1/24 (RU12) : rack1celadm06 port-1 :
OK rack1celadm06 port-2 : OK
Eth1/25 (RU10) : rack1celadm05 port-1 :
OK rack1celadm05 port-2 : OK
Eth1/26 (RU08) : rack1celadm04 port-1 :
OK rack1celadm04 port-2 : OK
Eth1/27 (RU06) : rack1celadm03 port-1 :
OK rack1celadm03 port-2 : OK
Eth1/28 (RU04) : rack1celadm02 port-1 :
OK rack1celadm02 port-2 : OK
Eth1/29 (RU02) : rack1celadm01 port-1 :
OK rack1celadm01 port-2 : OK
Eth1/30 (ISL peer switch) : rack1sw-rocea0 Ethernet1/30 :
OK rack1sw-roceb0 Ethernet1/30 : OK
Eth1/31 (ISL peer switch) : rack1sw-rocea0 Ethernet1/31 :
OK rack1sw-roceb0 Ethernet1/31 : OK
Eth1/32 (ISL peer switch) : rack1sw-rocea0 Ethernet1/32 :
```

```

OK                               rack1sw-roceb0 Ethernet1/32 : OK
Eth1/33 (ISL peer switch)      : rack1sw-rocea0 Ethernet1/33 :
OK                               rack1sw-roceb0 Ethernet1/33 : OK

```

- k. For each rack (R1 and R2), verify the RoCE Network Fabric operation by using the `infinicheck` command.
- Use `infinicheck` with the `-z` option to clear the files that were created during the last run of the `infinicheck` command.
 - Use `infinicheck` with the `-s` option to set up user equivalence for password-less SSH across the RoCE Network Fabric.
 - Finally, verify the RoCE Network Fabric operation by using `infinicheck` with the `-b` option, which is recommended on newly imaged machines where it is acceptable to suppress the `cellip.ora` and `cellinit.ora` configuration checks.

In each command, the hosts input file (`hosts.rack1` and `hosts.rack2`) contains a list of database server RoCE Network Fabric IP addresses (2 RoCE Network Fabric IP addresses for each database server), and the cells input file (`cells.rack1` and `cells.rack2`) contains a list of RoCE Network Fabric IP addresses for the storage servers (2 RoCE Network Fabric IP addresses for each storage server).

- i. Use the following recommended command sequence on the existing rack (R1):

```

i. # cd /opt/oracle.SupportTools/ibdiagtools
   # ./infinicheck -g hosts.rack1 -c cells.rack1 -z

ii. # cd /opt/oracle.SupportTools/ibdiagtools
    # ./infinicheck -g hosts.rack1 -c cells.rack1 -s

iii. # cd /opt/oracle.SupportTools/ibdiagtools
     # ./infinicheck -g hosts.rack1 -c cells.rack1 -b

```

- ii. Use the following recommended command sequence on the new rack (R2):

```

i. # cd /opt/oracle.SupportTools/ibdiagtools
   # ./infinicheck -g hosts.rack2 -c cells.rack2 -z

ii. # cd /opt/oracle.SupportTools/ibdiagtools
    # ./infinicheck -g hosts.rack2 -c cells.rack2 -s

iii. # cd /opt/oracle.SupportTools/ibdiagtools
     # ./infinicheck -g hosts.rack2 -c cells.rack2 -b

```

The following example shows the expected command results for the final command in the sequence:

```

# cd /opt/oracle.SupportTools/ibdiagtools
# ./infinicheck -g hosts.rackN -c cells.rackN -b

INFINICHECK
[Network Connectivity, Configuration and Performance]

#### FABRIC TYPE TESTS ####
System type identified: RoCE

```

```

Verifying User Equivalence of user=root from all DBs to all CELLS.
#### RoCE CONFIGURATION TESTS ####
Checking for presence of RoCE devices on all DBs and CELLS
[SUCCESS].... RoCE devices on all DBs and CELLS look good
Checking for RoCE Policy Routing settings on all DBs and CELLS
[SUCCESS].... RoCE Policy Routing settings look good
Checking for RoCE DSCP ToS mapping on all DBs and CELLS
[SUCCESS].... RoCE DSCP ToS settings look good
Checking for RoCE PFC settings and DSCP mapping on all DBs and
CELLs
[SUCCESS].... RoCE PFC and DSCP settings look good
Checking for RoCE interface MTU settings. Expected value : 2300
[SUCCESS].... RoCE interface MTU settings look good
Verifying switch advertised DSCP on all DBs and CELLS ports ( )
[SUCCESS].... Advertised DSCP settings from RoCE switch looks good
#### CONNECTIVITY TESTS ####
[COMPUTE NODES -> STORAGE CELLS]
(60 seconds approx.)
(Will walk through QoS values: 0-6) [SUCCESS].....Results OK
[SUCCESS]..... All can talk to all storage cells
[COMPUTE NODES -> COMPUTE NODES]
...

```

2. Install the spine switches (R1SS and R2SS).

- a. Physically install and power up the spine switches in the existing rack (R1SS) and the new rack (R2SS).
 - i. Physically install each spine switch in RU1.
 - ii. For each spine switch, ensure that the management Ethernet interface is connected to the management network and then supply power.
 - iii. On each spine switch, perform the initial configuration steps outlined in *Configuring the Cisco Nexus C9336C-FX2 Switch*. Skip the step for applying the golden configuration settings as you will do this later.
 - iv. For each spine switch, perform a ping test to the management Ethernet interface to ensure that the switch is online and accessible.
- b. Apply the golden configuration settings to the new spine switches.

See *Applying Golden Configuration Settings on RoCE Network Fabric Switches* in *Oracle Exadata Database Machine Maintenance Guide*.

You can use the instance of `patchmgr` that you previously used to update the firmware on the leaf switches (in step 1.h).

Use a switch list file (`spines.lst`) to apply the golden configuration settings to both spine switches using one `patchmgr` command:

```

# cat spines.lst
R1SS_IP:mspine.201
R2SS_IP:mspine.202

# ./patchmgr --roceswitches spines.lst --apply-config -log_dir /tmp/
spinelogs

```

 Note:

In the switch list file, *R1SS_IP* is the management IP address or host name for the spine switch on the existing rack (R1SS) and *R2SS_IP* is the management IP address or host name for the spine switch on the new rack (R2SS).

- c. Upgrade the firmware on the spine switches.

See Updating RoCE Network Fabric Switch Firmware in *Oracle Exadata Database Machine Maintenance Guide*.

You can use the instance of `patchmgr` that you used in the previous step.

Use a switch list file (`spines.lst`) to perform the firmware upgrade on both spine switches using one `patchmgr` command:

```
# cat spines.lst
R1SS_IP:mspine.201
R2SS_IP:mspine.202

# ./patchmgr --roceswitches spines.lst --upgrade -log_dir /tmp/spinelogs
```

 Note:

In the switch list file, *R1SS_IP* is the management IP address or host name for the spine switch on the existing rack (R1SS) and *R2SS_IP* is the management IP address or host name for the spine switch on the new rack (R2SS).

- d. Connect the RoCE Network Fabric cables to the spine switches (R1SS and R2SS).

 WARNING:

At this stage, only connect the cables to the spine switches.

To avoid later complications, ensure that each cable connects to the correct switch and port.

DO NOT CONNECT ANY OF THE CABLES TO THE LEAF SWITCHES.

Use the cables that you prepared earlier (in step 1.c).

For the required cross-rack cabling information, see [Two-Rack Cabling for X8M Racks](#).

3. Perform the first round of configuration on the lower leaf switches (R1LL and R2LL).

Perform this step on the lower leaf switches (R1LL and R2LL) only.

 Note:

During this step, the lower leaf switch ports are shut down. While the R1LL ports are down, R1UL exclusively supports the RoCE Network Fabric. During this time, there is no redundancy in the RoCE Network Fabric, and availability cannot be maintained if R1UL goes down.

a. Shut down the switch ports on the lower leaf switches (R1LL and R2LL).

i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/4-33
R1LL(config-if-range)# shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

ii. Repeat the command sequence on R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/4-33
R2LL(config-if-range)# shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

b. Reconfigure the lower leaf switch ports (R1LL and R2LL) .

For each switch, you must use the correct corresponding switch configuration file, which you earlier copied to the switch (in step 1.g).

i. On R1LL, the switch configuration file name must end with step3_R1_LL.cfg:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# run-script bootflash:///
roce_multi_online_step3_R1_LL.cfg | grep 'none'
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

- ii. On R2LL, the switch configuration file name must end with `step3_R2_LL.cfg`:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# run-script bootflash:///
roce_multi_online_step3_R2_LL.cfg | grep 'none'
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

 Note:

This step can take approximately 5 to 8 minutes on each switch.

- c. Start the inter-switch ports on the lower leaf switches (R1LL and R2LL) .

- i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/4-7, ethernet 1/30-33
R1LL(config-if-range)# no shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

- ii. Repeat the command sequence on R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/4-7, ethernet 1/30-33
R2LL(config-if-range)# no shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

- d. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- e. Verify the status of the inter-switch ports on the lower leaf switches (R1LL and R2LL) .
Run the `show interface status` command on each lower leaf switch:

i. R1LL# **show interface status**

ii. R2LL# **show interface status**

Examine the output to ensure that the inter-switch ports are connected.

For example:

R1LL# **show interface status**

```

-----
-----
Port              Name              Status      Vlan      Duplex Speed  Type
-----
mgmt0             --                connected   routed    full   1000   --
-----
-----
Port              Name              Status      Vlan      Duplex Speed  Type
-----
Eth1/1            --                xcvrAbsen  1          auto   auto   --
Eth1/2            --                xcvrAbsen  1          auto   auto   --
Eth1/3            --                xcvrAbsen  1          auto   auto   --
Eth1/4            ISL1              connected   trunk     full   100G
QSFP-100G-CR4
Eth1/5            ISL2              connected   trunk     full   100G
QSFP-100G-CR4
Eth1/6            ISL3              connected   trunk     full   100G
QSFP-100G-CR4
Eth1/7            ISL4              connected   trunk     full   100G
QSFP-100G-CR4
Eth1/8            celadm14          disabled    3888      full   100G
QSFP-100G-CR4
Eth1/9            celadm13          disabled    3888      full   100G
QSFP-100G-CR4
Eth1/10           celadm12          disabled    3888      full   100G
QSFP-100G-CR4
Eth1/11           celadm11          disabled    3888      full   100G
QSFP-100G-CR4
Eth1/12           celadm10          disabled    3888      full   100G
QSFP-100G-CR4
Eth1/13           celadm09          disabled    3888      full   100G
QSFP-100G-CR4
Eth1/14           celadm08          disabled    3888      full   100G
QSFP-100G-CR4
Eth1/15           adm08             disabled    3888      full   100G
QSFP-100G-CR4
Eth1/16           adm07             disabled    3888      full   100G
QSFP-100G-CR4
Eth1/17           adm06             disabled    3888      full   100G
QSFP-100G-CR4
Eth1/18           adm05             disabled    3888      full   100G
QSFP-100G-CR4

```

Eth1/19	adm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

- f. Start the storage server ports on the lower leaf switches (R1LL and R2LL) .

i. On R1LL:

```

R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R1LL(config-if-range)# no shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#

```

ii. Repeat the command sequence on R2LL:

```

R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R2LL(config-if-range)# no shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#

```

- g. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- h. Verify the status of the storage server ports on the lower leaf switches (R1LL and R2LL).

Run the `show interface status` command on each lower leaf switch:

i. R1LL# `show interface status`

ii. R2LL# `show interface status`

Examine the output to ensure that the storage server ports are connected.

For example:

R1LL# `show interface status`

```

-----
-----
Port          Name          Status      Vlan      Duplex Speed  Type
-----
-----
mgmt0         --            connected   routed    full   1000   --
-----
-----
Port          Name          Status      Vlan      Duplex Speed  Type
-----
-----
Eth1/1        --            xcvrAbsen   1          auto   auto   --
Eth1/2        --            xcvrAbsen   1          auto   auto   --
Eth1/3        --            xcvrAbsen   1          auto   auto   --
Eth1/4        ISL1          connected   trunk      full   100G
QSFP-100G-CR4
Eth1/5        ISL2          connected   trunk      full   100G
QSFP-100G-CR4
Eth1/6        ISL3          connected   trunk      full   100G
QSFP-100G-CR4
Eth1/7        ISL4          connected   trunk      full   100G
QSFP-100G-CR4
Eth1/8        celadm14      connected   3888      full   100G
QSFP-100G-CR4

```

Eth1/9	celadm13	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/10	celadm12	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/11	celadm11	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/12	celadm10	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/13	celadm09	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/14	celadm08	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/15	adm08	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/16	adm07	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/17	adm06	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/18	adm05	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/19	adm04	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/20	adm03	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/21	adm02	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/22	adm01	disabled 3888	full	100G	
QSFP-100G-CR4					
Eth1/23	celadm07	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/24	celadm06	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/25	celadm05	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/26	celadm04	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/27	celadm03	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/28	celadm02	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/29	celadm01	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/30	ISL5	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/31	ISL6	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/32	ISL7	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/33	ISL8	connected trunk	full	100G	
QSFP-100G-CR4					
Eth1/34	--	xcvrAbsen 1	auto	auto	--
Eth1/35	--	xcvrAbsen 1	auto	auto	--
Eth1/36	--	xcvrAbsen 1	auto	auto	--
Po100	--	connected trunk	full	100G	--
Lo0	Routing loopback i	connected routed	auto	auto	--
Lo1	VTEP loopback inte	connected routed	auto	auto	--

```
Vlan1      --          down      routed      auto      auto      --
nve1       --          connected --          auto      auto      --
```

- i. Start the database server ports on the lower leaf switches (R1LL and R2LL).

- i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# interface ethernet 1/15-22
R1LL(config-if-range)# no shut
R1LL(config-if-range)# exit
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

- ii. Repeat the command sequence on R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# interface ethernet 1/15-22
R2LL(config-if-range)# no shut
R2LL(config-if-range)# exit
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

- j. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
 - k. Verify the status of the database server ports on the lower leaf switches (R1LL and R2LL).

Run the `show interface status` command on each lower leaf switch:

- i. R1LL# `show interface status`

- ii. R2LL# `show interface status`

Examine the output to ensure that the database server ports are `connected`.

For example:

```
R1LL# show interface status

-----
-----
Port          Name          Status      Vlan      Duplex Speed  Type
-----
-----
mgmt0         --            connected   routed    full   1000   --
```

Port	Name	Status	Vlan	Duplex	Speed	Type
Eth1/1	--	xcvrAbsen	1	auto	auto	--
Eth1/2	--	xcvrAbsen	1	auto	auto	--
Eth1/3	--	xcvrAbsen	1	auto	auto	--
Eth1/4	ISL1	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/5	ISL2	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/6	ISL3	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/7	ISL4	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/8	celadm14	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/9	celadm13	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/10	celadm12	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/11	celadm11	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/12	celadm10	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/13	celadm09	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/14	celadm08	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/15	adm08	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/16	adm07	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/17	adm06	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/18	adm05	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/19	adm04	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	connected	3888	full	100G	

QSFP-100G-CR4						
Eth1/28	celadm02	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	connected	trunk	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

 Note:

Before proceeding, ensure that you have completed all of the actions in step 3 on both lower leaf switches (R1LL and R2LL). If not, then ensure that you go back and perform the missing actions.

4. Perform the first round of configuration on the upper leaf switches (R1UL and R2UL).
Perform this step on the upper leaf switches (R1UL and R2UL) only.

 Note:

At the start of this step, the upper leaf switch ports are shut down. While the R1UL ports are down, R1LL exclusively supports the RoCE Network Fabric on the existing rack. During this time, there is no redundancy in the RoCE Network Fabric, and availability cannot be maintained if R1LL goes down.

- a. Shut down the upper leaf switch ports (R1UL and R2UL).
 - i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/4-33
R1UL(config-if-range)# shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
```

```
R1UL(config)# <Ctrl-Z>
R1UL#
```

- ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/4-33
R2UL(config-if-range)# shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- b. On both racks, remove the inter-switch links between the leaf switches (R1LL to R1UL, and R2LL to R2UL).

On every leaf switch, remove the cables for the inter-switch links:

- i. On R1LL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
 - ii. On R1UL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
 - iii. On R2LL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
 - iv. On R2UL, disconnect the inter-switch links from ports 04, 05, 06, 07, 30, 31, 32, and 33.
- c. On both racks, cable the upper leaf switch to both of the spine switches (R1UL and R2UL to R1SS and R2SS).

Connect the cables from the spine switches that you prepared earlier (in step 2.d).

Cable the switches as described in [Two-Rack Cabling for X8M Racks](#):

- i. On R1UL, cable ports 04, 05, 06, 07, 30, 31, 32, and 33 to R1SS and R2SS.
- ii. On R2UL, cable ports 04, 05, 06, 07, 30, 31, 32, and 33 to R1SS and R2SS.

Note:

Ensure that each cable connects to the correct switch and port at both ends. In addition to physically checking each connection, you can run the `show lldp neighbors` command on each network switch and examine the output to confirm correct connections. You can individually check each cable connection to catch and correct errors quickly.

- d. Reconfigure the upper leaf switch ports (R1UL and R2UL).

For each switch, you must use the correct corresponding switch configuration file, which you earlier copied to the switch (in step 1.g):

- i. On R1UL, the switch configuration file name must end with `step4_R1_UL.cfg`:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# run-script bootflash:///
roce_multi_online_step4_R1_UL.cfg | grep 'none'
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

- ii. On R2UL, the switch configuration file name must end with `step4_R2_UL.cfg`:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# run-script bootflash:///
roce_multi_online_step4_R2_UL.cfg | grep 'none'
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```



Note:

This step can take approximately 5 to 8 minutes on each switch.

- e. Check the status of the RoCE Network Fabric ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

- i. R1UL# `show interface status`

- ii. R2UL# `show interface status`

Examine the output to ensure that all of the cabled ports are disabled.

For example:

```
R1UL# show interface status
```

```
-----
-----
Port          Name          Status    Vlan      Duplex  Speed
Type
-----
-----
mgmt0         --            connected routed    full    1000    --
```

Port Type	Name	Status	Vlan	Duplex	Speed
...					
Eth1/4	RouterPort1	<i>disabled</i>	routed	full	100G
QSFP-100G-CR4					
Eth1/5	RouterPort2	<i>disabled</i>	routed	full	100G
QSFP-100G-CR4					
Eth1/6	RouterPort3	<i>disabled</i>	routed	full	100G
QSFP-100G-CR4					
Eth1/7	RouterPort4	<i>disabled</i>	routed	full	100G
QSFP-100G-CR4					
Eth1/8	celadm14	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/9	celadm13	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/10	celadm12	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/11	celadm11	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/12	celadm10	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/13	celadm09	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/14	celadm08	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/15	adm08	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/16	adm07	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/17	adm06	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
Eth1/18	adm05	<i>disabled</i>	3888	full	100G
QSFP-100G-CR4					
...					

 Note:

Before proceeding, ensure that you have completed all of the actions to this point in step 4 on both upper leaf switches (R1UL and R2UL). If not, then ensure that you go back and perform the missing actions.

- f. Verify the configuration of the upper leaf switches.

You can use the instance of `patchmgr` that you previously used to update the switch firmware (in step 1.h).

Use a switch list file (`ul.lst`) to check both upper leaf switches using one `patchmgr` command:

```
# cat ul.lst
R1UL_IP:mleaf.102
R2UL_IP:mleaf.104
```

On a system with Secure Fabric enabled, use the `msfleaf` tag in the switch list file:

```
# cat ul.lst
R1UL_IP:msfleaf.102
R2UL_IP:msfleaf.104
```

The following shows the recommended command and an example of the expected results:

```
# ./patchmgr --roceswitches ul.lst --verify-config -log_dir /tmp/log
2020-08-10 13:40:09 -0700      :Working: Initiating config
verification... Expect up to 6 minutes for each switch
Mon Aug 10 13:40:13 PDT 2020 1 of 4 :Verifying config on switch ...
...
Mon Aug 10 13:40:32 PDT 2020:      [INFO      ] Config matches
template: ...
Mon Aug 10 13:40:32 PDT 2020:      [SUCCESS   ] Config validation
successful!
2020-08-10 13:40:32 -0700
  Config check on RoCE switch(es)
2020-08-10 13:40:32 -0700
  Completed run of command: ./patchmgr --roceswitches ul.lst --verify-
config -log_dir /tmp/log
2020-08-10 13:40:32 -0700      :INFO      : config attempted on nodes in
file ul.lst: [R1UL_IP R2UL_IP]
2020-08-10 13:40:32 -0700      :INFO      : For details, check the
following files in /tmp/log:
2020-08-10 13:40:32 -0700      :INFO      : - updateRoceSwitch.log
2020-08-10 13:40:32 -0700      :INFO      : - updateRoceSwitch.trc
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.stdout
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.stderr
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.log
2020-08-10 13:40:32 -0700      :INFO      : - patchmgr.trc
2020-08-10 13:40:32 -0700      :INFO      : Exit status:0
2020-08-10 13:40:32 -0700      :INFO      : Exiting.
```

In the command output, verify that the switch configuration is good for both upper leaf switches. You can ignore messages about the ports that are down.

5. Finalize the configuration of the lower leaf switches (R1LL and R2LL).

Perform this step on the lower leaf switches (R1LL and R2LL) only.

a. Reconfigure the lower leaf switch ports (R1LL and R2LL).

Run the following command sequence on both of the lower leaf switches (R1LL and R2LL).

You must use the correct switch configuration file, which you earlier copied to the switch (in step 1.g). In this step, the configuration file name must end with `step5.cfg`.

i. On R1LL:

```
R1LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1LL(config)# run-script bootflash:///roce_multi_online_step5.cfg |
grep 'none'
R1LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1LL(config)# <Ctrl-Z>
R1LL#
```

ii. Repeat the command sequence on R2LL:

```
R2LL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2LL(config)# run-script bootflash:///roce_multi_online_step5.cfg |
grep 'none'
R2LL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2LL(config)# <Ctrl-Z>
R2LL#
```

 Note:

This step can take approximately 5 to 8 minutes on each switch.

b. On both racks, cable the lower leaf switch to both of the spine switches (R1LL and R2LL to R1SS and R2SS).

Connect the cables from the spine switches that you prepared earlier (in step 2.d).

Cable the switches as described in [Two-Rack Cabling for X8M Racks](#):

- i. On R1LL, cable ports 04, 05, 06, 07, 30, 31, 32, and 33 to R1SS and R2SS.
- ii. On R2LL, cable ports 04, 05, 06, 07, 30, 31, 32, and 33 to R1SS and R2SS.

 Note:

Ensure that each cable connects to the correct switch and port at both ends. In addition to physically checking each connection, you can run the `show lldp neighbors` command on each network switch and examine the output to confirm correct connections. You can individually check each cable connection to catch and correct errors quickly.

c. On the lower leaf switches, verify that all of the cabled RoCE Network Fabric ports are connected (R1LL and R2LL).

Run the `show interface status` command on each lower leaf switch:

i. R1LL# `show interface status`

ii. R2LL# `show interface status`

Examine the output to ensure that all of the cabled ports are connected.

For example:

R1LL# `show interface status`

```
-----
-----
Port          Name          Status    Vlan      Duplex  Speed
Type
-----
-----
mgmt0         --           connected routed     full    1000   --
-----
-----
Port          Name          Status    Vlan      Duplex  Speed
Type
-----
-----
...
Eth1/4        RouterPort1    connected routed     full    100G
QSFP-100G-CR4
Eth1/5        RouterPort2    connected routed     full    100G
QSFP-100G-CR4
Eth1/6        RouterPort3    connected routed     full    100G
QSFP-100G-CR4
Eth1/7        RouterPort4    connected routed     full    100G
QSFP-100G-CR4
Eth1/8        celadm14       connected 3888      full    100G
QSFP-100G-CR4
Eth1/9        celadm13       connected 3888      full    100G
QSFP-100G-CR4
Eth1/10       celadm12       connected 3888      full    100G
QSFP-100G-CR4
Eth1/11       celadm11       connected 3888      full    100G
QSFP-100G-CR4
Eth1/12       celadm10       connected 3888      full    100G
QSFP-100G-CR4
Eth1/13       celadm09       connected 3888      full    100G
QSFP-100G-CR4
Eth1/14       celadm08       connected 3888      full    100G
QSFP-100G-CR4
Eth1/15       adm08          connected 3888      full    100G
QSFP-100G-CR4
Eth1/16       adm07          connected 3888      full    100G
QSFP-100G-CR4
Eth1/17       adm06          connected 3888      full    100G
QSFP-100G-CR4
Eth1/18       adm05          connected 3888      full    100G
```

QSFP-100G-CR4

...

 Note:

Before proceeding, ensure that you have completed all of the actions to this point in step 5 on both lower leaf switches (R1LL and R2LL). If not, then ensure that you go back and perform the missing actions.

d. Verify the configuration of the lower leaf switches.

You can use the instance of `patchmgr` that you previously used to update the switch firmware (in step 1.h).

Use a switch list file (`ll.lst`) to check both lower leaf switches using one `patchmgr` command:

```
# cat ll.lst
R1LL_IP:mleaf.101
R2LL_IP:mleaf.103
```

On a system with Secure Fabric enabled, use the `msfleaf` tag in the switch list file:

```
# cat ll.lst
R1LL_IP:msfleaf.101
R2LL_IP:msfleaf.103
```

The following shows the recommended command and an example of the expected results:

```
# ./patchmgr --roceswitches ll.lst --verify-config -log_dir /tmp/log
2020-08-10 13:45:09 -0700      :Working: Initiating config
verification... Expect up to 6 minutes for each switch
Mon Aug 10 13:45:13 PDT 2020 1 of 4 :Verifying config on switch ...
...
Mon Aug 10 13:45:32 PDT 2020:      [INFO      ] Config matches
template: ...
Mon Aug 10 13:45:32 PDT 2020:      [SUCCESS   ] Config validation
successful!
2020-08-10 13:45:32 -0700
Config check on RoCE switch(es)
2020-08-10 13:45:32 -0700
Completed run of command: ./patchmgr --roceswitches ll.lst --verify-
config -log_dir /tmp/log
2020-08-10 13:45:32 -0700      :INFO      : config attempted on nodes in
file ll.lst: [R1LL_IP R2LL_IP]
2020-08-10 13:45:32 -0700      :INFO      : For details, check the
following files in /tmp/log:
2020-08-10 13:45:32 -0700      :INFO      : - updateRoceSwitch.log
2020-08-10 13:45:32 -0700      :INFO      : - updateRoceSwitch.trc
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.stdout
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.stderr
2020-08-10 13:45:32 -0700      :INFO      : - patchmgr.log
```



```

2020-08-10 13:45:32 -0700      :INFO   : - patchmgr.trc
2020-08-10 13:45:32 -0700      :INFO   : Exit status:0
2020-08-10 13:45:32 -0700      :INFO   : Exiting.

```

In the command output, verify that the switch configuration is good for both lower leaf switches.

- e. Verify that nve is up on the lower leaf switches (R1LL and R2LL).

Run the following command on each lower leaf switch and examine the output:

- i. R1LL# **show nve peers**

- ii. R2LL# **show nve peers**

At this point, you should see one nve peer with State=Up.

For example:

```

R1LL# show nve peers
Interface          Peer-IP State LearnType  Uptime          Router-Mac
-----
nve1               100.64.1.103 Up        CP 00:04:29      n/a

```

- f. Verify that BGP is up on the lower leaf switches (R1LL and R2LL).

Run the following command on each lower leaf switch and examine the output:

- i. R1LL# **show logging log | grep BGP**

- ii. R2LL# **show logging log | grep BGP**

Look for two entries with Up in the rightmost column that are associated with different IP addresses.

For example:

```

R1LL# show logging log | grep BGP
2020 Aug 10 13:47:13 R1LL %BGP-5-ADJCHANGE: bgp- [29342] (default)
neighbor 100.64.0.201 Up
2020 Aug 10 13:47:24 R1LL %BGP-5-ADJCHANGE: bgp- [29342] (default)
neighbor 100.64.0.202 Up

```

6. Finalize the configuration of the upper leaf switches (R1UL and R2UL).

Perform this step on the upper leaf switches (R1UL and R2UL) only.

- a. Start the inter-switch ports on the upper leaf switches (R1UL and R2UL).

- i. On R1UL:

```

R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/4-7, ethernet 1/30-33
R1UL(config-if-range)# no shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%

```

```
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/4-7, ethernet 1/30-33
R2UL(config-if-range) # no shut
R2UL(config-if-range) # exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- b. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- c. Verify the status of the inter-switch ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

i. R1UL# `show interface status`

ii. R2UL# `show interface status`

Examine the output to ensure that the inter-switch ports are connected.

For example:

```
R1UL# show interface status
```

```
-----
-----
Port          Name                Status      Vlan      Duplex Speed  Type
-----
-----
mgmt0         --                  connected   routed    full   1000   --
-----
-----
Port          Name                Status      Vlan      Duplex Speed  Type
-----
-----
Eth1/1        --                  xcvrAbsen  1          auto   auto   --
Eth1/2        --                  xcvrAbsen  1          auto   auto   --
Eth1/3        --                  xcvrAbsen  1          auto   auto   --
Eth1/4        ISL1                connected   routed    full   100G
QSFP-100G-CR4
Eth1/5        ISL2                connected   routed    full   100G
QSFP-100G-CR4
Eth1/6        ISL3                connected   routed    full   100G
QSFP-100G-CR4
```

Eth1/7	ISL4	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/8	celadm14	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/9	celadm13	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/10	celadm12	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/11	celadm11	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/12	celadm10	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/13	celadm09	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/14	celadm08	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/15	adm08	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/16	adm07	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/17	adm06	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/18	adm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/19	adm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/26	celadm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	<i>connected</i>	routed	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--

Eth1/36	--	xcvrAbsen 1	auto	auto	--
Po100	--	connected trunk	full	100G	--
Lo0	Routing loopback i	connected routed	auto	auto	--
Lo1	VTEP loopback inte	connected routed	auto	auto	--
Vlan1	--	down routed	auto	auto	--
nve1	--	connected --	auto	auto	--

- d. Start the storage server ports on the upper leaf switches (R1UL and R2UL).

i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R1UL(config-if-range)# no shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/8-14, ethernet 1/23-29
R2UL(config-if-range)# no shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- e. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- f. Verify the status of the storage server ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

i. R1UL# `show interface status`

ii. R2UL# `show interface status`

Examine the output to ensure that the storage server ports are connected.

For example:

```
R1UL# show interface status
```

```
-----
-----
```

Port	Name	Status	Vlan	Duplex	Speed	Type

mgmt0	--	connected	routed	full	1000	--

Port	Name	Status	Vlan	Duplex	Speed	Type

Eth1/1	--	xcvrAbsen	1	auto	auto	--
Eth1/2	--	xcvrAbsen	1	auto	auto	--
Eth1/3	--	xcvrAbsen	1	auto	auto	--
Eth1/4	ISL1	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/5	ISL2	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/6	ISL3	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/7	ISL4	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/8	celadm14	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/9	celadm13	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/10	celadm12	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/11	celadm11	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/12	celadm10	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/13	celadm09	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/14	celadm08	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/15	adm08	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/16	adm07	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/17	adm06	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/18	adm05	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/19	adm04	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/20	adm03	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/21	adm02	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/22	adm01	disabled	3888	full	100G	
QSFP-100G-CR4						
Eth1/23	celadm07	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/24	celadm06	<i>connected</i>	3888	full	100G	
QSFP-100G-CR4						
Eth1/25	celadm05	<i>connected</i>	3888	full	100G	

QSFP-100G-CR4						
Eth1/26	celadm04	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/27	celadm03	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/28	celadm02	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/29	celadm01	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/30	ISL5	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/31	ISL6	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/32	ISL7	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/33	ISL8	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/34	--	xcvrAbsen	1	auto	auto	--
Eth1/35	--	xcvrAbsen	1	auto	auto	--
Eth1/36	--	xcvrAbsen	1	auto	auto	--
Po100	--	connected	trunk	full	100G	--
Lo0	Routing loopback i	connected	routed	auto	auto	--
Lo1	VTEP loopback inte	connected	routed	auto	auto	--
Vlan1	--	down	routed	auto	auto	--
nve1	--	connected	--	auto	auto	--

- g. Start the database server ports on the upper leaf switches (R1UL and R2UL).

- i. On R1UL:

```
R1UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1UL(config)# interface ethernet 1/15-22
R1UL(config-if-range)# no shut
R1UL(config-if-range)# exit
R1UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R1UL(config)# <Ctrl-Z>
R1UL#
```

- ii. Repeat the command sequence on R2UL:

```
R2UL# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2UL(config)# interface ethernet 1/15-22
R2UL(config-if-range)# no shut
R2UL(config-if-range)# exit
R2UL(config)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete
R2UL(config)# <Ctrl-Z>
R2UL#
```

- h. Wait for 5 minutes to ensure that the ports you just started are fully operational before continuing.
- i. Verify the status of the database server ports on the upper leaf switches (R1UL and R2UL).

Run the `show interface status` command on each upper leaf switch:

i. R1UL# `show interface status`

ii. R2UL# `show interface status`

Examine the output to ensure that the database server ports are connected.

For example:

R1UL# `show interface status`

Port	Name	Status	Vlan	Duplex	Speed	Type
mgmt0	--	connected	routed	full	1000	--
Eth1/1	--	xcvrAbsen	1	auto	auto	--
Eth1/2	--	xcvrAbsen	1	auto	auto	--
Eth1/3	--	xcvrAbsen	1	auto	auto	--
Eth1/4	ISL1	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/5	ISL2	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/6	ISL3	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/7	ISL4	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/8	celadm14	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/9	celadm13	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/10	celadm12	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/11	celadm11	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/12	celadm10	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/13	celadm09	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/14	celadm08	connected	3888	full	100G	
QSFP-100G-CR4						
Eth1/15	adm08	connected	3888	full	100G	
QSFP-100G-CR4						

Eth1/16	adm07	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/17	adm06	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/18	adm05	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/19	adm04	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/20	adm03	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/21	adm02	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/22	adm01	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/23	celadm07	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/24	celadm06	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/25	celadm05	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/26	celadm04	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/27	celadm03	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/28	celadm02	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/29	celadm01	connected 3888	full	100G	
QSFP-100G-CR4					
Eth1/30	ISL5	connected routed	full	100G	
QSFP-100G-CR4					
Eth1/31	ISL6	connected routed	full	100G	
QSFP-100G-CR4					
Eth1/32	ISL7	connected routed	full	100G	
QSFP-100G-CR4					
Eth1/33	ISL8	connected routed	full	100G	
QSFP-100G-CR4					
Eth1/34	--	xcvrAbsen 1	auto	auto	--
Eth1/35	--	xcvrAbsen 1	auto	auto	--
Eth1/36	--	xcvrAbsen 1	auto	auto	--
Po100	--	connected trunk	full	100G	--
Lo0	Routing loopback i	connected routed	auto	auto	--
Lo1	VTEP loopback inte	connected routed	auto	auto	--
Vlan1	--	down routed	auto	auto	--
nve1	--	connected --	auto	auto	--

- j. Verify that nve is up on the leaf switches (R1LL, R1UL, R2LL, and R2UL).

Run the following command on each leaf switch and examine the output:

i. R1LL# **show nve peers**

ii. R1UL# **show nve peers**

iii. R2LL# **show nve peers**

iv. R2UL# show nve peers

In the output, you should see three nve peers with State=Up.

For example:

```
R1UL# show nve peers
Interface      Peer-IP State LearnType  Uptime      Router-Mac
-----
nve1           100.64.1.101 Up        CP 00:04:29 n/a
nve1           100.64.1.103 Up        CP 00:07:48 n/a
nve1           100.64.1.104 Up        CP 00:04:10 n/a
```

k. Verify that BGP is up on the upper leaf switches (R1UL and R2UL).

Run the following command on each upper leaf switch and examine the output:

i. R1UL# show logging log | grep BGP**ii. R2UL# show logging log | grep BGP**

In the output, look for two entries with Up in the rightmost column that are associated with different IP addresses.

For example:

```
R1UL# show logging log | grep BGP
2020 Aug 10 13:57:13 R1UL %BGP-5-ADJCHANGE: bgp- [32782] (default)
neighbor 100.64.0.201 Up
2020 Aug 10 13:57:24 R1UL %BGP-5-ADJCHANGE: bgp- [32782] (default)
neighbor 100.64.0.202 Up
```

7. For each rack (R1 and R2), confirm the multi-rack cabling by running the verify_roce_cables.py script.

The `verify_roce_cables.py` script uses two input files; one for database servers and storage servers (`nodes.rackN`), and another for switches (`switches.rackN`). In each file, every server or switch must be listed on separate lines. Use fully qualified domain names or IP addresses for each server and switch.

See My Oracle Support document 2587717.1 for download and detailed usage instructions.

Run the `verify_roce_cables.py` script against both of the racks:

```
a. # cd /opt/oracle.SupportTools/ibdiagtools
   # ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
```

```
b. # cd /opt/oracle.SupportTools/ibdiagtools
   # ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2
```

Check the output of the `verify_roce_cables.py` script against the tables in [Two-Rack Cabling for X8M Racks](#). Also, check that output in the `CABLE OK?` columns contains the OK status.

The following examples show extracts of the expected command results:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
SWITCH PORT (EXPECTED PEER)  LOWER LEAF (rack1sw-rocea0) : CABLE OK?  UPPER LEAF (rack1sw-roceb0) : CABLE OK?
-----
Eth1/4 (ISL peer switch)      : rack1sw-roces0 Ethernet1/17 : OK          rack1sw-roces0
Ethernet1/9 : OK
Eth1/5 (ISL peer switch)      : rack1sw-roces0 Ethernet1/13 : OK          rack1sw-roces0
Ethernet1/5 : OK
Eth1/6 (ISL peer switch)      : rack1sw-roces0 Ethernet1/19 : OK          rack1sw-roces0
Ethernet1/11 : OK
Eth1/7 (ISL peer switch)      : rack1sw-roces0 Ethernet1/15 : OK          rack1sw-roces0
Ethernet1/7 : OK
Eth1/12 (celadm10)            : rack1celadm10 port-1      : OK          rack1celadm10
port-2      : OK
Eth1/13 (celadm09)            : rack1celadm09 port-1      : OK          rack1celadm09
port-2      : OK
Eth1/14 (celadm08)            : rack1celadm08 port-1      : OK          rack1celadm08
port-2      : OK
...
Eth1/15 (adm08)                : rack1dbadm08 port-1      : OK          rack1dbadm08
port-2      : OK
Eth1/16 (adm07)                : rack1dbadm07 port-1      : OK          rack1dbadm07
port-2      : OK
Eth1/17 (adm06)                : rack1dbadm06 port-1      : OK          rack1dbadm06
port-2      : OK
...
Eth1/30 (ISL peer switch)      : rack2sw-roces0 Ethernet1/17 : OK          rack2sw-roces0
Ethernet1/9 : OK
Eth1/31 (ISL peer switch)      : rack2sw-roces0 Ethernet1/13 : OK          rack2sw-roces0
Ethernet1/5 : OK
Eth1/32 (ISL peer switch)      : rack2sw-roces0 Ethernet1/19 : OK          rack2sw-roces0
Ethernet1/11 : OK
Eth1/33 (ISL peer switch)      : rack2sw-roces0 Ethernet1/15 : OK          rack2sw-roces0
Ethernet1/7 : OK

# ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2
SWITCH PORT (EXPECTED PEER)  LOWER LEAF (rack2sw-rocea0) : CABLE OK?  UPPER LEAF (rack2sw-roceb0) : CABLE OK?
-----
Eth1/4 (ISL peer switch)      : rack1sw-roces0 Ethernet1/18 : OK          rack1sw-roces0
Ethernet1/10 : OK
...
```

8. Verify the RoCE Network Fabric operation across both interconnected racks by using the `infinicheck` command.

Use the following recommended command sequence to verify the RoCE Network Fabric operation across both racks.

In each command, `hosts.all` contains a list of database server RoCE Network Fabric IP addresses from both racks (2 RoCE Network Fabric IP addresses for each database

server), and `cells.all` contains a list of RoCE Network Fabric IP addresses for the storage servers from both racks (2 RoCE Network Fabric IP addresses for each storage server).

- a.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./infinicheck -g hosts.all -c cells.all -z
```
- b.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./infinicheck -g hosts.all -c cells.all -s
```
- c.

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./infinicheck -g hosts.all -c cells.all -b
```

See step 1.k for most information about each `infinicheck` command.

The following example shows the expected command results for the final command in the sequence:

```
# cd /opt/oracle.SupportTools/ibdiagtools
# ./infinicheck -g hosts.all -c cells.all -b

INFINICHECK
    [Network Connectivity, Configuration and Performance]

        ##### FABRIC TYPE TESTS #####
System type identified: RoCE
Verifying User Equivalence of user=root from all DBs to all CELLS.
        ##### RoCE CONFIGURATION TESTS #####
            Checking for presence of RoCE devices on all DBs and CELLS
[SUCCESS].... RoCE devices on all DBs and CELLS look good
            Checking for RoCE Policy Routing settings on all DBs and CELLS
[SUCCESS].... RoCE Policy Routing settings look good
            Checking for RoCE DSCP ToS mapping on all DBs and CELLS
[SUCCESS].... RoCE DSCP ToS settings look good
            Checking for RoCE PFC settings and DSCP mapping on all DBs and CELLS
[SUCCESS].... RoCE PFC and DSCP settings look good
            Checking for RoCE interface MTU settings. Expected value : 2300
[SUCCESS].... RoCE interface MTU settings look good
            Verifying switch advertised DSCP on all DBs and CELLS ports ( )
[SUCCESS].... Advertised DSCP settings from RoCE switch looks good
        ##### CONNECTIVITY TESTS #####
            [COMPUTE NODES -> STORAGE CELLS]
                (60 seconds approx.)
                (Will walk through QoS values: 0-6) [SUCCESS].....Results OK
[SUCCESS]..... All can talk to all storage cells
            [COMPUTE NODES -> COMPUTE NODES]
...

```

At this point, both racks share the RoCE Network Fabric, and the combined system is ready for further configuration.

Cabling Two RoCE Network Fabric Racks Together with Down Time Allowed

Use this simpler procedure to cable together two racks with RoCE Network Fabric where some down-time can be tolerated.

This procedure is for systems with RoCE Network Fabric (X8M or later) using Oracle Exadata System Software Release 20.1.0 or later.

In this procedure, the existing rack is R1, and the new rack is R2.

Use the applicable cabling tables depending on your system configuration:

- [Two-Rack Cabling for RA21 and Later Model Racks](#)
- [Two-Rack Cabling for X8M Racks](#)
- [Two-Rack Cabling for a System Combining an X8M Rack and a Later Model Rack](#)

1. Ensure the new rack is near the existing rack.

The RDMA Network Fabric cables must be able to reach the servers in each rack.

2. Ensure you have a backup of the current switch configuration for each switch in the existing and new rack.

See *Backing Up Settings on the RoCE Network Fabric Switch* in *Oracle Exadata Database Machine Maintenance Guide*.

3. Shut down all servers on both the new rack (R2) and the existing rack (R1).

The switches should remain available.

4. Update the firmware to the latest available release on all of the RoCE Network Fabric switches.

For this step, treat all of the switches as if they belong to a single rack system.

See *Updating RoCE Network Fabric Switch Firmware* in *Oracle Exadata Database Machine Maintenance Guide*.

5. Apply the multi-rack golden configuration settings on the RoCE Network Fabric switches.

Use the procedure described in *Applying Golden Configuration Settings on RoCE Network Fabric Switches*, in *Oracle Exadata Database Machine Maintenance Guide*.

6. Enable the leaf switch server ports.

The leaf switch server ports may be disabled as a consequence of applying the multi-rack golden configuration settings in the previous step.

To ensure that the leaf switch server ports are enabled, log in to each of the four leaf switches and run the following commands on each leaf switch:

```
rack1sw-rocea0# config term
rack1sw-rocea0# int eth1/8-30
rack1sw-rocea0# no shut
rack1sw-rocea0# copy running-config startup-config
```

7. Perform the physical cabling of the switches.

- a. In Rack 2, remove the existing inter-switch connections between the two leaf switches, R2UL and R2LL.

- b. In Rack 2, cable each leaf switch to the spine switches using the applicable cabling tables.
 - c. In Rack 1, remove the existing inter-switch connections between the two leaf switches, R1UL and R1LL.
 - d. In Rack 1, cable each leaf switch to the spine switches using the applicable cabling tables.
8. Confirm each switch is available and connected.

For each of the six switches (two spine switches and four leaf switches), confirm the output from the `show interface status` command shows `connected` and `100G` for each connected inter-switch port. Use the appropriate cabling tables to identify the ports that should be connected.

In the following examples, the leaf switches are ports Eth1/4 to Eth1/7, and Eth1/30 to Eth1/33. The spine switches are ports Eth1/5 to Eth1/20.

When run from a spine switch, the output should be similar to the following:

```
rack1sw-roces0# show interface status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
mgmt0	--	connected	routed	full	1000	--
Port	Name	Status	Vlan	Duplex	Speed	Type
...						
Eth1/5	RouterPort5	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/6	RouterPort6	connected	routed	full	100G	
QSFP-100G-SR4						
Eth1/7	RouterPort7	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/8	RouterPort8	connected	routed	full	100G	
QSFP-100G-SR4						
Eth1/9	RouterPort9	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/10	RouterPort10	connected	routed	full	100G	
QSFP-100G-SR4						
Eth1/11	RouterPort11	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/12	RouterPort12	connected	routed	full	100G	
QSFP-100G-SR4						
Eth1/13	RouterPort13	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/14	RouterPort14	connected	routed	full	100G	
QSFP-100G-SR4						
Eth1/15	RouterPort15	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/16	RouterPort16	connected	routed	full	100G	
QSFP-100G-SR4						
Eth1/17	RouterPort17	connected	routed	full	100G	

```

QSFP-100G-CR4
Eth1/18      RouterPort18    connected routed    full    100G
QSFP-100G-SR4
Eth1/19      RouterPort19    connected routed    full    100G
QSFP-100G-CR4
Eth1/20      RouterPort20    connected routed    full    100G
QSFP-100G-SR4
Eth1/21      RouterPort21    xcvrAbsen    routed    full    100G
--
...

```

When run from a leaf switch, the output should be similar to the following:

```
rack1sw-rocea0# show interface status
```

```

-----
Port          Name          Status      Vlan      Duplex  Speed  Type
-----
mgmt0         --            connected   routed    full    1000   --
-----
Port          Name          Status      Vlan      Duplex  Speed  Type
-----
...
Eth1/4        RouterPort1    connected routed    full    100G
QSFP-100G-CR4
Eth1/5        RouterPort2    connected routed    full    100G
QSFP-100G-CR4
Eth1/6        RouterPort3    connected routed    full    100G
QSFP-100G-CR4
Eth1/7        RouterPort4    connected routed    full    100G
QSFP-100G-CR4
Eth1/8        celadm14       connected   3888      full    100G
QSFP-100G-CR4
...
Eth1/29       celadm01       connected   3888      full    100G
QSFP-100G-CR4
Eth1/30       RouterPort5    connected routed    full    100G
QSFP-100G-SR4
Eth1/31       RouterPort6    connected routed    full    100G
QSFP-100G-SR4
Eth1/32       RouterPort7    connected routed    full    100G
QSFP-100G-SR4
Eth1/33       RouterPort8    connected routed    full    100G
QSFP-100G-SR4
...

```

9. Check the neighbor discovery for every switch in racks R1 and R2.

Log in to each switch and use the `show lldp neighbors` command. Make sure that all switches are visible and check the switch ports assignment against the applicable cabling tables.

A spine switch should see the two leaf switches in each rack, but not the other spine switch. The output for a spine switch should be similar to the following:

 Note:

The interfaces output in the `Port ID` column are different for each switch based on the applicable cabling tables.

```
rack1sw-roces0# show lldp neighbors
...
Device ID           Local Intf      Hold-time  Capability  Port ID
rack1-adm0          mgmt0          120        BR          Ethernet1/47
rack1sw-roceb0      Eth1/5         120        BR          Ethernet1/5
rack2sw-roceb0      Eth1/6         120        BR          Ethernet1/5
rack1sw-roceb0      Eth1/7         120        BR          Ethernet1/7
rack2sw-roceb0      Eth1/8         120        BR          Ethernet1/7
rack1sw-roceb0      Eth1/9         120        BR          Ethernet1/4
rack2sw-roceb0      Eth1/10        120        BR          Ethernet1/4
rack1sw-roceb0      Eth1/11        120        BR          Ethernet1/6
rack2sw-roceb0      Eth1/12        120        BR          Ethernet1/6
rack1sw-rocea0      Eth1/13        120        BR          Ethernet1/5
rack2sw-rocea0      Eth1/14        120        BR          Ethernet1/5
rack1sw-rocea0      Eth1/15        120        BR          Ethernet1/7
rack2sw-rocea0      Eth1/16        120        BR          Ethernet1/7
rack1sw-rocea0      Eth1/17        120        BR          Ethernet1/4
rack2sw-rocea0      Eth1/18        120        BR          Ethernet1/4
rack1sw-rocea0      Eth1/19        120        BR          Ethernet1/6
rack2sw-rocea0      Eth1/20        120        BR          Ethernet1/6
Total entries displayed: 17
```

Each leaf switch should see the two spine switches, but not the other leaf switches. The output for a leaf switch should be similar to the following:

 Note:

The interfaces output in the `Port ID` column are different for each switch based on the applicable cabling tables.

```
rack1sw-rocea0# show lldp neighbors
...
Device ID           Local Intf      Hold-time  Capability  Port ID
switch              mgmt0          120        BR          Ethernet1/46
rack1sw-roces0      Eth1/4         120        BR          Ethernet1/17
rack1sw-roces0      Eth1/5         120        BR          Ethernet1/13
rack1sw-roces0      Eth1/6         120        BR          Ethernet1/19
rack1sw-roces0      Eth1/7         120        BR          Ethernet1/15
rack2sw-roces0      Eth1/30        120        BR          Ethernet1/17
rack2sw-roces0      Eth1/31        120        BR          Ethernet1/13
rack2sw-roces0      Eth1/32        120        BR          Ethernet1/19
rack2sw-roces0      Eth1/33        120        BR          Ethernet1/15
```

```
rocetoi-ext-sw      Eth1/36      120      BR      Ethernet1/49
Total entries displayed: 10
```

10. Power on all servers in racks R1 and R2.

11. For each rack, confirm the multi-rack cabling by running the `verify_roce_cables.py` script.

Refer to My Oracle Support Doc ID 2587717.1 for download and usage instructions.

Check the output of the `verify_roce_cables.py` script against the applicable cabling tables. Also, check that output in the `CABLE OK?` columns contains the `OK` status.

When running the script, two input files are used, one for nodes and one for switches. Each file should contain the servers or switches on separate lines. Use fully qualified domain names or IP addresses for each server and switch.

The following output is a partial example of the command results:

```
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
SWITCH PORT (EXPECTED PEER)  LEAF-1 (rack1sw-rocea0)      : CABLE OK?  LEAF-2 (rack1sw-
roceb0)      : CABLE OK?
-----
Eth1/4 (ISL peer switch)    : rack1sw-roces0 Ethernet1/17 : OK          rack1sw-roces0
Ethernet1/9 : OK
Eth1/5 (ISL peer switch)    : rack1sw-roces0 Ethernet1/13 : OK          rack1sw-roces0
Ethernet1/5 : OK
Eth1/6 (ISL peer switch)    : rack1sw-roces0 Ethernet1/19 : OK          rack1sw-roces0
Ethernet1/11: OK
Eth1/7 (ISL peer switch)    : rack1sw-roces0 Ethernet1/15 : OK          rack1sw-roces0
Ethernet1/7 : OK
Eth1/12 (celadm10)          : rack1celadm10 port-1      : OK          rack1celadm10
port-2      : OK
Eth1/13 (celadm09)          : rack1celadm09 port-1      : OK          rack1celadm09
port-2      : OK
Eth1/14 (celadm08)          : rack1celadm08 port-1      : OK          rack1celadm08
port-2      : OK
...
Eth1/15 (adm08)             : rack1dbadm08 port-1      : OK          rack1dbadm08
port-2      : OK
Eth1/16 (adm07)             : rack1dbadm07 port-1      : OK          rack1dbadm07
port-2      : OK
Eth1/17 (adm06)             : rack1dbadm06 port-1      : OK          rack1dbadm06
port-2      : OK
...
Eth1/30 (ISL peer switch)   : rack2sw-roces0 Ethernet1/17 : OK          rack2sw-roces0
Ethernet1/9 : OK
Eth1/31 (ISL peer switch)   : rack2sw-roces0 Ethernet1/13 : OK          rack2sw-roces0
Ethernet1/5 : OK
Eth1/32 (ISL peer switch)   : rack2sw-roces0 Ethernet1/19 : OK          rack2sw-roces0
Ethernet1/11: OK
Eth1/33 (ISL peer switch)   : rack2sw-roces0 Ethernet1/15 : OK          rack2sw-roces0
Ethernet1/7 : OK

# ./verify_roce_cables.py -n nodes.rack2 -s switches.rack2
SWITCH PORT (EXPECTED PEER)  LEAF-1 (rack2sw-rocea0)      : CABLE OK?  LEAF-2 (rack2sw-
```



```

roceb0)      : CABLE OK?
-----
Eth1/4 (ISL peer switch) : rack1sw-roces0 Ethernet1/18 : OK      rack1sw-roces0
Ethernet1/10: OK
...

```

12. Verify the RoCE Network Fabric operation by using the `infinicheck` command.

Use the following recommended command sequence. In each command, `hosts.lst` contains a list of database server RoCE Network Fabric IP addresses (2 RoCE Network Fabric IP addresses for each database server), and `cells.lst` contains a list of RoCE Network Fabric IP addresses for the storage servers (2 RoCE Network Fabric IP addresses for each storage server).

- Use `infinicheck` with the `-z` option to clear the files that were created during the last run of the `infinicheck` command. For example:

```
# /opt/oracle.SupportTools/ibdiagtools/infinicheck -g hosts.lst -c
cells.lst -z
```

- Use `infinicheck` with the `-s` option to set up user equivalence for password-less SSH across the RoCE Network Fabric. For example:

```
# /opt/oracle.SupportTools/ibdiagtools/infinicheck -g hosts.lst -c
cells.lst -s
```

- Finally, verify the RoCE Network Fabric operation by using `infinicheck` with the `-b` option, which is recommended on newly imaged machines where it is acceptable to suppress the `cellip.ora` and `cellinit.ora` configuration checks. For example:

```
# /opt/oracle.SupportTools/ibdiagtools/infinicheck -g hosts.lst -c
cells.lst -b
```

```

INFINICHECK
      [Network Connectivity, Configuration and Performance]

      ##### FABRIC TYPE TESTS #####
System type identified: RoCE
Verifying User Equivalence of user=root from all DBs to all CELLS.
      ##### RoCE CONFIGURATION TESTS #####
      Checking for presence of RoCE devices on all DBs and CELLS
[SUCCESS].... RoCE devices on all DBs and CELLS look good
      Checking for RoCE Policy Routing settings on all DBs and CELLS
[SUCCESS].... RoCE Policy Routing settings look good
      Checking for RoCE DSCP ToS mapping on all DBs and CELLS
[SUCCESS].... RoCE DSCP ToS settings look good
      Checking for RoCE PFC settings and DSCP mapping on all DBs and
CELLs
[SUCCESS].... RoCE PFC and DSCP settings look good
      Checking for RoCE interface MTU settings. Expected value : 2300
[SUCCESS].... RoCE interface MTU settings look good
      Verifying switch advertised DSCP on all DBs and CELLS ports ( )
[SUCCESS].... Advertised DSCP settings from RoCE switch looks good
      ##### CONNECTIVITY TESTS #####
[COMPUTE NODES -> STORAGE CELLS]

```

```
(60 seconds approx.)
(Will walk through QoS values: 0-6) [SUCCESS].....Results OK
[SUCCESS]..... All can talk to all storage cells
[COMPUTE NODES -> COMPUTE NODES]
...
```

Cabling Two InfiniBand Network Fabric Racks Together

Use this procedure to cable together two racks with InfiniBand Network Fabric.

This procedure assumes that the racks are adjacent to each other.

In the procedure, the existing rack is R1, and the new rack is R2.

1. Set the priority of the current, active Subnet Manager Master to 10 on the spine switch, as follows:
 - a. Log in to any RDMA Network Fabric switch on the active system.
 - b. Use the `getmaster` command to determine that the Subnet Manager Master is running on the spine switch.
 - c. Log in to the spine switch.
 - d. Use the `disablesm` command to stop Subnet Manager.
 - e. Use the `setsmpriority 10` command to set the priority to 10.
 - f. Use the `enablesm` command to restart Subnet Manager.
 - g. Repeat the `getmaster` command to ensure the Subnet Manager Master is running on the spine switch.
2. Ensure the new rack is near the existing rack. The RDMA Network Fabric cables must be able to reach the servers in each rack.
3. Completely shut down the new rack (R2).
4. Cable the two leaf switches R2 IB2 and R2 IB3 in the new rack according to [Two-Rack Cabling with InfiniBand Network Fabric](#). Note that you need to first remove the seven existing inter-switch connections between each leaf switch, as well as the two connections between the leaf switches and the spine switch in the new rack R2, not in the existing rack R1.
5. Verify both RDMA Network Fabric interfaces are up on all database nodes and storage cells. You can do this by running the `ibstat` command on each node and verifying both interfaces are up.
6. Power off leaf switch R1 IB2. This causes all the database servers and Exadata Storage Servers to fail over their RDMA Network Fabric traffic to R1 IB3.
7. Disconnect all seven inter-switch links between R1 IB2 and R1 IB3, as well as the one connection between R1 IB2 and the spine switch R1 IB1.
8. Cable leaf switch R1 IB2 according to [Two-Rack Cabling with InfiniBand Network Fabric](#).
9. Power on leaf switch R1 IB2.
10. Wait for three minutes for R1 IB2 to become completely operational.

To check the switch, log in to the switch and run the `ibswitches` command. The output should show three switches, R1 IB1, R1 IB2, and R1 IB3.

11. Verify both RDMA Network Fabric interfaces are up on all database nodes and storage cells. You can do this by running the `ibstat` command on each node and verifying both interfaces are up.
12. Power off leaf switch R1 IB3. This causes all the database servers and storage servers to fail over their RDMA Network Fabric traffic to R1 IB2.
13. Disconnect the one connection between R1 IB3 and the spine switch R1 IB1.
14. Cable leaf switch R1 IB3 according to [Two-Rack Cabling with InfiniBand Network Fabric](#).
15. Power on leaf switch R1 IB3.
16. Wait for three minutes for R1 IB3 to become completely operational.
To check the switch, log in to the switch and run the `ibswitches` command. The output should show three switches, R1 IB1, R1 IB2, and R1 IB3.
17. Power on all the InfiniBand switches in R2.
18. Wait for three minutes for the switches to become completely operational.
To check the switch, log in to the switch and run the `ibswitches` command. The output should show six switches, R1 IB1, R1 IB2, R1 IB3, R2 IB1, R2 IB2, and R2 IB3.
19. Ensure the Subnet Manager Master is running on R1 IB1 by running the `getmaster` command from any switch.
20. Power on all servers in R2.
21. Log in to spine switch R1 IB1, and lower its priority to 8 as follows:
 - a. Use the `disableesm` command to stop Subnet Manager.
 - b. Use the `setsmpriority 8` command to set the priority to 8.
 - c. Use the `enableesm` command to restart Subnet Manager.
22. Ensure Subnet Manager Master is running on one of the spine switches.

Cabling Several Racks Together

By using the following procedures, you can add another rack to an existing multi-rack system.

The procedures assume that the racks are adjacent to each other. The existing racks are R1, R2, ... R_n, and the new rack is R_{n+1}. For example, if you have four racks and you are adding a fifth rack, the existing racks are designated R1, R2, R3, and R4 and the new rack is designated R5.

Starting with Recovery Appliance RA21 models, you can cable up to 12 racks together without additional switches. For prior models (up to X8M), you can cable up to 18 racks together without additional switches.

The procedures differ for systems that use RoCE Network Fabric (X8M and later) and InfiniBand Network Fabric (X8 and earlier).

Cabling Several RoCE Network Fabric Racks Together

Use this procedure to add another rack to an existing multi-rack system with RoCE Network Fabric.

This procedure is for systems with RoCE Network Fabric (X8M or later) using Oracle Exadata System Software Release 20.1.0 or later.

 WARNING:

Take time to read and understand this procedure before implementation. Pay careful attention to all the instructions, not just the command examples. A system outage may occur if the instructions are not applied correctly.

In this procedure, the existing racks are R1, R2, ... ,Rn, and the new rack is Rn+1.

 Note:

Cabling three or more racks together requires no downtime for the existing racks R1, R2, ..., Rn. Only the new rack, Rn+1, is powered down.

Use the applicable cabling tables depending on your system:

- [Multi-Rack Cabling Tables for Recovery Appliance RA21 and Later Models](#)
- [Multi-Rack Cabling Tables for Recovery Appliance X8M Models](#)

In the following steps, these example switch names are used for the new rack (Rn+1):

- rack5sw-roces0: Rack 5 spine switch (R5SS)
- rack5sw-rocea0: Rack 5 lower leaf switch (R5LL)
- rack5sw-roceb0: Rack 5 upper leaf switch (R5UL)

1. Ensure the new rack is near the existing racks (R1, R2, ..., Rn).

The RDMA Network Fabric cables must be able to reach the servers in each rack.

2. Ensure you have a backup of the current switch configuration for each switch in the existing racks and the new rack.

For each switch, complete the steps in the Oracle Exadata Database Machine Maintenance Guide, section Backing Up Settings on the RoCE Network Fabric Switch.

3. Shut down all servers in the new rack (Rn+1).

Refer to Powering Off Oracle Exadata Rack. The switches must remain online and available.

4. Verify the configuration of the existing RoCE Network Fabric switches.

Before you configure the RoCE Network Fabric switches in the new rack (Rn+1), check the configuration of the RoCE Network Fabric switches in the existing racks (R1, R2, ..., Rn). You must do this to ensure that every switch uses a unique loopback octet. The loopback octet is the last octet of the switch loopback IP address.

- a. Connect to an existing RoCE Network Fabric leaf switches and determine the loopback octet for the switch.

Use the command shown in the following example.

```
rack1sw-rocea0# show interface loopback 1 | grep Address
```

```
Internet Address is 192.128.10.101/32
```

In the example, the loopback octet is 101.

- b. Determine the loopback octet for every other leaf switch.

Use the command shown in the following example.

```
rack1sw-rocea0# show nve peers
```

Interface	Peer-IP	State	LearnType
nve1	192.128.10.102	Up	CP
nve1	192.128.10.103	Up	CP
nve1	192.128.10.104	Up	CP
nve1	192.128.10.105	Up	CP
nve1	192.128.10.106	Up	CP
nve1	192.128.10.107	Up	CP
nve1	192.128.10.108	Up	CP

In the example, the output shows seven other leaf switches having loopback octet values from 102 to 108. This output is consistent with an existing system containing four racks.

- c. Determine the loopback octet for every spine switch.

Use the command shown in the following example.

```
rack1sw-rocea0# show bgp l2vpn evpn summary | egrep -v
"BGP|Idle|I|Neighbor|memory"
```

```
192.128.10.201 4 65502 9161 581 75716 0 0 08:53:23 3687
192.128.10.202 4 65502 9160 582 75716 0 0 08:34:20 3687
192.128.10.203 4 65502 9162 582 75716 0 0 08:41:22 3687
192.128.10.204 4 65502 9163 582 75716 0 0 08:50:27 3687
```

In the example, the output shows four spine switches having loopback octet values from 201 to 204. This output is also consistent with an existing system containing four racks.

- d. Validate the configuration of the existing RoCE Network Fabric switches.

Check the information gathered from the existing RoCE Network Fabric switches to ensure that every switch uses a unique loopback octet value and that all the values are as expected.

Verify that the information gathered from the existing RoCE Network Fabric switches conforms to the following conventions:

- On the leaf switches, the overall range of loopback octet values should start with 101 and increase incrementally (by 1) for each leaf switch.

According to the best-practice convention, the loopback octet value for each leaf switch should be configured as follows:

- 101 - Rack 1 lower leaf switch (R1LL)
- 102 - Rack 1 upper leaf switch (R1UL)
- 103 - Rack 2 lower leaf switch (R2LL)
- 104 - Rack 2 upper leaf switch (R2UL)

- 105 - Rack 3 lower leaf switch (R3LL)
- 106 - Rack 3 upper leaf switch (R3UL), and so on.
- On the spine switches, the range of loopback octet values should start with 201 and increase incrementally (by 1) for each spine switch.

According to the best-practice convention, the loopback octet value for each spine switch should be configured as follows:

- 201 - Rack 1 spine switch (R1SS)
- 202 - Rack 2 spine switch (R2SS)
- 203 - Rack 3 spine switch (R3SS)
- 204 - Rack 4 spine switch (R4SS), and so on.

Caution:

If the switches in the existing racks (R1, R2, ..., Rn) don't conform to the above conventions, then you must take special care to assign unique loopback octet values to the switches in the new rack (Rn+1) as part of applying their golden configuration settings (in the next step).

If multiple switches use the same loopback octet, the RoCE Network Fabric cannot function correctly, resulting in a system outage.

5. Apply the golden configuration settings on the RoCE Network Fabric switches in the new rack (Rn+1).

Combine the information about the existing RoCE Network Fabric switches you gathered in the previous step and the procedure described in Applying Golden Configuration Settings on RoCE Network Fabric Switches (in *Oracle Exadata Database Machine Maintenance Guide*).

Caution:

Take care when performing this step, as misconfiguration of the RoCE Network Fabric will likely cause a system outage.

For example, every switch in a multi-rack configuration must have a unique loopback octet. If multiple switches use the same loopback octet, the RoCE Network Fabric cannot function correctly, resulting in a system outage.

6. Enable the leaf switch server ports on the RoCE Network Fabric leaf switches in the new rack (Rn+1).

The leaf switch server ports may be disabled as a consequence of applying the multi-rack golden configuration settings in the previous step.

To ensure that the leaf switch server ports are enabled, log in to each of the leaf switches in the new rack and run the following commands on each switch:

```
rack5sw-rocea0# config term
rack5sw-rocea0# int eth1/8-30
```

```
rack5sw-rocea0# no shut
rack5sw-rocea0# copy running-config startup-config
```

7. Perform the physical cabling of the switches in the new rack (R_{n+1}).

Caution:

Cabling within a live network must be done carefully in order to avoid potentially serious disruptions.

- a. Remove the eight existing inter-switch connections (ports 4, 5, 6, 7 and 30, 31, 32, 33) between each leaf switch in the new rack (R_{n+1}).
- b. Cable the leaf switches in the new rack according to the applicable cabling table.
For example, if you are adding a 5th rack to a system using Exadata X9M (or later model) racks, then use "Table 22-17 Leaf Switch Connections for the Fifth Rack in a Five-Rack System".
8. Add the new rack to the switches in the existing racks (R_1 to R_n).
 - a. For an existing rack (R_x), cable the lower leaf switch R_xLL according to the applicable cabling table.
 - b. For the same rack, cable the upper leaf switch R_xUL according to the applicable cabling table.
 - c. Repeat these steps for each existing rack, R_1 to R_n .
9. Confirm each switch is available and connected.

For each switch in racks R_1 , R_2 , ..., R_n , R_{n+1} , confirm the output for the switch `show interface status` command shows `connected` and `100G`.

When run from a spine switch, the output should be similar to the following:

```
rack1sw-roces0# show interface status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
mgmt0	--	connected	routed	full	1000	--
...						
Eth1/5	RouterPort5	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/6	RouterPort6	connected	routed	full	100G	
QSFP-100G-SR4						
Eth1/7	RouterPort7	connected	routed	full	100G	
QSFP-100G-CR4						
Eth1/8	RouterPort8	connected	routed	full	100G	
QSFP-100G-SR4						
Eth1/9	RouterPort9	connected	routed	full	100G	
QSFP-100G-CR4						

```

Eth1/10      RouterPort10    connected routed    full    100G
QSFP-100G-SR4
Eth1/11      RouterPort11    connected routed    full    100G
QSFP-100G-CR4
Eth1/12      RouterPort12    connected routed    full    100G
QSFP-100G-SR4
Eth1/13      RouterPort13    connected routed    full    100G
QSFP-100G-CR4
Eth1/14      RouterPort14    connected routed    full    100G
QSFP-100G-SR4
Eth1/15      RouterPort15    connected routed    full    100G
QSFP-100G-CR4
Eth1/16      RouterPort16    connected routed    full    100G
QSFP-100G-SR4
Eth1/17      RouterPort17    connected routed    full    100G
QSFP-100G-CR4
Eth1/18      RouterPort18    connected routed    full    100G
QSFP-100G-SR4
Eth1/19      RouterPort19    connected routed    full    100G
QSFP-100G-CR4
Eth1/20      RouterPort20    connected routed    full    100G
QSFP-100G-SR4
Eth1/21      RouterPort21    xcvrAbsen    routed    full    100G
--
...

```

When run from a leaf switch, the output should be similar to the following:

```
rack1sw-rocea0# show interface status
```

```

-----
-----
Port          Name                Status    Vlan    Duplex  Speed  Type
-----
-----
mgmt0         --                  connected routed    full    1000   --
-----
-----
Port          Name                Status    Vlan    Duplex  Speed  Type
-----
-----
...
Eth1/4        RouterPort1         connected routed    full    100G
QSFP-100G-CR4
Eth1/5        RouterPort2         connected routed    full    100G
QSFP-100G-CR4
Eth1/6        RouterPort3         connected routed    full    100G
QSFP-100G-CR4
Eth1/7        RouterPort4         connected routed    full    100G
QSFP-100G-CR4
Eth1/8        celadm14            connected 3888     full    100G
QSFP-100G-CR4
...
Eth1/29       celadm01            connected 3888     full    100G
QSFP-100G-CR4
Eth1/30       RouterPort5         connected routed    full    100G

```



```

QSFP-100G-SR4
Eth1/31      RouterPort6      connected routed    full    100G
QSFP-100G-SR4
Eth1/32      RouterPort7      connected routed    full    100G
QSFP-100G-SR4
Eth1/33      RouterPort8      connected routed    full    100G
QSFP-100G-SR4
...

```

10. Check the neighbor discovery for every switch in racks R1, R2, ..., Rn, Rn+1.

Log in to each switch and use the `show lldp neighbors` command. Make sure that all switches are visible and check the switch ports assignment (leaf switches: ports Eth1/4 - Eth1/7, Eth1/30 - Eth1/33; spine switches: ports Eth1/5 - Eth1/20) against the applicable cabling tables.

Each spine switch should see all the leaf switches in each rack, but not the other spine switches. The output for a spine switch should be similar to the following:

 Note:

The interfaces in the rightmost output column (for example, Ethernet1/5) are different for each switch based on the applicable cabling tables.

```

rack1sw-roces0# show lldp neighbors | grep roce
rack1sw-roceb0 Eth1/5 120 BR Ethernet1/5
rack2sw-roceb0 Eth1/6 120 BR Ethernet1/5
rack1sw-roceb0 Eth1/7 120 BR Ethernet1/7
rack2sw-roceb0 Eth1/8 120 BR Ethernet1/7
rack1sw-roceb0 Eth1/9 120 BR Ethernet1/4
rack2sw-roceb0 Eth1/10 120 BR Ethernet1/4
rack3sw-roceb0 Eth1/11 120 BR Ethernet1/5
rack3sw-roceb0 Eth1/12 120 BR Ethernet1/7
rack1sw-rocea0 Eth1/13 120 BR Ethernet1/5
rack2sw-rocea0 Eth1/14 120 BR Ethernet1/5
rack1sw-rocea0 Eth1/15 120 BR Ethernet1/7
rack2sw-rocea0 Eth1/16 120 BR Ethernet1/7
rack3sw-rocea0 Eth1/17 120 BR Ethernet1/5
rack2sw-rocea0 Eth1/18 120 BR Ethernet1/4
rack3sw-rocea0 Eth1/19 120 BR Ethernet1/7
rack3sw-rocea0 Eth1/20 120 BR Ethernet1/4
...

```

Each leaf switch should see the spine switch in every rack, but not the other leaf switches. The output for a leaf switch should be similar to the following:

 Note:

The interfaces in the rightmost output column (for example, Ethernet1/13) are different for each switch based on the applicable cabling tables.

```
rack1sw-rocea0# show lldp neighbors | grep roce
rack3sw-roces0 Eth1/4 120 BR Ethernet1/13
rack1sw-roces0 Eth1/5 120 BR Ethernet1/13
rack3sw-roces0 Eth1/6 120 BR Ethernet1/15
rack1sw-roces0 Eth1/7 120 BR Ethernet1/15
rack2sw-roces0 Eth1/30 120 BR Ethernet1/17
rack2sw-roces0 Eth1/31 120 BR Ethernet1/13
rack3sw-roces0 Eth1/32 120 BR Ethernet1/17
rack2sw-roces0 Eth1/33 120 BR Ethernet1/15
...
```

11. Power on all the servers in the new rack ($Rn+1$).
12. For each rack, confirm the multi-rack cabling by running the `verify_roce_cables.py` script.

Refer to My Oracle Support Doc ID 2587717.1 for download and usage instructions.

Check the output of the `verify_roce_cables.py` script against the applicable cabling tables. Also, check that output in the `CABLE OK?` columns contains the OK status.

When running the script, two input files are used, one for nodes and one for switches. Each file should contain the servers or switches on separate lines. Use fully qualified domain names or IP addresses for each server and switch.

The following output is a partial example of the command results:

```
# ./verify_roce_cables.py -n nodes.rack1 -s switches.rack1
SWITCH PORT (EXPECTED PEER)  LEAF-1 (rack1sw-rocea0)      : CABLE OK?  LEAF-2 (rack1sw-roceb0)
: CABLE OK?
-----
Eth1/4 (ISL peer switch)    : rack1sw-roces0 Ethernet1/17 : OK          rack1sw-roces0
Ethernet1/9 : OK
Eth1/5 (ISL peer switch)    : rack1sw-roces0 Ethernet1/13 : OK          rack1sw-roces0
Ethernet1/5 : OK
Eth1/6 (ISL peer switch)    : rack1sw-roces0 Ethernet1/19 : OK          rack1sw-roces0
Ethernet1/11: OK
Eth1/7 (ISL peer switch)    : rack1sw-roces0 Ethernet1/15 : OK          rack1sw-roces0
Ethernet1/7 : OK
Eth1/12 (celadm10)          : rack1celadm10 port-1        : OK          rack1celadm10
port-2 : OK
Eth1/13 (celadm09)          : rack1celadm09 port-1        : OK          rack1celadm09
port-2 : OK
Eth1/14 (celadm08)          : rack1celadm08 port-1        : OK          rack1celadm08
port-2 : OK
...
Eth1/15 (adm08)             : rack1dbadm08 port-1         : OK          rack1dbadm08
port-2 : OK
Eth1/16 (adm07)             : rack1dbadm07 port-1         : OK          rack1dbadm07
port-2 : OK
```

```

Eth1/17 (adm06)          : rack1dbadm06 port-1          : OK          rack1dbadm06
port-2                  : OK
...
Eth1/30 (ISL peer switch) : rack2sw-roces0 Ethernet1/17 : OK          rack2sw-roces0
Ethernet1/9 : OK
Eth1/31 (ISL peer switch) : rack2sw-roces0 Ethernet1/13 : OK          rack2sw-roces0
Ethernet1/5 : OK
Eth1/32 (ISL peer switch) : rack2sw-roces0 Ethernet1/19 : OK          rack2sw-roces0
Ethernet1/11: OK
Eth1/33 (ISL peer switch) : rack2sw-roces0 Ethernet1/15 : OK          rack2sw-roces0
Ethernet1/7 : OK

```

13. Verify the RoCE Network Fabric operation by using the `infinicheck` command.

Use the following recommended command sequence. In each command, `hosts.lst` contains a list of database server RoCE Network Fabric IP addresses (2 RoCE Network Fabric IP addresses for each database server), and `cells.lst` contains a list of RoCE Network Fabric IP addresses for the storage servers (2 RoCE Network Fabric IP addresses for each storage server).

- Use `infinicheck` with the `-z` option to clear the files that were created during the last run of the `infinicheck` command. For example:

```
# /opt/oracle.SupportTools/ibdiagtools/infinicheck -g hosts.lst -c
cells.lst -z
```

- Use `infinicheck` with the `-s` option to set up user equivalence for password-less SSH across the RoCE Network Fabric. For example:

```
# /opt/oracle.SupportTools/ibdiagtools/infinicheck -g hosts.lst -c
cells.lst -s
```

- Finally, verify the RoCE Network Fabric operation by using `infinicheck` with the `-b` option, which is recommended on newly imaged machines where it is acceptable to suppress the `cellip.ora` and `cellinit.ora` configuration checks. For example:

```
# /opt/oracle.SupportTools/ibdiagtools/infinicheck -g hosts.lst -c
cells.lst -b
```

```

INFINICHECK
    [Network Connectivity, Configuration and Performance]

    ##### FABRIC TYPE TESTS #####
System type identified: RoCE
Verifying User Equivalence of user=root from all DBs to all CELLS.
    ##### RoCE CONFIGURATION TESTS #####
    Checking for presence of RoCE devices on all DBs and CELLS
[SUCCESS].... RoCE devices on all DBs and CELLS look good
    Checking for RoCE Policy Routing settings on all DBs and CELLS
[SUCCESS].... RoCE Policy Routing settings look good
    Checking for RoCE DSCP ToS mapping on all DBs and CELLS
[SUCCESS].... RoCE DSCP ToS settings look good
    Checking for RoCE PFC settings and DSCP mapping on all DBs and
CELLs
[SUCCESS].... RoCE PFC and DSCP settings look good
    Checking for RoCE interface MTU settings. Expected value : 2300

```

```
[SUCCESS].... RoCE interface MTU settings look good
    Verifying switch advertised DSCP on all DBs and CELLS ports ( )
[SUCCESS].... Advertised DSCP settings from RoCE switch looks good
#### CONNECTIVITY TESTS ####
[COMPUTE NODES -> STORAGE CELLS]
    (60 seconds approx.)
    (Will walk through QoS values: 0-6) [SUCCESS].....Results OK
[SUCCESS]..... All can talk to all storage cells
[COMPUTE NODES -> COMPUTE NODES]
...
```

Cabling Several InfiniBand Network Fabric Racks Together

Use this procedure to add another rack to an existing multi-rack system with InfiniBand Network Fabric.

This procedure is for systems with InfiniBand Network Fabric (X8, or earlier).

1. Set the priority of the current active Subnet Manager Master to 10 on a spine switch.
 - a. Log in to any InfiniBand switch on the active system.
 - b. Use the `getmaster` command to determine that the Subnet Manager Master is running on the spine switch.

The following example shows that the Subnet Manager Master is running on the spine switch `dm01sw-ib1`.

```
# getmaster
20100701 11:46:38 OpenSM Master on Switch : 0x0021283a8516a0a0 ports 36
Sun DCS 36
QDR switch dm01sw-ib1.example.com enhanced port 0 lid 1 lmc 0
```

If the Subnet Manager Master is not running on a spine switch, then perform the following steps:

- i. Use the `getmaster` command to identify the current location of the Subnet Manager Master.
- ii. Log in as the `root` user on the leaf switch that is the Subnet Manager Master.
- iii. Disable Subnet Manager on the switch. The Subnet Manager Master relocates to another switch.
- iv. Use the `getmaster` command to identify the current location of the Subnet Manager Master. If a spine switch is not the Subnet Manager Master, then repeat steps 1.b.ii and 1.b.iii until a spine switch is the Subnet Manager Master.
- v. Enable Subnet Manager on the leaf switches that were disabled during this procedure.
- c. Log in to the Subnet Manager Master 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 to ensure that 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 (R_{n+1}).
4. Cable the leaf switch in the new rack according to the appropriate table in [Multi-Rack Cabling Tables for ZDLRA Rack Models with InfiniBand Network Fabric \(X2 to X8\)](#).

For example, if rack R_{n+1} is R4, then use [Table 24-9](#).

▲ Caution:

Cabling within a live network must be done carefully in order to avoid potentially serious disruptions.

The cabling table that you use for your new InfiniBand topology tells you how to connect ports on the leaf switches to ports on spine switches in order to connect the racks. Some of these ports on the spine switches might be already in use to support the existing InfiniBand topology. In these cases, connect only the cable on the leaf switch in the new rack and stop there for now. Make note of which cables you were not able to terminate.

Do not unplug any cables on the spine switch in the existing rack at this point. Step 5 describes how to re-cable the leaf switches on the existing racks (one leaf switch after the other - while the leaf switch being re-cabled will be powered off), which will free up these currently in-use ports. At that point, you can connect the other end of the cable from the leaf switch in the new rack to the spine switch in the existing rack as indicated in the table.

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 causes all servers in the rack to fail over their InfiniBand traffic to Rx IB3.

- b. Cable leaf switch Rx IB2 according to [Multi-Rack Cabling Tables for ZDLRA Rack Models with InfiniBand Network Fabric \(X2 to X8\)](#).

- c. Power on leaf switch Rx IB2.

- d. Wait at least three minutes for Rx IB2 to become completely operational.

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 .

- e. Power off leaf switch Rx IB3.

This causes all servers in the rack to fail over their InfiniBand traffic to Rx IB2.

- f. Cable leaf switch Rx IB3 according to [Multi-Rack Cabling Tables for ZDLRA Rack Models with InfiniBand Network Fabric \(X2 to X8\)](#).

- g. Power on leaf switch Rx IB3.

- h. Wait at least three minutes for Rx IB3 to become completely operational.

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 .

All racks should now be rewired according to [Multi-Rack Cabling Tables for ZDLRA Rack Models with InfiniBand Network Fabric \(X2 to X8\)](#).

6. Power on all the InfiniBand switches in the new rack.
7. Wait three minutes for the switches to become completely operational.
To check the switch, log in to the switch and run the `ibswitches` command. The output should show $(n+1)*3$ switches for IB1, IB2, and IB3 in racks R1, R2, ... R $n+1$.
8. Ensure that the Subnet Manager Master is running on R1 IB1 by running the `getmaster` command from any switch.
9. Power on all servers in the new rack (R $n+1$).
10. Log in to spine switch R1 IB1, and lower its priority to 8.
 - a. Use the `disable-sm` command to stop Subnet Manager.
 - b. Use the `set-sm-priority 8` command to set the priority to 8.
 - c. Use the `enable-sm` command to restart Subnet Manager.
11. Ensure that the Subnet Manager Master is running on one of the spine switches using the `getmaster` command from any switch.
12. Ensure that the Subnet Manager is running on every spine switch by entering the following command from any switch:

```
ibdiagnet -r
```

Each spine switch should show as running in the `Summary Fabric SM-state-priority` section of the output. If a spine switch is not running, then log in to the switch and enable the Subnet Manager using the `enable-sm` command.

13. If there are now four or more racks, then log in to the leaf switches in each rack and disable Subnet Manager using the `disable-sm` command.

Multi-Rack Cabling Tables for Recovery Appliance RA21 and Later Models

This section contains multi-rack cabling tables for Recovery Appliance RA21 and later models, which use RoCE Network Fabric.

Understanding Multi-Rack Cabling for RA21 and Later Model Racks

Up to 14 racks (RA21 and later models) can be cabled together without external RDMA Network Fabric switches.

To interconnect multiple racks, each rack must contain three RDMA Network Fabric switches.

Two of the switches, known as leaf switches, connect to all of the database servers and storage servers in the rack. The leaf switches are located as follows:

- At rack unit 20 (U20) and 22 (U22) in 2-socket systems (Recovery Appliance X9M and later models).

The third switch, known as the spine switch, interconnects the leaf switches in a multi-rack configuration. The spine switch is located at rack unit 1 (U1). If any rack does not include a spine switch, then order the expansion switch kit.


The procedures in this section assume the racks are adjacent to each other, standard routing in raised floor is used, and there are no obstacles in the raised floor. If these assumptions are not correct for your environment, then longer cables may be required for the connections.



Note:

By default, Oracle Exadata Database Machine racks do not include spare cables or a third RoCE Network Fabric switch. To extend these racks, you must order the required cables and RoCE Network Fabric switch.

The following diagram shows the default RDMA Network Fabric architecture for a single-rack system. Each rack has two leaf switches, with eight connections between the leaf switches. The database servers and storage servers are each connected to both leaf switches. Each server contains a dual-port RDMA Network Fabric card, with port 1 connected to the lower leaf switch and port 2 connected to the upper leaf switch.



To connect up to 14 racks (RA21 and later models) together, use the following general approach:


1. Remove the eight existing inter-switch connections between the leaf switches on each rack.
2. From each leaf switch, evenly distribute 14 connections to the spine switches in all of the interconnected racks.

The 14 connections use the 8 ports that were previously used for the inter-switch connections and 6 additional free ports on each leaf switch.

 Note:

For X9M-8 systems with three database servers and 11 storage servers only, the database servers and storage server require 23 leaf switch ports, which leaves only 13 inter-switch links on each leaf switch. Consequently, these systems are limited to a maximum of 13 interconnected racks.

The resulting RoCE Network Fabric for a typical 2 rack system is illustrated in the following diagram:



As shown in the preceding diagram, every leaf switch has 7 connections to every spine switch. The leaf switches are not directly interconnected with other leaf switches, and the spine switches are not directly interconnected with each other.

As the number of racks increases, the inter-switch connections from every leaf switch are evenly distributed to all of spine switches.

Preparing for Multi-Rack Cabling with RA21 and Later Model Racks

Racks can be added together to increase system capacity and performance. When cabling racks together, note the following:

- The cable lengths shown in this document assume the racks are adjacent to each other. If the racks are not adjacent, or there are obstacles in the raised floor, or if you use overhead cabling, then longer cables may be required. For optical cables, the maximum supported cable length is 100 meters. For copper cables, the maximum supported cable length is 5 meters.
- Oracle recommends that the names for the servers include the rack unit number. This helps identify the server during diagnostics.

- When completing Oracle Exadata Deployment Assistant (OEDA) for the additional rack, you are prompted for SCAN addresses. However, these SCAN addresses are not used because the SCAN address from the original rack are used. Manually remove the new SCAN addresses from the generated installation files.
- The software owner account names and group names, as well as their identifiers, must match the names and identifiers of the original rack.
- If the additional grid disks are used with existing disk groups, then ensure the grid disk sizes for the new rack are the same as the original rack.
- ***For multi-rack configurations containing up to 14 racks, a spine switch must exist in each rack in order to interconnect the RoCE Network Fabric.***

Perform the following tasks before cabling racks together:

1. Determine the number of racks that will be cabled together.
2. Order the parts needed to connect the racks.

To extend Recovery Appliance racks with RoCE Network Fabric, for each rack being added you must order extra cables, transceivers for longer cables, and a RoCE Network Fabric spine switch, if one is required.

When connecting four or more racks, or if you need longer cables for your environment, you must purchase additional 10 meter or 15 meter fiber cables with two QSFP28 SR transceivers to connect each end. The QSFP28 SR transceivers are needed for fiber cables over 5 meters in length.

For multi-rack configurations containing up to 14 racks, the following table outlines the cables needed to interconnect the racks:

Number of Racks	3 Meter Cables Required	5 Meter Cables Required	10 Meter Cables Required	15 Meter Cables Required
2	28	28	0	0
3	28	56	0	0
4	28	71	13	0
5	28	78	34	0
6	28	87	53	0
7	28	88	80	0
8	28	91	105	0
9	28	94	130	0
10	28	95	151	6
11	28	97	168	15
12	28	99	184	25
13	28	100	196	40
14	28	100	204	60

3. Gather the following tools:
 - Pen
 - Note pad
 - Flashlight
 - Labels for cables and rack

- Cable ties (Oracle recommends hook and loop straps)
- 4. Determine a naming method for the rack prefixes. For example, if the original rack has the prefix `dbm01`, then use the prefix `dbm02` for the second rack, the prefix `dbm03` for the third rack, and so on.
- 5. Verify that the racks have unique host names and IP addresses. All servers interconnected in the racks must have unique names and IP addresses.

Server name and IP address conventions may differ in the following cases:

- **Initial installation of all Recovery Appliance Racks:** System address assignments and host names should be complete.
 - **New Recovery Appliance Racks are added to an existing cluster:** The new rack configuration should require unique host names and IP addresses for the new Recovery Appliance. The IP addresses on the same subnet cannot conflict with the existing systems.
 - **Two existing Recovery Appliance Racks are clustered together:** You can assign host names and IP addresses only if Recovery Appliance racks are already assigned unique host names and IP addresses, or the entire cluster must be reconfigured. The machines must be on the same subnet and not have conflicting IP addresses.
6. Ensure the IP addresses for the new servers are in the same subnet, and do not overlap with the currently-installed servers.
 7. Check that the firmware on the original switches is at the same level as the new switches by using the `sh ver` command.

It is highly recommended, though not mandatory, to use the same firmware version on all of the switches. If the firmware is not at the same level, you can apply a firmware patch to bring the switches up to the same firmware level.

The remaining topics in this section provide detailed cabling information for cabling up to 14 racks together.

Two-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect two racks, both of which use RoCE Network Fabric.



Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, for these systems only, ignore the connections to port number 30 on every leaf switch in the following tables. This adjustment leaves only 13 inter-switch links on each leaf switch and only applies to X9M-8 systems with three database servers and 11 storage servers.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a two-rack system.

Table 22-1 Leaf Switch Connections for the First Rack in a Two-Rack System

Switch Pair	Connections	Cable Length
R1UL to R1SS	R1UL-P1 to R1SS-P21 R1UL-P2 to R1SS-P23 R1UL-P3 to R1SS-P25 R1UL-P4 to R1SS-P9 R1UL-P5 to R1SS-P5 R1UL-P6 to R1SS-P11 R1UL-P7 to R1SS-P7	3 meters

Table 22-1 (Cont.) Leaf Switch Connections for the First Rack in a Two-Rack System

Switch Pair	Connections	Cable Length
R1UL to R2SS	R1UL-P30 to R2SS-P9 R1UL-P31 to R2SS-P5 R1UL-P32 to R2SS-P11 R1UL-P33 to R2SS-P7 R1UL-P34 to R2SS-P21 R1UL-P35 to R2SS-P23 R1UL-P36 to R2SS-P25	5 meters
R1LL to R1SS	R1LL-P1 to R1SS-P27 R1LL-P2 to R1SS-P29 R1LL-P3 to R1SS-P31 R1LL-P4 to R1SS-P17 R1LL-P5 to R1SS-P13 R1LL-P6 to R1SS-P19 R1LL-P7 to R1SS-P15	3 meters
R1LL to R2SS	R1LL-P30 to R2SS-P17 R1LL-P31 to R2SS-P13 R1LL-P32 to R2SS-P19 R1LL-P33 to R2SS-P15 R1LL-P34 to R2SS-P27 R1LL-P35 to R2SS-P29 R1LL-P36 to R2SS-P31	5 meters

Table 22-2 Leaf Switch Connections for the Second Rack in a Two-Rack System

Switch Pair	Connections	Cable Length
R2UL to R1SS	R2UL-P1 to R1SS-P22 R2UL-P2 to R1SS-P24 R2UL-P3 to R1SS-P26 R2UL-P4 to R1SS-P10 R2UL-P5 to R1SS-P6 R2UL-P6 to R1SS-P12 R2UL-P7 to R1SS-P8	5 meters
R2UL to R2SS	R2UL-P30 to R2SS-P10 R2UL-P31 to R2SS-P6 R2UL-P32 to R2SS-P12 R2UL-P33 to R2SS-P8 R2UL-P34 to R2SS-P22 R2UL-P35 to R2SS-P24 R2UL-P36 to R2SS-P26	3 meters

Table 22-2 (Cont.) Leaf Switch Connections for the Second Rack in a Two-Rack System

Switch Pair	Connections	Cable Length
R2LL to R1SS	R2LL-P1 to R1SS-P28 R2LL-P2 to R1SS-P30 R2LL-P3 to R1SS-P32 R2LL-P4 to R1SS-P18 R2LL-P5 to R1SS-P14 R2LL-P6 to R1SS-P20 R2LL-P7 to R1SS-P16	5 meters
R2LL to R2SS	R2LL-P30 to R2SS-P18 R2LL-P31 to R2SS-P14 R2LL-P32 to R2SS-P20 R2LL-P33 to R2SS-P16 R2LL-P34 to R2SS-P28 R2LL-P35 to R2SS-P30 R2LL-P36 to R2SS-P32	3 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-3 Two-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R1UL-P7	3 meters
R1SS-P8	R2UL-P7	5 meters
R1SS-P9	R1UL-P4	3 meters
R1SS-P10	R2UL-P4	5 meters
R1SS-P11	R1UL-P6	3 meters
R1SS-P12	R2UL-P6	5 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R1LL-P7	3 meters
R1SS-P16	R2LL-P7	5 meters
R1SS-P17	R1LL-P4	3 meters
R1SS-P18	R2LL-P4	5 meters
R1SS-P19	R1LL-P6	3 meters
R1SS-P20	R2LL-P6	5 meters
R1SS-P21	R1UL-P1	3 meters
R1SS-P22	R2UL-P1	5 meters
R1SS-P23	R1UL-P2	3 meters
R1SS-P24	R2UL-P2	5 meters
R1SS-P25	R1UL-P3	3 meters
R1SS-P26	R2UL-P3	5 meters

Table 22-3 (Cont.) Two-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P27	R1LL-P1	3 meters
R1SS-P28	R2LL-P1	5 meters
R1SS-P29	R1LL-P2	3 meters
R1SS-P30	R2LL-P2	5 meters
R1SS-P31	R1LL-P3	3 meters
R1SS-P32	R2LL-P3	5 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R1UL-P33	5 meters
R2SS-P8	R2UL-P33	3 meters
R2SS-P9	R1UL-P30	5 meters
R2SS-P10	R2UL-P30	3 meters
R2SS-P11	R1UL-P32	5 meters
R2SS-P12	R2UL-P32	3 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R1LL-P33	5 meters
R2SS-P16	R2LL-P33	3 meters
R2SS-P17	R1LL-P30	5 meters
R2SS-P18	R2LL-P30	3 meters
R2SS-P19	R1LL-P32	5 meters
R2SS-P20	R2LL-P32	3 meters
R2SS-P21	R1UL-P34	5 meters
R2SS-P22	R2UL-P34	3 meters
R2SS-P23	R1UL-P35	5 meters
R2SS-P24	R2UL-P35	3 meters
R2SS-P25	R1UL-P36	5 meters
R2SS-P26	R2UL-P36	3 meters
R2SS-P27	R1LL-P34	5 meters
R2SS-P28	R2LL-P34	3 meters
R2SS-P29	R1LL-P35	5 meters
R2SS-P30	R2LL-P35	3 meters
R2SS-P31	R1LL-P36	5 meters
R2SS-P32	R2LL-P36	3 meters

Three-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect three racks using RoCE Network Fabric.



Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, for these systems only, ignore the connections to port number 30 on every leaf switch in the following tables. This adjustment leaves only 13 inter-switch links on each leaf switch and only applies to X9M-8 systems with three database servers and 11 storage servers.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a three-rack system.

Table 22-4 Leaf Switch Connections for the First Rack in a Three-Rack System

Switch Pair	Connections	Cable Length
R1UL to R1SS	R1UL-P2 to R1SS-P21 R1UL-P3 to R1SS-P23 R1UL-P4 to R1SS-P9 R1UL-P5 to R1SS-P5 R1UL-P7 to R1SS-P7	3 meters
R1UL to R2SS	R1UL-P30 to R2SS-P9 R1UL-P31 to R2SS-P5 R1UL-P33 to R2SS-P7 R1UL-P34 to R2SS-P21 R1UL-P35 to R2SS-P23	5 meters

Table 22-4 (Cont.) Leaf Switch Connections for the First Rack in a Three-Rack System

Switch Pair	Connections	Cable Length
R1UL to R3SS	R1UL-P1 to R3SS-P21 R1UL-P6 to R3SS-P5 R1UL-P32 to R3SS-P7 R1UL-P36 to R3SS-P23	5 meters
R1LL to R1SS	R1LL-P2 to R1SS-P27 R1LL-P3 to R1SS-P29 R1LL-P5 to R1SS-P13 R1LL-P7 to R1SS-P15	3 meters
R1LL to R2SS	R1LL-P30 to R2SS-P17 R1LL-P31 to R2SS-P13 R1LL-P33 to R2SS-P15 R1LL-P34 to R2SS-P27 R1LL-P35 to R2SS-P29	5 meters
R1LL to R3SS	R1LL-P1 to R3SS-P27 R1LL-P4 to R3SS-P13 R1LL-P6 to R3SS-P15 R1LL-P32 to R3SS-P17 R1LL-P36 to R3SS-P29	5 meters

Table 22-5 Leaf Switch Connections for the Second Rack in a Three-Rack System

Switch Pair	Connections	Cable Length
R2UL to R1SS	R2UL-P2 to R1SS-P22 R2UL-P3 to R1SS-P24 R2UL-P4 to R1SS-P10 R2UL-P5 to R1SS-P6 R2UL-P7 to R1SS-P8	5 meters
R2UL to R2SS	R2UL-P31 to R2SS-P6 R2UL-P33 to R2SS-P8 R2UL-P34 to R2SS-P22 R2UL-P35 to R2SS-P24	3 meters
R2UL to R3SS	R2UL-P1 to R3SS-P22 R2UL-P6 to R3SS-P6 R2UL-P30 to R3SS-P8 R2UL-P32 to R3SS-P10 R2UL-P36 to R3SS-P24	5 meters
R2LL to R1SS	R2LL-P2 to R1SS-P28 R2LL-P3 to R1SS-P30 R2LL-P4 to R1SS-P18 R2LL-P5 to R1SS-P14 R2LL-P7 to R1SS-P16	5 meters

Table 22-5 (Cont.) Leaf Switch Connections for the Second Rack in a Three-Rack System

Switch Pair	Connections	Cable Length
R2LL to R2SS	R2LL-P30 to R2SS-P18 R2LL-P31 to R2SS-P14 R2LL-P33 to R2SS-P16 R2LL-P34 to R2SS-P28 R2LL-P35 to R2SS-P30	3 meters
R2LL to R3	R2LL-P1 to R3SS-P28 R2LL-P6 to R3SS-P14 R2LL-P32 to R3SS-P16 R2LL-P36 to R3SS-P30	5 meters

Table 22-6 Leaf Switch Connections for the Third Rack in a Three-Rack System

Switch Pair	Connections	Cable Length
R3UL to R1SS	R3UL-P2 to R1SS-P25 R3UL-P3 to R1SS-P31 R3UL-P5 to R1SS-P11 R3UL-P7 to R1SS-P12	5 meters
R3UL to R2SS	R3UL-P4 to R2SS-P10 R3UL-P6 to R2SS-P11 R3UL-P31 to R2SS-P12 R3UL-P34 to R2SS-P25 R3UL-P35 to R2SS-P31	5 meters
R3UL to R3SS	R3UL-P1 to R3SS-P25 R3UL-P30 to R3SS-P11 R3UL-P32 to R3SS-P12 R3UL-P33 to R3SS-P9 R3UL-P36 to R3SS-P31	3 meters
R3LL to R1SS	R3LL-P2 to R1SS-P26 R3LL-P3 to R1SS-P32 R3LL-P4 to R1SS-P20 R3LL-P5 to R1SS-P17 R3LL-P7 to R1SS-P19	5 meters
R3LL to R2SS	R3LL-P6 to R2SS-P19 R3LL-P31 to R2SS-P20 R3LL-P34 to R2SS-P26 R3LL-P35 to R2SS-P32	5 meters
R3LL to R3SS	R3LL-P1 to R3SS-P26 R3LL-P30 to R3SS-P19 R3LL-P32 to R3SS-P20 R3LL-P33 to R3SS-P18 R3LL-P36 to R3SS-P32	3 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-7 Three-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R1UL-P7	3 meters
R1SS-P8	R2UL-P7	5 meters
R1SS-P9	R1UL-P4	3 meters
R1SS-P10	R2UL-P4	5 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R3UL-P7	5 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R1LL-P7	3 meters
R1SS-P16	R2LL-P7	5 meters
R1SS-P17	R3LL-P5	5 meters
R1SS-P18	R2LL-P4	5 meters
R1SS-P19	R3LL-P7	5 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R1UL-P2	3 meters
R1SS-P22	R2UL-P2	5 meters
R1SS-P23	R1UL-P3	3 meters
R1SS-P24	R2UL-P3	5 meters
R1SS-P25	R3UL-P2	5 meters
R1SS-P26	R3LL-P2	5 meters
R1SS-P27	R1LL-P2	3 meters
R1SS-P28	R2LL-P2	5 meters
R1SS-P29	R1LL-P3	3 meters
R1SS-P30	R2LL-P3	5 meters
R1SS-P31	R3UL-P3	5 meters
R1SS-P32	R3LL-P3	5 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R1UL-P33	5 meters
R2SS-P8	R2UL-P33	3 meters
R2SS-P9	R1UL-P30	5 meters
R2SS-P10	R3UL-P4	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R3UL-P31	5 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R1LL-P33	5 meters
R2SS-P16	R2LL-P33	3 meters
R2SS-P17	R1LL-P30	5 meters

Table 22-7 (Cont.) Three-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R2SS-P18	R2LL-P30	3 meters
R2SS-P19	R3LL-P6	5 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R1UL-P34	5 meters
R2SS-P22	R2UL-P34	3 meters
R2SS-P23	R1UL-P35	5 meters
R2SS-P24	R2UL-P35	3 meters
R2SS-P25	R3UL-P34	5 meters
R2SS-P26	R3LL-P34	5 meters
R2SS-P27	R1LL-P34	5 meters
R2SS-P28	R2LL-P34	3 meters
R2SS-P29	R1LL-P35	5 meters
R2SS-P30	R2LL-P35	3 meters
R2SS-P31	R3UL-P35	5 meters
R2SS-P32	R3LL-P35	5 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R1UL-P32	5 meters
R3SS-P8	R2UL-P30	5 meters
R3SS-P9	R3UL-P33	3 meters
R3SS-P10	R2UL-P32	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R3UL-P32	3 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R1LL-P6	5 meters
R3SS-P16	R2LL-P32	5 meters
R3SS-P17	R1LL-P32	5 meters
R3SS-P18	R3LL-P33	3 meters
R3SS-P19	R3LL-P30	3 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R1UL-P1	5 meters
R3SS-P22	R2UL-P1	5 meters
R3SS-P23	R1UL-P36	5 meters
R3SS-P24	R2UL-P36	5 meters
R3SS-P25	R3UL-P1	3 meters
R3SS-P26	R3LL-P1	3 meters
R3SS-P27	R1LL-P1	5 meters
R3SS-P28	R2LL-P1	5 meters
R3SS-P29	R1LL-P36	5 meters
R3SS-P30	R2LL-P36	5 meters
R3SS-P31	R3UL-P36	3 meters

Table 22-7 (Cont.) Three-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R3SS-P32	R3LL-P36	3 meters

Four-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect four racks using RoCE Network Fabric.

Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, for these systems only, ignore the connections to port number 30 on every leaf switch in the following tables. This adjustment leaves only 13 inter-switch links on each leaf switch and only applies to X9M-8 systems with three database servers and 11 storage servers.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a four-rack system.

Table 22-8 Leaf Switch Connections for the First Rack in a Four-Rack System

Switch Pair	Connections	Cable Length
R1UL to R1SS	R1UL-P2 to R1SS-P21 R1UL-P3 to R1SS-P23 R1UL-P5 to R1SS-P5 R1UL-P7 to R1SS-P7	3 meters
R1UL to R2SS	R1UL-P31 to R2SS-P5 R1UL-P33 to R2SS-P7 R1UL-P34 to R2SS-P21 R1UL-P35 to R2SS-P23	5 meters
R1UL to R3SS	R1UL-P1 to R3SS-P21 R1UL-P6 to R3SS-P5 R1UL-P32 to R3SS-P7	5 meters
R1UL to R4SS	R1UL-P4 to R4SS-P5 R1UL-P30 to R4SS-P7 R1UL-P36 to R4SS-P21	10 meters
R1LL to R1SS	R1LL-P3 to R1SS-P27 R1LL-P5 to R1SS-P13 R1LL-P7 to R1SS-P15	3 meters
R1LL to R2SS	R1LL-P31 to R2SS-P13 R1LL-P33 to R2SS-P15 R1LL-P34 to R2SS-P27 R1LL-P35 to R2SS-P29	5 meters
R1LL to R3SS	R1LL-P1 to R3SS-P27 R1LL-P2 to R3SS-P29 R1LL-P4 to R3SS-P13 R1LL-P6 to R3SS-P15	5 meters
R1LL to R4SS	R1LL-P30 to R4SS-P13 R1LL-P32 to R4SS-P15 R1LL-P36 to R4SS-P27	10 meters

Table 22-9 Leaf Switch Connections for the Second Rack in a Four-Rack System

Switch Pair	Connections	Cable Length
R2UL to R1SS	R2UL-P3 to R1SS-P22 R2UL-P5 to R1SS-P6 R2UL-P7 to R1SS-P8	5 meters
R2UL to R2SS	R2UL-P31 to R2SS-P6 R2UL-P33 to R2SS-P8 R2UL-P35 to R2SS-P22	3 meters
R2UL to R3SS	R2UL-P1 to R3SS-P22 R2UL-P2 to R3SS-P24 R2UL-P6 to R3SS-P6 R2UL-P30 to R3SS-P8	5 meters

Table 22-9 (Cont.) Leaf Switch Connections for the Second Rack in a Four-Rack System

Switch Pair	Connections	Cable Length
R2UL to R4SS	R2UL-P4 to R4SS-P6 R2UL-P32 to R4SS-P8 R2UL-P34 to R4SS-P22 R2UL-P36 to R4SS-P24	5 meters
R2LL to R1SS	R2LL-P1 to R1SS-P24 R2LL-P3 to R1SS-P28 R2LL-P5 to R1SS-P14 R2LL-P7 to R1SS-P16	5 meters
R2LL to R2SS	R2LL-P31 to R2SS-P14 R2LL-P33 to R2SS-P16 R2LL-P35 to R2SS-P28	3 meters
R2LL to R3SS	R2LL-P6 to R3SS-P14 R2LL-P32 to R3SS-P16 R2LL-P34 to R3SS-P28	5 meters
R2LL to R4SS	R2LL-P2 to R4SS-P28 R2LL-P4 to R4SS-P14 R2LL-P30 to R4SS-P16 R2LL-P36 to R4SS-P30	5 meters

Table 22-10 Leaf Switch Connections for the Third Rack in a Four-Rack System

Switch Pair	Connections	Cable Length
R3UL to R1SS	R3UL-P1 to R1SS-P25 R3UL-P3 to R1SS-P29 R3UL-P5 to R1SS-P11 R3UL-P7 to R1SS-P12	5 meters
R3UL to R2SS	R3UL-P2 to R2SS-P25 R3UL-P6 to R2SS-P11 R3UL-P31 to R2SS-P12 R3UL-P35 to R2SS-P31	5 meters
R3UL to R3SS	R3UL-P30 to R3SS-P11 R3UL-P32 to R3SS-P12 R3UL-P34 to R3SS-P25	3 meters
R3UL to R4SS	R3UL-P4 to R4SS-P11 R3UL-P33 to R4SS-P12 R3UL-P36 to R4SS-P25	5 meters
R3LL to R1SS	R3LL-P3 to R1SS-P26 R3LL-P4 to R1SS-P20 R3LL-P7 to R1SS-P19	5 meters

Table 22-10 (Cont.) Leaf Switch Connections for the Third Rack in a Four-Rack System

Switch Pair	Connections	Cable Length
R3LL to R2SS	R3LL-P2 to R2SS-P24 R3LL-P6 to R2SS-P19 R3LL-P31 to R2SS-P20 R3LL-P35 to R2SS-P26	5 meters
R3LL to R3SS	R3LL-P1 to R3SS-P26 R3LL-P30 to R3SS-P19 R3LL-P32 to R3SS-P20 R3LL-P34 to R3SS-P23	3 meters
R3LL to R4SS	R3LL-P5 to R4SS-P19 R3LL-P33 to R4SS-P20 R3LL-P36 to R4SS-P26	5 meters

Table 22-11 Leaf Switch Connections for the Fourth Rack in a Four-Rack System

Switch Pair	Connections	Cable Length
R4UL to R1SS	R4UL-P3 to R1SS-P31 R4UL-P5 to R1SS-P9 R4UL-P7 to R1SS-P10	10 meters
R4UL to R2SS	R4UL-P2 to R2SS-P30 R4UL-P4 to R2SS-P9 R4UL-P6 to R2SS-P10	5 meters
R4UL to R3SS	R4UL-P31 to R3SS-P9 R4UL-P33 to R3SS-P10 R4UL-P34 to R3SS-P30 R4UL-P35 to R3SS-P31	5 meters
R4UL to R4SS	R4UL-P1 to R4SS-P23 R4UL-P30 to R4SS-P9 R4UL-P32 to R4SS-P10 R4UL-P36 to R4SS-P29	3 meters
R4LL to R1SS	R4LL-P1 to R1SS-P32 R4LL-P3 to R1SS-P30 R4LL-P5 to R1SS-P17 R4LL-P7 to R1SS-P18	10 meters
R4LL to R2SS	R4LL-P2 to R2SS-P32 R4LL-P4 to R2SS-P17 R4LL-P6 to R2SS-P18	5 meters
R4LL to R3SS	R4LL-P31 to R3SS-P17 R4LL-P33 to R3SS-P18 R4LL-P34 to R3SS-P32	5 meters

Table 22-11 (Cont.) Leaf Switch Connections for the Fourth Rack in a Four-Rack System

Switch Pair	Connections	Cable Length
R4LL to R4SS	R4LL-P30 to R4SS-P17 R4LL-P32 to R4SS-P18 R4LL-P35 to R4SS-P31 R4LL-P36 to R4SS-P32	3 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-12 Four-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R1UL-P7	3 meters
R1SS-P8	R2UL-P7	5 meters
R1SS-P9	R4UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R3UL-P7	5 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R1LL-P7	3 meters
R1SS-P16	R2LL-P7	5 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R4LL-P7	10 meters
R1SS-P19	R3LL-P7	5 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R1UL-P2	3 meters
R1SS-P22	R2UL-P3	5 meters
R1SS-P23	R1UL-P3	3 meters
R1SS-P24	R2LL-P1	5 meters
R1SS-P25	R3UL-P1	5 meters
R1SS-P26	R3LL-P3	5 meters
R1SS-P27	R1LL-P3	3 meters
R1SS-P28	R2LL-P3	5 meters
R1SS-P29	R3UL-P3	5 meters
R1SS-P30	R4LL-P3	10 meters
R1SS-P31	R4UL-P3	10 meters
R1SS-P32	R4LL-P1	10 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R1UL-P33	5 meters

Table 22-12 (Cont.) Four-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R2SS-P8	R2UL-P33	3 meters
R2SS-P9	R4UL-P4	5 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R3UL-P31	5 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R1LL-P33	5 meters
R2SS-P16	R2LL-P33	3 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R4LL-P6	5 meters
R2SS-P19	R3LL-P6	5 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R1UL-P34	5 meters
R2SS-P22	R2UL-P35	3 meters
R2SS-P23	R1UL-P35	5 meters
R2SS-P24	R3LL-P2	5 meters
R2SS-P25	R3UL-P2	5 meters
R2SS-P26	R3LL-P35	5 meters
R2SS-P27	R1LL-P34	5 meters
R2SS-P28	R2LL-P35	3 meters
R2SS-P29	R1LL-P35	5 meters
R2SS-P30	R4UL-P2	5 meters
R2SS-P31	R3UL-P35	5 meters
R2SS-P32	R4LL-P2	5 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R1UL-P32	5 meters
R3SS-P8	R2UL-P30	5 meters
R3SS-P9	R4UL-P31	5 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R3UL-P32	3 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R1LL-P6	5 meters
R3SS-P16	R2LL-P32	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R4LL-P33	5 meters
R3SS-P19	R3LL-P30	3 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R1UL-P1	5 meters

Table 22-12 (Cont.) Four-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R3SS-P22	R2UL-P1	5 meters
R3SS-P23	R3LL-P34	3 meters
R3SS-P24	R2UL-P2	5 meters
R3SS-P25	R3UL-P34	3 meters
R3SS-P26	R3LL-P1	3 meters
R3SS-P27	R1LL-P1	5 meters
R3SS-P28	R2LL-P34	5 meters
R3SS-P29	R1LL-P2	5 meters
R3SS-P30	R4UL-P34	5 meters
R3SS-P31	R4UL-P35	5 meters
R3SS-P32	R4LL-P34	5 meters
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R1UL-P30	10 meters
R4SS-P8	R2UL-P32	5 meters
R4SS-P9	R4UL-P30	3 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R3UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R1LL-P32	10 meters
R4SS-P16	R2LL-P30	5 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R4LL-P32	3 meters
R4SS-P19	R3LL-P5	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R1UL-P36	10 meters
R4SS-P22	R2UL-P34	5 meters
R4SS-P23	R4UL-P1	3 meters
R4SS-P24	R2UL-P36	5 meters
R4SS-P25	R3UL-P36	5 meters
R4SS-P26	R3LL-P36	5 meters
R4SS-P27	R1LL-P36	10 meters
R4SS-P28	R2LL-P2	5 meters
R4SS-P29	R4UL-P36	3 meters
R4SS-P30	R2LL-P36	5 meters
R4SS-P31	R4LL-P35	3 meters
R4SS-P32	R4LL-P36	3 meters

Five-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect five racks using RoCE Network Fabric.



Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, for these systems only, ignore the connections to port number 30 on every leaf switch in the following tables. This adjustment leaves only 13 inter-switch links on each leaf switch and only applies to X9M-8 systems with three database servers and 11 storage servers.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a five-rack system.

Table 22-13 Leaf Switch Connections for the First Rack in a Five-Rack System

Switch Pair	Connections	Cable Length
R1UL to R1SS	R1UL-P3 to R1SS-P21 R1UL-P5 to R1SS-P5 R1UL-P7 to R1SS-P7	3 meters

Table 22-13 (Cont.) Leaf Switch Connections for the First Rack in a Five-Rack System

Switch Pair	Connections	Cable Length
R1UL to R2SS	R1UL-P31 to R2SS-P5 R1UL-P33 to R2SS-P7 R1UL-P35 to R2SS-P21	5 meters
R1UL to R3SS	R1UL-P1 to R3SS-P21 R1UL-P6 to R3SS-P5 R1UL-P32 to R3SS-P7	5 meters
R1UL to R4SS	R1UL-P2 to R4SS-P21 R1UL-P4 to R4SS-P5 R1UL-P34 to R4SS-P23	10 meters
R1UL to R5SS	R1UL-P30 to R5SS-P5 R1UL-P36 to R5SS-P21	10 meters
R1LL to R1SS	R1LL-P3 to R1SS-P27 R1LL-P5 to R1SS-P13	3 meters
R1LL to R2SS	R1LL-P31 to R2SS-P13 R1LL-P33 to R2SS-P15 R1LL-P35 to R2SS-P27	5 meters
R1LL to R3SS	R1LL-P2 to R3SS-P27 R1LL-P4 to R3SS-P13 R1LL-P6 to R3SS-P15	5 meters
R1LL to R4SS	R1LL-P30 to R4SS-P13 R1LL-P32 to R4SS-P15 R1LL-P34 to R4SS-P27	10 meters
R1LL to R5SS	R1LL-P1 to R5SS-P25 R1LL-P7 to R5SS-P13 R1LL-P36 to R5SS-P27	10 meters

Table 22-14 Leaf Switch Connections for the Second Rack in a Five-Rack System

Switch Pair	Connections	Cable Length
R2UL to R1SS	R2UL-P1 to R1SS-P24 R2UL-P3 to R1SS-P22 R2UL-P5 to R1SS-P6	5 meters
R2UL to R2SS	R2UL-P31 to R2SS-P6 R2UL-P35 to R2SS-P22	3 meters
R2UL to R3SS	R2UL-P2 to R3SS-P22 R2UL-P6 to R3SS-P6 R2UL-P30 to R3SS-P8	5 meters
R2UL to R4SS	R2UL-P4 to R4SS-P6 R2UL-P32 to R4SS-P8 R2UL-P34 to R4SS-P22	5 meters

Table 22-14 (Cont.) Leaf Switch Connections for the Second Rack in a Five-Rack System

Switch Pair	Connections	Cable Length
R2UL to R5SS	R2UL-P7 to R5SS-P6 R2UL-P33 to R5SS-P8 R2UL-P36 to R5SS-P22	10 meters
R2LL to R1SS	R2LL-P3 to R1SS-P28 R2LL-P5 to R1SS-P14 R2LL-P7 to R1SS-P16	5 meters
R2LL to R2SS	R2LL-P31 to R2SS-P14 R2LL-P34 to R2SS-P24 R2LL-P35 to R2SS-P28	3 meters
R2LL to R3SS	R2LL-P1 to R3SS-P28 R2LL-P6 to R3SS-P14	5 meters
R2LL to R4SS	R2LL-P2 to R4SS-P28 R2LL-P4 to R4SS-P14 R2LL-P30 to R4SS-P16	5 meters
R2LL to R5SS	R2LL-P32 to R5SS-P16 R2LL-P33 to R5SS-P14 R2LL-P36 to R5SS-P28	10 meters

Table 22-15 Leaf Switch Connections for the Third Rack in a Five-Rack System

Switch Pair	Connections	Cable Length
R3UL to R1SS	R3UL-P3 to R1SS-P25 R3UL-P5 to R1SS-P11 R3UL-P7 to R1SS-P12	5 meters
R3UL to R2SS	R3UL-P2 to R2SS-P25 R3UL-P6 to R2SS-P11 R3UL-P31 to R2SS-P12	5 meters
R3UL to R3SS	R3UL-P30 to R3SS-P11 R3UL-P34 to R3SS-P25 R3UL-P36 to R3SS-P24	3 meters
R3UL to R4SS	R3UL-P1 to R4SS-P25 R3UL-P4 to R4SS-P11	5 meters
R3UL to R5SS	R3UL-P32 to R5SS-P11 R3UL-P33 to R5SS-P12 R3UL-P35 to R5SS-P24	5 meters
R3LL to R1SS	R3LL-P3 to R1SS-P26 R3LL-P4 to R1SS-P20 R3LL-P7 to R1SS-P19	5 meters
R3LL to R2SS	R3LL-P2 to R2SS-P26 R3LL-P6 to R2SS-P19 R3LL-P31 to R2SS-P20	5 meters

Table 22-15 (Cont.) Leaf Switch Connections for the Third Rack in a Five-Rack System

Switch Pair	Connections	Cable Length
R3LL to R3SS	R3LL-P30 to R3SS-P19 R3LL-P32 to R3SS-P20 R3LL-P34 to R3SS-P26	3 meters
R3LL to R4SS	R3LL-P33 to R4SS-P20 R3LL-P35 to R4SS-P24 R3LL-P36 to R4SS-P26	5 meters
R3LL to R5SS	R3LL-P1 to R5SS-P26 R3LL-P5 to R5SS-P20	5 meters

Table 22-16 Leaf Switch Connections for the Fourth Rack in a Five-Rack System

Switch Pair	Connections	Cable Length
R4UL to R1SS	R4UL-P3 to R1SS-P31 R4UL-P7 to R1SS-P10	10 meters
R4UL to R2SS	R4UL-P2 to R2SS-P31 R4UL-P4 to R2SS-P9 R4UL-P6 to R2SS-P10	5 meters
R4UL to R3SS	R4UL-P31 to R3SS-P9 R4UL-P33 to R3SS-P10 R4UL-P35 to R3SS-P31	5 meters
R4UL to R4SS	R4UL-P30 to R4SS-P9 R4UL-P32 to R4SS-P10 R4UL-P34 to R4SS-P31	3 meters
R4UL to R5SS	R4UL-P1 to R5SS-P31 R4UL-P5 to R5SS-P10 R4UL-P36 to R5SS-P32	5 meters
R4LL to R1SS	R4LL-P1 to R1SS-P32 R4LL-P3 to R1SS-P30 R4LL-P5 to R1SS-P17	10 meters
R4LL to R2SS	R4LL-P2 to R2SS-P30 R4LL-P4 to R2SS-P17	5 meters
R4LL to R3SS	R4LL-P31 to R3SS-P17 R4LL-P33 to R3SS-P18 R4LL-P35 to R3SS-P30	5 meters
R4LL to R4SS	R4LL-P30 to R4SS-P17 R4LL-P32 to R4SS-P18 R4LL-P34 to R4SS-P30	3 meters
R4LL to R5SS	R4LL-P7 to R5SS-P17 R4LL-P6 to R5SS-P18 R4LL-P36 to R5SS-P30	5 meters

Table 22-17 Leaf Switch Connections for the Fifth Rack in a Five-Rack System

Switch Pair	Connections	Cable Length
R5UL to R1SS	R5UL-P3 to R1SS-P23 R5UL-P5 to R1SS-P8 R5UL-P7 to R1SS-P9	10 meters
R5UL to R2SS	R5UL-P1 to R2SS-P32 R5UL-P2 to R2SS-P23 R5UL-P4 to R2SS-P8	10 meters
R5UL to R3SS	R5UL-P6 to R3SS-P12 R5UL-P34 to R3SS-P23	5 meters
R5UL to R4SS	R5UL-P31 to R4SS-P7 R5UL-P33 to R4SS-P12 R5UL-P35 to R4SS-P32	5 meters
R5UL to R5SS	R5UL-P30 to R5SS-P7 R5UL-P32 to R5SS-P9 R5UL-P36 to R5SS-P23	3 meters
R5LL to R1SS	R5LL-P3 to R1SS-P29 R5LL-P5 to R1SS-P15 R5LL-P7 to R1SS-P18	10 meters
R5LL to R2SS	R5LL-P2 to R2SS-P29 R5LL-P4 to R2SS-P16 R5LL-P6 to R2SS-P18	10 meters
R5LL to R3SS	R5LL-P1 to R3SS-P32 R5LL-P31 to R3SS-P16 R5LL-P35 to R3SS-P29	5 meters
R5LL to R4SS	R5LL-P33 to R4SS-P19 R5LL-P34 to R4SS-P29	5 meters
R5LL to R5SS	R5LL-P30 to R5SS-P15 R5LL-P32 to R5SS-P19 R5LL-P36 to R5SS-P29	3 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-18 Five-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R1UL-P7	3 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R5UL-P7	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R3UL-P7	5 meters

Table 22-18 (Cont.) Five-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R5LL-P5	10 meters
R1SS-P16	R2LL-P7	5 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R3LL-P7	5 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R1UL-P3	3 meters
R1SS-P22	R2UL-P3	5 meters
R1SS-P23	R5UL-P3	10 meters
R1SS-P24	R2UL-P1	5 meters
R1SS-P25	R3UL-P3	5 meters
R1SS-P26	R3LL-P3	5 meters
R1SS-P27	R1LL-P3	3 meters
R1SS-P28	R2LL-P3	5 meters
R1SS-P29	R5LL-P3	10 meters
R1SS-P30	R4LL-P3	10 meters
R1SS-P31	R4UL-P3	10 meters
R1SS-P32	R4LL-P1	10 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R1UL-P33	5 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R4UL-P4	5 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R3UL-P31	5 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R1LL-P33	5 meters
R2SS-P16	R5LL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters
R2SS-P19	R3LL-P6	5 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R1UL-P35	5 meters
R2SS-P22	R2UL-P35	3 meters
R2SS-P23	R5UL-P2	10 meters
R2SS-P24	R2LL-P34	3 meters
R2SS-P25	R3UL-P2	5 meters
R2SS-P26	R3LL-P2	5 meters

Table 22-18 (Cont.) Five-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R2SS-P27	R1LL-P35	5 meters
R2SS-P28	R2LL-P35	3 meters
R2SS-P29	R5LL-P2	10 meters
R2SS-P30	R4LL-P2	5 meters
R2SS-P31	R4UL-P2	5 meters
R2SS-P32	R5UL-P1	10 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R1UL-P32	5 meters
R3SS-P8	R2UL-P30	5 meters
R3SS-P9	R4UL-P31	5 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R5UL-P6	5 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R1LL-P6	5 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R4LL-P33	5 meters
R3SS-P19	R3LL-P30	3 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R1UL-P1	5 meters
R3SS-P22	R2UL-P2	5 meters
R3SS-P23	R5UL-P34	5 meters
R3SS-P24	R3UL-P36	3 meters
R3SS-P25	R3UL-P34	3 meters
R3SS-P26	R3LL-P34	3 meters
R3SS-P27	R1LL-P2	5 meters
R3SS-P28	R2LL-P1	5 meters
R3SS-P29	R5LL-P35	5 meters
R3SS-P30	R4LL-P35	5 meters
R3SS-P31	R4UL-P35	5 meters
R3SS-P32	R5LL-P1	5 meters
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R5UL-P31	5 meters
R4SS-P8	R2UL-P32	5 meters
R4SS-P9	R4UL-P30	3 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R5UL-P33	5 meters

Table 22-18 (Cont.) Five-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R1LL-P32	10 meters
R4SS-P16	R2LL-P30	5 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R4LL-P32	3 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R1UL-P2	10 meters
R4SS-P22	R2UL-P34	5 meters
R4SS-P23	R1UL-P34	10 meters
R4SS-P24	R3LL-P35	5 meters
R4SS-P25	R3UL-P1	5 meters
R4SS-P26	R3LL-P36	5 meters
R4SS-P27	R1LL-P34	10 meters
R4SS-P28	R2LL-P2	5 meters
R4SS-P29	R5LL-P34	5 meters
R4SS-P30	R4LL-P34	3 meters
R4SS-P31	R4UL-P34	3 meters
R4SS-P32	R5UL-P35	5 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R2UL-P33	10 meters
R5SS-P9	R5UL-P32	3 meters
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R3UL-P33	5 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R2LL-P32	10 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R4LL-P6	5 meters
R5SS-P19	R5LL-P32	3 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R1UL-P36	10 meters
R5SS-P22	R2UL-P36	10 meters
R5SS-P23	R5UL-P36	3 meters
R5SS-P24	R3UL-P35	5 meters
R5SS-P25	R1LL-P1	10 meters
R5SS-P26	R3LL-P1	5 meters

Table 22-18 (Cont.) Five-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R5SS-P27	R1LL-P36	10 meters
R5SS-P28	R2LL-P36	10 meters
R5SS-P29	R5LL-P36	3 meters
R5SS-P30	R4LL-P36	5 meters
R5SS-P31	R4UL-P1	5 meters
R5SS-P32	R4UL-P36	5 meters

Six-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect six racks using RoCE Network Fabric.



Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, for these systems only, ignore the connections to port number 30 on every leaf switch in the following tables. This adjustment leaves only 13 inter-switch links on each leaf switch and only applies to X9M-8 systems with three database servers and 11 storage servers.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a six-rack system.

Table 22-19 Leaf Switch Connections for the First Rack in a Six-Rack System

Switch Pair	Connections	Cable Length
R1UL to R1SS	R1UL-P3 to R1SS-P21 R1UL-P5 to R1SS-P5 R1UL-P7 to R1SS-P7	3 meters
R1UL to R2SS	R1UL-P31 to R2SS-P5 R1UL-P33 to R2SS-P7 R1UL-P35 to R2SS-P21	5 meters
R1UL to R3SS	R1UL-P1 to R3SS-P21 R1UL-P6 to R3SS-P5	5 meters
R1UL to R4SS	R1UL-P2 to R4SS-P21 R1UL-P4 to R4SS-P5	10 meters
R1UL to R5SS	R1UL-P30 to R5SS-P5 R1UL-P34 to R5SS-P21	10 meters
R1UL to R6SS	R1UL-P32 to R6SS-P5 R1UL-P36 to R6SS-P21	10 meters
R1LL to R1SS	R1LL-P3 to R1SS-P27 R1LL-P5 to R1SS-P13	3 meters
R1LL to R2SS	R1LL-P31 to R2SS-P13 R1LL-P33 to R2SS-P15 R1LL-P35 to R2SS-P27	5 meters
R1LL to R3SS	R1LL-P2 to R3SS-P27 R1LL-P4 to R3SS-P13 R1LL-P6 to R3SS-P15	5 meters
R1LL to R4SS	R1LL-P30 to R4SS-P13 R1LL-P34 to R4SS-P27	10 meters
R1LL to R5SS	R1LL-P7 to R5SS-P13 R1LL-P1 to R5SS-P27	10 meters
R1LL to R6SS	R1LL-P32 to R6SS-P13 R1LL-P36 to R6SS-P27	10 meters

Table 22-20 Leaf Switch Connections for the Second Rack in a Six-Rack System

Switch Pair	Connections	Cable Length
R2UL to R1SS	R2UL-P3 to R1SS-P22 R2UL-P5 to R1SS-P6	5 meters
R2UL to R2SS	R2UL-P31 to R2SS-P6 R2UL-P35 to R2SS-P22	3 meters
R2UL to R3SS	R2UL-P6 to R3SS-P6 R2UL-P30 to R3SS-P8 R2UL-P2 to R3SS-P22	5 meters

Table 22-20 (Cont.) Leaf Switch Connections for the Second Rack in a Six-Rack System

Switch Pair	Connections	Cable Length
R2UL to R4SS	R2UL-P4 to R4SS-P6 R2UL-P32 to R4SS-P8 R2UL-P34 to R4SS-P22	5 meters
R2UL to R5SS	R2UL-P1 to R5SS-P22 R2UL-P7 to R5SS-P6	10 meters
R2UL to R6SS	R2UL-P33 to R6SS-P6 R2UL-P36 to R6SS-P22	10 meters
R2LL to R1SS	R2LL-P3 to R1SS-P28 R2LL-P5 to R1SS-P14	5 meters
R2LL to R2SS	R2LL-P31 to R2SS-P14 R2LL-P35 to R2SS-P28	3 meters
R2LL to R3SS	R2LL-P2 to R3SS-P28 R2LL-P6 to R3SS-P14	5 meters
R2LL to R4SS	R2LL-P4 to R4SS-P14 R2LL-P30 to R4SS-P16 R2LL-P34 to R4SS-P28	5 meters
R2LL to R5SS	R2LL-P32 to R5SS-P16 R2LL-P33 to R5SS-P14 R2LL-P36 to R5SS-P28	10 meters
R2LL to R6SS	R2LL-P1 to R6SS-P28 R2LL-P7 to R6SS-P14	10 meters

Table 22-21 Leaf Switch Connections for the Third Rack in a Six-Rack System

Switch Pair	Connections	Cable Length
R3UL to R1SS	R3UL-P3 to R1SS-P25 R3UL-P5 to R1SS-P11	5 meters
R3UL to R2SS	R3UL-P2 to R2SS-P25 R3UL-P6 to R2SS-P11	5 meters
R3UL to R3SS	R3UL-P30 to R3SS-P11 R3UL-P34 to R3SS-P25	3 meters
R3UL to R4SS	R3UL-P1 to R4SS-P25 R3UL-P4 to R4SS-P11	5 meters
R3UL to R5SS	R3UL-P32 to R5SS-P11 R3UL-P33 to R5SS-P12 R3UL-P35 to R5SS-P25	5 meters
R3UL to R6SS	R3UL-P7 to R6SS-P11 R3UL-P31 to R6SS-P12 R3UL-P36 to R6SS-P25	10 meters

Table 22-21 (Cont.) Leaf Switch Connections for the Third Rack in a Six-Rack System

Switch Pair	Connections	Cable Length
R3LL to R1SS	R3LL-P3 to R1SS-P26 R3LL-P4 to R1SS-P20 R3LL-P7 to R1SS-P19	5 meters
R3LL to R2SS	R3LL-P31 to R2SS-P20 R3LL-P35 to R2SS-P26	5 meters
R3LL to R3SS	R3LL-P32 to R3SS-P20 R3LL-P34 to R3SS-P26	3 meters
R3LL to R4SS	R3LL-P33 to R4SS-P20 R3LL-P36 to R4SS-P26	5 meters
R3LL to R5SS	R3LL-P1 to R5SS-P26 R3LL-P5 to R5SS-P20	5 meters
R3LL to R6SS	R3LL-P2 to R6SS-P26 R3LL-P6 to R6SS-P19 R3LL-P30 to R6SS-P20	10 meters

Table 22-22 Leaf Switch Connections for the Fourth Rack in a Six-Rack System

Switch Pair	Connections	Cable Length
R4UL to R1SS	R4UL-P1 to R1SS-P32 R4UL-P3 to R1SS-P31 R4UL-P7 to R1SS-P10	10 meters
R4UL to R2SS	R4UL-P2 to R2SS-P31 R4UL-P4 to R2SS-P9 R4UL-P6 to R2SS-P10	5 meters
R4UL to R3SS	R4UL-P31 to R3SS-P9 R4UL-P33 to R3SS-P10	5 meters
R4UL to R4SS	R4UL-P32 to R4SS-P10 R4UL-P34 to R4SS-P31	3 meters
R4UL to R5SS	R4UL-P5 to R5SS-P10 R4UL-P35 to R5SS-P31	5 meters
R4UL to R6SS	R4UL-P30 to R6SS-P10 R4UL-P36 to R6SS-P31	5 meters
R4LL to R1SS	R4LL-P3 to R1SS-P30 R4LL-P5 to R1SS-P17	10 meters
R4LL to R2SS	R4LL-P1 to R2SS-P32 R4LL-P2 to R2SS-P30 R4LL-P4 to R2SS-P17	5 meters
R4LL to R3SS	R4LL-P31 to R3SS-P17 R4LL-P33 to R3SS-P18 R4LL-P35 to R3SS-P30	5 meters

Table 22-22 (Cont.) Leaf Switch Connections for the Fourth Rack in a Six-Rack System

Switch Pair	Connections	Cable Length
R4LL to R4SS	R4LL-P30 to R4SS-P17 R4LL-P32 to R4SS-P18	3 meters
R4LL to R5SS	R4LL-P7 to R5SS-P17 R4LL-P34 to R5SS-P30	5 meters
R4LL to R6SS	R4LL-P6 to R6SS-P17 R4LL-P36 to R6SS-P30	5 meters

Table 22-23 Leaf Switch Connections for the Fifth Rack in a Six-Rack System

Switch Pair	Connections	Cable Length
R5UL to R1SS	R5UL-P3 to R1SS-P23 R5UL-P5 to R1SS-P8	10 meters
R5UL to R2SS	R5UL-P2 to R2SS-P23 R5UL-P4 to R2SS-P8	10 meters
R5UL to R3SS	R5UL-P6 to R3SS-P12 R5UL-P1 to R3SS-P23 R5UL-P34 to R3SS-P24	5 meters
R5UL to R4SS	R5UL-P31 to R4SS-P7 R5UL-P33 to R4SS-P12 R5UL-P35 to R4SS-P23	5 meters
R5UL to R5SS	R5UL-P30 to R5SS-P7 R5UL-P32 to R5SS-P9	3 meters
R5UL to R6SS	R5UL-P7 to R6SS-P7 R5UL-P36 to R6SS-P23	5 meters
R5LL to R1SS	R5LL-P3 to R1SS-P29 R5LL-P7 to R1SS-P18	10 meters
R5LL to R2SS	R5LL-P6 to R2SS-P18 R5LL-P34 to R2SS-P29	10 meters
R5LL to R3SS	R5LL-P31 to R3SS-P16 R5LL-P35 to R3SS-P29	5 meters
R5LL to R4SS	R5LL-P1 to R4SS-P29 R5LL-P2 to R4SS-P30 R5LL-P33 to R4SS-P19	5 meters
R5LL to R5SS	R5LL-P30 to R5SS-P15 R5LL-P32 to R5SS-P19 R5LL-P36 to R5SS-P29	3 meters
R5LL to R6SS	R5LL-P4 to R6SS-P15 R5LL-P5 to R6SS-P9	5 meters

Table 22-24 Leaf Switch Connections for the Sixth Rack in a Six-Rack System

Switch Pair	Connections	Cable Length
R6UL to R1SS	R6UL-P5 to R1SS-P9 R6UL-P7 to R1SS-P16	10 meters
R6UL to R2SS	R6UL-P2 to R2SS-P24 R6UL-P4 to R2SS-P16	10 meters
R6UL to R3SS	R6UL-P3 to R3SS-P31 R6UL-P6 to R3SS-P19	10 meters
R6UL to R4SS	R6UL-P1 to R4SS-P32 R6UL-P31 to R4SS-P15	5 meters
R6UL to R5SS	R6UL-P33 to R5SS-P18 R6UL-P34 to R5SS-P24 R6UL-P35 to R5SS-P23	5 meters
R6UL to R6SS	R6UL-P30 to R6SS-P16 R6UL-P32 to R6SS-P18 R6UL-P36 to R6SS-P24	3 meters
R6LL to R1SS	R6LL-P3 to R1SS-P24 R6LL-P5 to R1SS-P12 R6LL-P7 to R1SS-P15	10 meters
R6LL to R2SS	R6LL-P4 to R2SS-P12 R6LL-P6 to R2SS-P19	10 meters
R6LL to R3SS	R6LL-P1 to R3SS-P32 R6LL-P31 to R3SS-P7	10 meters
R6LL to R4SS	R6LL-P2 to R4SS-P24 R6LL-P33 to R4SS-P9	5 meters
R6LL to R5SS	R6LL-P30 to R5SS-P8 R6LL-P34 to R5SS-P32	5 meters
R6LL to R6SS	R6LL-P32 to R6SS-P8 R6LL-P35 to R6SS-P29 R6LL-P36 to R6SS-P32	3 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-25 Six-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R1UL-P7	3 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R6UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters

Table 22-25 (Cont.) Six-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P12	R6LL-P5	10 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R6LL-P7	10 meters
R1SS-P16	R6UL-P7	10 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R3LL-P7	5 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R1UL-P3	3 meters
R1SS-P22	R2UL-P3	5 meters
R1SS-P23	R5UL-P3	10 meters
R1SS-P24	R6LL-P3	10 meters
R1SS-P25	R3UL-P3	5 meters
R1SS-P26	R3LL-P3	5 meters
R1SS-P27	R1LL-P3	3 meters
R1SS-P28	R2LL-P3	5 meters
R1SS-P29	R5LL-P3	10 meters
R1SS-P30	R4LL-P3	10 meters
R1SS-P31	R4UL-P3	10 meters
R1SS-P32	R4UL-P1	10 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R1UL-P33	5 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R4UL-P4	5 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R6LL-P4	10 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R1LL-P33	5 meters
R2SS-P16	R6UL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters
R2SS-P19	R6LL-P6	10 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R1UL-P35	5 meters
R2SS-P22	R2UL-P35	3 meters
R2SS-P23	R5UL-P2	10 meters
R2SS-P24	R6UL-P2	10 meters
R2SS-P25	R3UL-P2	5 meters

Table 22-25 (Cont.) Six-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R2SS-P26	R3LL-P35	5 meters
R2SS-P27	R1LL-P35	5 meters
R2SS-P28	R2LL-P35	3 meters
R2SS-P29	R5LL-P34	10 meters
R2SS-P30	R4LL-P2	5 meters
R2SS-P31	R4UL-P2	5 meters
R2SS-P32	R4LL-P1	5 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R6LL-P31	10 meters
R3SS-P8	R2UL-P30	5 meters
R3SS-P9	R4UL-P31	5 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R5UL-P6	5 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R1LL-P6	5 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R4LL-P33	5 meters
R3SS-P19	R6UL-P6	10 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R1UL-P1	5 meters
R3SS-P22	R2UL-P2	5 meters
R3SS-P23	R5UL-P1	5 meters
R3SS-P24	R5UL-P34	5 meters
R3SS-P25	R3UL-P34	3 meters
R3SS-P26	R3LL-P34	3 meters
R3SS-P27	R1LL-P2	5 meters
R3SS-P28	R2LL-P2	5 meters
R3SS-P29	R5LL-P35	5 meters
R3SS-P30	R4LL-P35	5 meters
R3SS-P31	R6UL-P3	10 meters
R3SS-P32	R6LL-P1	10 meters
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R5UL-P31	5 meters
R4SS-P8	R2UL-P32	5 meters
R4SS-P9	R6LL-P33	5 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters

Table 22-25 (Cont.) Six-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R4SS-P12	R5UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R6UL-P31	5 meters
R4SS-P16	R2LL-P30	5 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R4LL-P32	3 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R1UL-P2	10 meters
R4SS-P22	R2UL-P34	5 meters
R4SS-P23	R5UL-P35	5 meters
R4SS-P24	R6LL-P2	5 meters
R4SS-P25	R3UL-P1	5 meters
R4SS-P26	R3LL-P36	5 meters
R4SS-P27	R1LL-P34	10 meters
R4SS-P28	R2LL-P34	5 meters
R4SS-P29	R5LL-P1	5 meters
R4SS-P30	R5LL-P2	5 meters
R4SS-P31	R4UL-P34	3 meters
R4SS-P32	R6UL-P1	5 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R6LL-P30	5 meters
R5SS-P9	R5UL-P32	3 meters
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R3UL-P33	5 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R2LL-P32	10 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R6UL-P33	5 meters
R5SS-P19	R5LL-P32	3 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R1UL-P34	10 meters
R5SS-P22	R2UL-P1	10 meters
R5SS-P23	R6UL-P35	5 meters
R5SS-P24	R6UL-P34	5 meters
R5SS-P25	R3UL-P35	5 meters

Table 22-25 (Cont.) Six-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R5SS-P26	R3LL-P1	5 meters
R5SS-P27	R1LL-P1	10 meters
R5SS-P28	R2LL-P36	10 meters
R5SS-P29	R5LL-P36	3 meters
R5SS-P30	R4LL-P34	5 meters
R5SS-P31	R4UL-P35	5 meters
R5SS-P32	R6LL-P34	5 meters
R6SS-P5	R1UL-P32	10 meters
R6SS-P6	R2UL-P33	10 meters
R6SS-P7	R5UL-P7	5 meters
R6SS-P8	R6LL-P32	3 meters
R6SS-P9	R5LL-P5	5 meters
R6SS-P10	R4UL-P30	5 meters
R6SS-P11	R3UL-P7	10 meters
R6SS-P12	R3UL-P31	10 meters
R6SS-P13	R1LL-P32	10 meters
R6SS-P14	R2LL-P7	10 meters
R6SS-P15	R5LL-P4	5 meters
R6SS-P16	R6UL-P30	3 meters
R6SS-P17	R4LL-P6	5 meters
R6SS-P18	R6UL-P32	3 meters
R6SS-P19	R3LL-P6	10 meters
R6SS-P20	R3LL-P30	10 meters
R6SS-P21	R1UL-P36	10 meters
R6SS-P22	R2UL-P36	10 meters
R6SS-P23	R5UL-P36	5 meters
R6SS-P24	R6UL-P36	3 meters
R6SS-P25	R3UL-P36	10 meters
R6SS-P26	R3LL-P2	10 meters
R6SS-P27	R1LL-P36	10 meters
R6SS-P28	R2LL-P1	10 meters
R6SS-P29	R6LL-P35	3 meters
R6SS-P30	R4LL-P36	5 meters
R6SS-P31	R4UL-P36	5 meters
R6SS-P32	R6LL-P36	3 meters

Seven-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect seven racks using RoCE Network Fabric.

 Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, for these systems only, ignore the connections to port number 30 on every leaf switch in the following tables. This adjustment leaves only 13 inter-switch links on each leaf switch and only applies to X9M-8 systems with three database servers and 11 storage servers.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a seven-rack system.

Table 22-26 Leaf Switch Connections for the First Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R1UL to R1SS	R1UL-P5 to R1SS-P5 R1UL-P7 to R1SS-P7	3 meters
R1UL to R2SS	R1UL-P3 to R2SS-P21 R1UL-P31 to R2SS-P5	5 meters
R1UL to R3SS	R1UL-P1 to R3SS-P27 R1UL-P6 to R3SS-P5	5 meters
R1UL to R4SS	R1UL-P2 to R4SS-P22 R1UL-P4 to R4SS-P5	10 meters
R1UL to R5SS	R1UL-P30 to R5SS-P5 R1UL-P34 to R5SS-P29	10 meters

Table 22-26 (Cont.) Leaf Switch Connections for the First Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R1UL to R6SS	R1UL-P32 to R6SS-P5 R1UL-P35 to R6SS-P23	10 meters
R1UL to R7SS	R1UL-P33 to R7SS-P5 R1UL-P36 to R7SS-P23	10 meters
R1LL to R1SS	R1LL-P3 to R1SS-P21 R1LL-P5 to R1SS-P13	3 meters
R1LL to R2SS	R1LL-P31 to R2SS-P13 R1LL-P33 to R2SS-P15	5 meters
R1LL to R3SS	R1LL-P2 to R3SS-P21 R1LL-P4 to R3SS-P13	5 meters
R1LL to R4SS	R1LL-P30 to R4SS-P13 R1LL-P34 to R4SS-P21	10 meters
R1LL to R5SS	R1LL-P7 to R5SS-P13 R1LL-P35 to R5SS-P21	10 meters
R1LL to R6SS	R1LL-P32 to R6SS-P13 R1LL-P36 to R6SS-P21	10 meters
R1LL to R7SS	R1LL-P1 to R7SS-P21 R1LL-P6 to R7SS-P13	10 meters

Table 22-27 Leaf Switch Connections for the Second Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R2UL to R1SS	R2UL-P3 to R1SS-P28 R2UL-P5 to R1SS-P6	5 meters
R2UL to R2SS	R2UL-P31 to R2SS-P6 R2UL-P34 to R2SS-P28	3 meters
R2UL to R3SS	R2UL-P6 to R3SS-P6 R2UL-P30 to R3SS-P8	5 meters
R2UL to R4SS	R2UL-P2 to R4SS-P28 R2UL-P4 to R4SS-P6	5 meters
R2UL to R5SS	R2UL-P1 to R5SS-P28 R2UL-P7 to R5SS-P6	10 meters
R2UL to R6SS	R2UL-P33 to R6SS-P6 R2UL-P35 to R6SS-P28	10 meters
R2UL to R7SS	R2UL-P32 to R7SS-P6 R2UL-P36 to R7SS-P28	10 meters
R2LL to R1SS	R2LL-P3 to R1SS-P22 R2LL-P5 to R1SS-P14	5 meters
R2LL to R2SS	R2LL-P2 to R2SS-P22 R2LL-P31 to R2SS-P14	3 meters

Table 22-27 (Cont.) Leaf Switch Connections for the Second Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R2LL to R3SS	R2LL-P1 to R3SS-P22 R2LL-P6 to R3SS-P14	5 meters
R2LL to R4SS	R2LL-P4 to R4SS-P14 R2LL-P30 to R4SS-P16	5 meters
R2LL to R5SS	R2LL-P33 to R5SS-P14 R2LL-P35 to R5SS-P22	10 meters
R2LL to R6SS	R2LL-P7 to R6SS-P14 R2LL-P36 to R6SS-P22	10 meters
R2LL to R7SS	R2LL-P32 to R7SS-P14 R2LL-P34 to R7SS-P22	10 meters

Table 22-28 Leaf Switch Connections for the Third Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R3UL to R1SS	R3UL-P3 to R1SS-P25 R3UL-P5 to R1SS-P11	5 meters
R3UL to R2SS	R3UL-P2 to R2SS-P25 R3UL-P6 to R2SS-P11	5 meters
R3UL to R3SS	R3UL-P30 to R3SS-P11 R3UL-P34 to R3SS-P25	3 meters
R3UL to R4SS	R3UL-P1 to R4SS-P25 R3UL-P4 to R4SS-P11	5 meters
R3UL to R5SS	R3UL-P32 to R5SS-P11 R3UL-P33 to R5SS-P12	5 meters
R3UL to R6SS	R3UL-P7 to R6SS-P11 R3UL-P36 to R6SS-P32	10 meters
R3UL to R7SS	R3UL-P31 to R7SS-P11 R3UL-P35 to R7SS-P12	10 meters
R3LL to R1SS	R3LL-P2 to R1SS-P32 R3LL-P4 to R1SS-P20	5 meters
R3LL to R2SS	R3LL-P1 to R2SS-P32 R3LL-P31 to R2SS-P20	5 meters
R3LL to R3SS	R3LL-P32 to R3SS-P20 R3LL-P34 to R3SS-P32	3 meters
R3LL to R4SS	R3LL-P33 to R4SS-P20 R3LL-P35 to R4SS-P32	5 meters
R3LL to R5SS	R3LL-P3 to R5SS-P32 R3LL-P5 to R5SS-P20	5 meters
R3LL to R6SS	R3LL-P6 to R6SS-P19 R3LL-P30 to R6SS-P20	10 meters

Table 22-28 (Cont.) Leaf Switch Connections for the Third Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R3LL to R7SS	R3LL-P7 to R7SS-P20 R3LL-P36 to R7SS-P32	10 meters

Table 22-29 Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R4UL to R1SS	R4UL-P3 to R1SS-P31 R4UL-P7 to R1SS-P10	10 meters
R4UL to R2SS	R4UL-P4 to R2SS-P9 R4UL-P6 to R2SS-P10	5 meters
R4UL to R3SS	R4UL-P1 to R3SS-P31 R4UL-P33 to R3SS-P10	5 meters
R4UL to R4SS	R4UL-P32 to R4SS-P10 R4UL-P34 to R4SS-P31	3 meters
R4UL to R5SS	R4UL-P2 to R5SS-P31 R4UL-P5 to R5SS-P10	5 meters
R4UL to R6SS	R4UL-P30 to R6SS-P10 R4UL-P35 to R6SS-P31	5 meters
R4UL to R7SS	R4UL-P31 to R7SS-P10 R4UL-P36 to R7SS-P31	10 meters
R4LL to R1SS	R4LL-P3 to R1SS-P27 R4LL-P5 to R1SS-P17	10 meters
R4LL to R2SS	R4LL-P2 to R2SS-P27 R4LL-P4 to R2SS-P17	5 meters
R4LL to R3SS	R4LL-P31 to R3SS-P17 R4LL-P33 to R3SS-P18	5 meters
R4LL to R4SS	R4LL-P30 to R4SS-P17 R4LL-P34 to R4SS-P27	3 meters
R4LL to R5SS	R4LL-P1 to R5SS-P27 R4LL-P7 to R5SS-P17	5 meters
R4LL to R6SS	R4LL-P6 to R6SS-P17 R4LL-P35 to R6SS-P27	5 meters
R4LL to R7SS	R4LL-P32 to R7SS-P17 R4LL-P36 to R7SS-P27	10 meters

Table 22-30 Leaf Switch Connections for the Fifth Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R5UL to R1SS	R5UL-P3 to R1SS-P24 R5UL-P5 to R1SS-P8	10 meters

Table 22-30 (Cont.) Leaf Switch Connections for the Fifth Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R5UL to R2SS	R5UL-P2 to R2SS-P24 R5UL-P4 to R2SS-P8	10 meters
R5UL to R3SS	R5UL-P1 to R3SS-P24 R5UL-P6 to R3SS-P12	5 meters
R5UL to R4SS	R5UL-P31 to R4SS-P7 R5UL-P33 to R4SS-P12	5 meters
R5UL to R5SS	R5UL-P30 to R5SS-P7 R5UL-P34 to R5SS-P25	3 meters
R5UL to R6SS	R5UL-P7 to R6SS-P7 R5UL-P35 to R6SS-P25	5 meters
R5UL to R7SS	R5UL-P32 to R7SS-P7 R5UL-P36 to R7SS-P25	5 meters
R5LL to R1SS	R5LL-P3 to R1SS-P29 R5LL-P7 to R1SS-P18	10 meters
R5LL to R2SS	R5LL-P2 to R2SS-P29 R5LL-P6 to R2SS-P18	10 meters
R5LL to R3SS	R5LL-P31 to R3SS-P16 R5LL-P34 to R3SS-P29	5 meters
R5LL to R4SS	R5LL-P33 to R4SS-P19 R5LL-P35 to R4SS-P29	5 meters
R5LL to R5SS	R5LL-P30 to R5SS-P15 R5LL-P32 to R5SS-P19	3 meters
R5LL to R6SS	R5LL-P1 to R6SS-P29 R5LL-P4 to R6SS-P15	5 meters
R5LL to R7SS	R5LL-P5 to R7SS-P15 R5LL-P36 to R7SS-P29	5 meters

Table 22-31 Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R6UL to R1SS	R6UL-P3 to R1SS-P23 R6UL-P7 to R1SS-P16	10 meters
R6UL to R2SS	R6UL-P2 to R2SS-P23 R6UL-P4 to R2SS-P16	10 meters
R6UL to R3SS	R6UL-P1 to R3SS-P23 R6UL-P6 to R3SS-P19	10 meters
R6UL to R4SS	R6UL-P31 to R4SS-P15 R6UL-P34 to R4SS-P23	5 meters
R6UL to R5SS	R6UL-P33 to R5SS-P18 R6UL-P35 to R5SS-P23	5 meters

Table 22-31 (Cont.) Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R6UL to R6SS	R6UL-P30 to R6SS-P16 R6UL-P32 to R6SS-P18	3 meters
R6UL to R7SS	R6UL-P5 to R7SS-P16 R6UL-P36 to R7SS-P26	5 meters
R6LL to R1SS	R6LL-P5 to R1SS-P12 R6LL-P7 to R1SS-P15	10 meters
R6LL to R2SS	R6LL-P2 to R2SS-P31 R6LL-P4 to R2SS-P12	10 meters
R6LL to R3SS	R6LL-P3 to R3SS-P28 R6LL-P31 to R3SS-P7	10 meters
R6LL to R4SS	R6LL-P33 to R4SS-P9 R6LL-P34 to R4SS-P24	5 meters
R6LL to R5SS	R6LL-P30 to R5SS-P8 R6LL-P35 to R5SS-P24	5 meters
R6LL to R6SS	R6LL-P32 to R6SS-P8 R6LL-P36 to R6SS-P24	3 meters
R6LL to R7SS	R6LL-P1 to R7SS-P24 R6LL-P6 to R7SS-P8	5 meters

Table 22-32 Leaf Switch Connections for the Seventh Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R7UL to R1SS	R7UL-P3 to R1SS-P30 R7UL-P5 to R1SS-P9	10 meters
R7UL to R2SS	R7UL-P1 to R2SS-P30 R7UL-P7 to R2SS-P7	10 meters
R7UL to R3SS	R7UL-P2 to R3SS-P30 R7UL-P4 to R3SS-P9	10 meters
R7UL to R4SS	R7UL-P6 to R4SS-P8 R7UL-P34 to R4SS-P30	10 meters
R7UL to R5SS	R7UL-P31 to R5SS-P9 R7UL-P35 to R5SS-P30	5 meters
R7UL to R6SS	R7UL-P33 to R6SS-P9 R7UL-P36 to R6SS-P30	5 meters
R7UL to R7SS	R7UL-P30 to R7SS-P9 R7UL-P32 to R7SS-P30	3 meters
R7LL to R1SS	R7LL-P3 to R1SS-P26 R7LL-P5 to R1SS-P19	10 meters
R7LL to R2SS	R7LL-P1 to R2SS-P26 R7LL-P7 to R2SS-P19	10 meters

Table 22-32 (Cont.) Leaf Switch Connections for the Seventh Rack in a Seven-Rack System

Switch Pair	Connections	Cable Length
R7LL to R3SS	R7LL-P2 to R3SS-P26 R7LL-P4 to R3SS-P15	10 meters
R7LL to R4SS	R7LL-P6 to R4SS-P18 R7LL-P34 to R4SS-P26	10 meters
R7LL to R5SS	R7LL-P31 to R5SS-P16 R7LL-P35 to R5SS-P26	5 meters
R7LL to R6SS	R7LL-P33 to R6SS-P12 R7LL-P36 to R6SS-P26	5 meters
R7LL to R7SS	R7LL-P30 to R7SS-P18 R7LL-P32 to R7SS-P19	3 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-33 Seven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R1UL-P7	3 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R7UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R6LL-P5	10 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R6LL-P7	10 meters
R1SS-P16	R6UL-P7	10 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R7LL-P5	10 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R1LL-P3	3 meters
R1SS-P22	R2LL-P3	5 meters
R1SS-P23	R6UL-P3	10 meters
R1SS-P24	R5UL-P3	10 meters
R1SS-P25	R3UL-P3	5 meters
R1SS-P26	R7LL-P3	10 meters
R1SS-P27	R4LL-P3	10 meters
R1SS-P28	R2UL-P3	5 meters

Table 22-33 (Cont.) Seven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P29	R5LL-P3	10 meters
R1SS-P30	R7UL-P3	10 meters
R1SS-P31	R4UL-P3	10 meters
R1SS-P32	R3LL-P2	5 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R7UL-P7	10 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R4UL-P4	5 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R6LL-P4	10 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R1LL-P33	5 meters
R2SS-P16	R6UL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters
R2SS-P19	R7LL-P7	10 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R1UL-P3	5 meters
R2SS-P22	R2LL-P2	3 meters
R2SS-P23	R6UL-P2	10 meters
R2SS-P24	R5UL-P2	10 meters
R2SS-P25	R3UL-P2	5 meters
R2SS-P26	R7LL-P1	10 meters
R2SS-P27	R4LL-P2	5 meters
R2SS-P28	R2UL-P34	3 meters
R2SS-P29	R5LL-P2	10 meters
R2SS-P30	R7UL-P1	10 meters
R2SS-P31	R6LL-P2	10 meters
R2SS-P32	R3LL-P1	5 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R6LL-P31	10 meters
R3SS-P8	R2UL-P30	5 meters
R3SS-P9	R7UL-P4	10 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R5UL-P6	5 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters

Table 22-33 (Cont.) Seven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R3SS-P15	R7LL-P4	10 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R4LL-P33	5 meters
R3SS-P19	R6UL-P6	10 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R1LL-P2	5 meters
R3SS-P22	R2LL-P1	5 meters
R3SS-P23	R6UL-P1	10 meters
R3SS-P24	R5UL-P1	5 meters
R3SS-P25	R3UL-P34	3 meters
R3SS-P26	R7LL-P2	10 meters
R3SS-P27	R1UL-P1	5 meters
R3SS-P28	R6LL-P3	10 meters
R3SS-P29	R5LL-P34	5 meters
R3SS-P30	R7UL-P2	10 meters
R3SS-P31	R4UL-P1	5 meters
R3SS-P32	R3LL-P34	3 meters
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R5UL-P31	5 meters
R4SS-P8	R7UL-P6	10 meters
R4SS-P9	R6LL-P33	5 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R5UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R6UL-P31	5 meters
R4SS-P16	R2LL-P30	5 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R7LL-P6	10 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R1LL-P34	10 meters
R4SS-P22	R1UL-P2	10 meters
R4SS-P23	R6UL-P34	5 meters
R4SS-P24	R6LL-P34	5 meters
R4SS-P25	R3UL-P1	5 meters
R4SS-P26	R7LL-P34	10 meters
R4SS-P27	R4LL-P34	3 meters
R4SS-P28	R2UL-P2	5 meters

Table 22-33 (Cont.) Seven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R4SS-P29	R5LL-P35	5 meters
R4SS-P30	R7UL-P34	10 meters
R4SS-P31	R4UL-P34	3 meters
R4SS-P32	R3LL-P35	5 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R6LL-P30	5 meters
R5SS-P9	R7UL-P31	5 meters
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R3UL-P33	5 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R7LL-P31	5 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R6UL-P33	5 meters
R5SS-P19	R5LL-P32	3 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R1LL-P35	10 meters
R5SS-P22	R2LL-P35	10 meters
R5SS-P23	R6UL-P35	5 meters
R5SS-P24	R6LL-P35	5 meters
R5SS-P25	R5UL-P34	3 meters
R5SS-P26	R7LL-P35	5 meters
R5SS-P27	R4LL-P1	5 meters
R5SS-P28	R2UL-P1	10 meters
R5SS-P29	R1UL-P34	10 meters
R5SS-P30	R7UL-P35	5 meters
R5SS-P31	R4UL-P2	5 meters
R5SS-P32	R3LL-P3	5 meters
R6SS-P5	R1UL-P32	10 meters
R6SS-P6	R2UL-P33	10 meters
R6SS-P7	R5UL-P7	5 meters
R6SS-P8	R6LL-P32	3 meters
R6SS-P9	R7UL-P33	5 meters
R6SS-P10	R4UL-P30	5 meters
R6SS-P11	R3UL-P7	10 meters
R6SS-P12	R7LL-P33	5 meters
R6SS-P13	R1LL-P32	10 meters
R6SS-P14	R2LL-P7	10 meters

Table 22-33 (Cont.) Seven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R6SS-P15	R5LL-P4	5 meters
R6SS-P16	R6UL-P30	3 meters
R6SS-P17	R4LL-P6	5 meters
R6SS-P18	R6UL-P32	3 meters
R6SS-P19	R3LL-P6	10 meters
R6SS-P20	R3LL-P30	10 meters
R6SS-P21	R1LL-P36	10 meters
R6SS-P22	R2LL-P36	10 meters
R6SS-P23	R1UL-P35	10 meters
R6SS-P24	R6LL-P36	3 meters
R6SS-P25	R5UL-P35	5 meters
R6SS-P26	R7LL-P36	5 meters
R6SS-P27	R4LL-P35	5 meters
R6SS-P28	R2UL-P35	10 meters
R6SS-P29	R5LL-P1	5 meters
R6SS-P30	R7UL-P36	5 meters
R6SS-P31	R4UL-P35	5 meters
R6SS-P32	R3UL-P36	10 meters
R7SS-P5	R1UL-P33	10 meters
R7SS-P6	R2UL-P32	10 meters
R7SS-P7	R5UL-P32	5 meters
R7SS-P8	R6LL-P6	5 meters
R7SS-P9	R7UL-P30	3 meters
R7SS-P10	R4UL-P31	10 meters
R7SS-P11	R3UL-P31	10 meters
R7SS-P12	R3UL-P35	10 meters
R7SS-P13	R1LL-P6	10 meters
R7SS-P14	R2LL-P32	10 meters
R7SS-P15	R5LL-P5	5 meters
R7SS-P16	R6UL-P5	5 meters
R7SS-P17	R4LL-P32	10 meters
R7SS-P18	R7LL-P30	3 meters
R7SS-P19	R7LL-P32	3 meters
R7SS-P20	R3LL-P7	10 meters
R7SS-P21	R1LL-P1	10 meters
R7SS-P22	R2LL-P34	10 meters
R7SS-P23	R1UL-P36	10 meters
R7SS-P24	R6LL-P1	5 meters
R7SS-P25	R5UL-P36	5 meters
R7SS-P26	R6UL-P36	5 meters
R7SS-P27	R4LL-P36	10 meters
R7SS-P28	R2UL-P36	10 meters

Table 22-33 (Cont.) Seven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R7SS-P29	R5LL-P36	5 meters
R7SS-P30	R7UL-P32	3 meters
R7SS-P31	R4UL-P36	10 meters
R7SS-P32	R3LL-P36	10 meters

Eight-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect eight racks using RoCE Network Fabric.

Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, for these systems only, ignore the connections to port number 30 on every leaf switch in the following tables. This adjustment leaves only 13 inter-switch links on each leaf switch and only applies to X9M-8 systems with three database servers and 11 storage servers.

The following tables contain details for all of the RoCE Network Fabric cabling connections in an eight-rack system.

Table 22-34 Leaf Switch Connections for the First Rack in a Eight-Rack System

Connection	Cable Length
R1UL-P1 to R3SS-P23	5 meters
R1UL-P2 to R4SS-P23	10 meters
R1UL-P3 to R1SS-P23	3 meters
R1UL-P4 to R4SS-P5	10 meters
R1UL-P5 to R1SS-P5	3 meters
R1UL-P6 to R3SS-P5	5 meters
R1UL-P7 to R8SS-P5	10 meters
R1UL-P30 to R5SS-P5	10 meters
R1UL-P31 to R2SS-P5	5 meters
R1UL-P32 to R6SS-P5	10 meters
R1UL-P33 to R7SS-P5	10 meters
R1UL-P34 to R2SS-P23	5 meters
R1UL-P35 to R5SS-P23	10 meters
R1UL-P36 to R6SS-P23	10 meters
R1LL-P1 to R2SS-P24	5 meters
R1LL-P2 to R3SS-P24	5 meters
R1LL-P3 to R5SS-P24	10 meters
R1LL-P4 to R3SS-P13	5 meters
R1LL-P5 to R1SS-P13	3 meters
R1LL-P6 to R7SS-P13	10 meters
R1LL-P7 to R5SS-P13	10 meters
R1LL-P30 to R4SS-P13	10 meters
R1LL-P31 to R2SS-P13	5 meters
R1LL-P32 to R6SS-P13	10 meters
R1LL-P33 to R8SS-P13	10 meters
R1LL-P34 to R4SS-P24	10 meters
R1LL-P35 to R6SS-P24	10 meters
R1LL-P36 to R7SS-P24	10 meters

Table 22-35 Leaf Switch Connections for the Second Rack in a Eight-Rack System

Connection	Cable Length
R2UL-P1 to R5SS-P25	10 meters
R2UL-P2 to R3SS-P25	5 meters
R2UL-P3 to R4SS-P25	5 meters
R2UL-P4 to R4SS-P6	5 meters
R2UL-P5 to R1SS-P6	5 meters
R2UL-P6 to R3SS-P6	5 meters
R2UL-P7 to R5SS-P6	10 meters
R2UL-P30 to R8SS-P6	10 meters
R2UL-P31 to R2SS-P6	3 meters
R2UL-P32 to R7SS-P6	10 meters
R2UL-P33 to R6SS-P6	10 meters

Table 22-35 (Cont.) Leaf Switch Connections for the Second Rack in a Eight-Rack System

Connection	Cable Length
R2UL-P34 to R7SS-P25	10 meters
R2UL-P35 to R6SS-P25	10 meters
R2UL-P36 to R8SS-P25	10 meters
R2LL-P1 to R6SS-P32	10 meters
R2LL-P2 to R4SS-P32	5 meters
R2LL-P3 to R1SS-P22	5 meters
R2LL-P4 to R4SS-P14	5 meters
R2LL-P5 to R1SS-P14	5 meters
R2LL-P6 to R3SS-P14	5 meters
R2LL-P7 to R6SS-P14	10 meters
R2LL-P30 to R8SS-P14	10 meters
R2LL-P31 to R2SS-P14	3 meters
R2LL-P32 to R7SS-P14	10 meters
R2LL-P33 to R5SS-P14	10 meters
R2LL-P34 to R7SS-P32	10 meters
R2LL-P35 to R5SS-P32	10 meters
R2LL-P36 to R8SS-P32	10 meters

Table 22-36 Leaf Switch Connections for the Third Rack in a Eight-Rack System

Connection	Cable Length
R3UL-P1 to R6SS-P31	10 meters
R3UL-P2 to R2SS-P25	5 meters
R3UL-P3 to R1SS-P25	5 meters
R3UL-P4 to R4SS-P11	5 meters
R3UL-P5 to R1SS-P11	5 meters
R3UL-P6 to R2SS-P11	5 meters
R3UL-P7 to R6SS-P11	10 meters
R3UL-P30 to R3SS-P11	3 meters
R3UL-P31 to R7SS-P11	10 meters
R3UL-P32 to R5SS-P11	5 meters
R3UL-P33 to R8SS-P11	10 meters
R3UL-P34 to R5SS-P31	5 meters
R3UL-P35 to R7SS-P31	10 meters
R3UL-P36 to R8SS-P31	10 meters
R3LL-P1 to R7SS-P27	10 meters
R3LL-P2 to R1SS-P32	5 meters
R3LL-P3 to R2SS-P32	5 meters
R3LL-P4 to R1SS-P20	5 meters
R3LL-P5 to R5SS-P20	5 meters
R3LL-P6 to R8SS-P20	10 meters
R3LL-P7 to R7SS-P20	10 meters

Table 22-36 (Cont.) Leaf Switch Connections for the Third Rack in a Eight-Rack System

Connection	Cable Length
R3LL-P30 to R6SS-P20	10 meters
R3LL-P31 to R2SS-P20	5 meters
R3LL-P32 to R3SS-P20	3 meters
R3LL-P33 to R4SS-P20	5 meters
R3LL-P34 to R3SS-P32	3 meters
R3LL-P35 to R6SS-P27	10 meters
R3LL-P36 to R8SS-P27	10 meters

Table 22-37 Leaf Switch Connections for the Fourth Rack in a Eight-Rack System

Connection	Cable Length
R4UL-P1 to R8SS-P23	10 meters
R4UL-P2 to R2SS-P31	5 meters
R4UL-P3 to R1SS-P31	10 meters
R4UL-P4 to R8SS-P10	10 meters
R4UL-P5 to R5SS-P10	5 meters
R4UL-P6 to R2SS-P10	5 meters
R4UL-P7 to R1SS-P10	10 meters
R4UL-P30 to R6SS-P10	5 meters
R4UL-P31 to R7SS-P10	10 meters
R4UL-P32 to R4SS-P10	3 meters
R4UL-P33 to R3SS-P10	5 meters
R4UL-P34 to R4SS-P31	3 meters
R4UL-P35 to R3SS-P31	5 meters
R4UL-P36 to R7SS-P23	10 meters
R4LL-P1 to R5SS-P27	5 meters
R4LL-P2 to R2SS-P27	5 meters
R4LL-P3 to R1SS-P27	10 meters
R4LL-P4 to R2SS-P17	5 meters
R4LL-P5 to R1SS-P17	10 meters
R4LL-P6 to R6SS-P17	5 meters
R4LL-P7 to R5SS-P17	5 meters
R4LL-P30 to R4SS-P17	3 meters
R4LL-P31 to R3SS-P17	5 meters
R4LL-P32 to R7SS-P17	10 meters
R4LL-P33 to R8SS-P17	10 meters
R4LL-P34 to R3SS-P27	5 meters
R4LL-P35 to R4SS-P27	3 meters
R4LL-P36 to R8SS-P24	10 meters

Table 22-38 Leaf Switch Connections for the Fifth Rack in a Eight-Rack System

Connection	Cable Length
R5UL-P1 to R3SS-P28	5 meters
R5UL-P2 to R2SS-P28	10 meters
R5UL-P3 to R1SS-P28	10 meters
R5UL-P4 to R2SS-P8	10 meters
R5UL-P5 to R1SS-P8	10 meters
R5UL-P6 to R3SS-P8	5 meters
R5UL-P7 to R6SS-P7	5 meters
R5UL-P30 to R5SS-P7	3 meters
R5UL-P31 to R8SS-P7	10 meters
R5UL-P32 to R7SS-P7	5 meters
R5UL-P33 to R4SS-P12	5 meters
R5UL-P34 to R4SS-P28	5 meters
R5UL-P35 to R5SS-P28	3 meters
R5UL-P36 to R6SS-P28	5 meters
R5LL-P1 to R7SS-P22	5 meters
R5LL-P2 to R6SS-P22	5 meters
R5LL-P3 to R2SS-P22	10 meters
R5LL-P4 to R6SS-P15	5 meters
R5LL-P5 to R7SS-P15	5 meters
R5LL-P6 to R2SS-P18	10 meters
R5LL-P7 to R1SS-P18	10 meters
R5LL-P30 to R5SS-P15	3 meters
R5LL-P31 to R3SS-P16	5 meters
R5LL-P32 to R8SS-P15	10 meters
R5LL-P33 to R4SS-P19	5 meters
R5LL-P34 to R3SS-P22	5 meters
R5LL-P35 to R4SS-P22	5 meters
R5LL-P36 to R5SS-P22	3 meters

Table 22-39 Leaf Switch Connections for the Sixth Rack in a Eight-Rack System

Connection	Cable Length
R6UL-P1 to R4SS-P30	5 meters
R6UL-P2 to R3SS-P30	10 meters
R6UL-P3 to R7SS-P30	5 meters
R6UL-P4 to R2SS-P16	10 meters
R6UL-P5 to R7SS-P16	5 meters
R6UL-P6 to R3SS-P19	10 meters
R6UL-P7 to R1SS-P16	10 meters
R6UL-P30 to R6SS-P16	3 meters
R6UL-P31 to R4SS-P15	5 meters
R6UL-P32 to R8SS-P16	5 meters
R6UL-P33 to R5SS-P18	5 meters

Table 22-39 (Cont.) Leaf Switch Connections for the Sixth Rack in a Eight-Rack System

Connection	Cable Length
R6UL-P34 to R5SS-P30	5 meters
R6UL-P35 to R6SS-P30	3 meters
R6UL-P36 to R8SS-P30	5 meters
R6LL-P1 to R8SS-P26	5 meters
R6LL-P2 to R7SS-P26	5 meters
R6LL-P3 to R1SS-P24	10 meters
R6LL-P4 to R2SS-P12	10 meters
R6LL-P5 to R1SS-P12	10 meters
R6LL-P6 to R7SS-P8	5 meters
R6LL-P7 to R8SS-P8	5 meters
R6LL-P30 to R5SS-P8	5 meters
R6LL-P31 to R3SS-P7	10 meters
R6LL-P32 to R6SS-P8	3 meters
R6LL-P33 to R4SS-P9	5 meters
R6LL-P34 to R4SS-P26	5 meters
R6LL-P35 to R5SS-P26	5 meters
R6LL-P36 to R6SS-P26	3 meters

Table 22-40 Leaf Switch Connections for the Seventh Rack in a Eight-Rack System

Connection	Cable Length
R7UL-P1 to R2SS-P30	10 meters
R7UL-P2 to R5SS-P21	5 meters
R7UL-P3 to R1SS-P30	10 meters
R7UL-P4 to R3SS-P9	10 meters
R7UL-P5 to R1SS-P9	10 meters
R7UL-P6 to R4SS-P8	10 meters
R7UL-P7 to R2SS-P7	10 meters
R7UL-P30 to R7SS-P9	3 meters
R7UL-P31 to R5SS-P9	5 meters
R7UL-P32 to R8SS-P9	5 meters
R7UL-P33 to R6SS-P9	5 meters
R7UL-P34 to R8SS-P21	5 meters
R7UL-P35 to R6SS-P21	5 meters
R7UL-P36 to R7SS-P21	3 meters
R7LL-P1 to R2SS-P26	10 meters
R7LL-P2 to R3SS-P26	10 meters
R7LL-P3 to R1SS-P26	10 meters
R7LL-P4 to R3SS-P15	10 meters
R7LL-P5 to R1SS-P19	10 meters
R7LL-P6 to R4SS-P18	10 meters
R7LL-P7 to R2SS-P19	10 meters
R7LL-P30 to R7SS-P18	3 meters

Table 22-40 (Cont.) Leaf Switch Connections for the Seventh Rack in a Eight-Rack System

Connection	Cable Length
R7LL-P31 to R5SS-P16	5 meters
R7LL-P32 to R8SS-P12	5 meters
R7LL-P33 to R6SS-P18	5 meters
R7LL-P34 to R6SS-P29	5 meters
R7LL-P35 to R7SS-P29	3 meters
R7LL-P36 to R8SS-P29	5 meters

Table 22-41 Leaf Switch Connections for the Eighth Rack in a Eight-Rack System

Connection	Cable Length
R8UL-P1 to R2SS-P21	10 meters
R8UL-P2 to R3SS-P21	10 meters
R8UL-P3 to R1SS-P21	10 meters
R8UL-P4 to R3SS-P18	10 meters
R8UL-P5 to R1SS-P15	10 meters
R8UL-P6 to R4SS-P16	10 meters
R8UL-P7 to R2SS-P15	10 meters
R8UL-P30 to R7SS-P19	5 meters
R8UL-P31 to R5SS-P19	10 meters
R8UL-P32 to R8SS-P19	3 meters
R8UL-P33 to R6SS-P19	5 meters
R8UL-P34 to R7SS-P28	5 meters
R8UL-P35 to R8SS-P28	3 meters
R8UL-P36 to R4SS-P21	10 meters
R8LL-P1 to R2SS-P29	10 meters
R8LL-P2 to R3SS-P29	10 meters
R8LL-P3 to R1SS-P29	10 meters
R8LL-P4 to R3SS-P12	10 meters
R8LL-P5 to R1SS-P7	10 meters
R8LL-P6 to R4SS-P7	10 meters
R8LL-P7 to R2SS-P9	10 meters
R8LL-P30 to R7SS-P12	5 meters
R8LL-P31 to R5SS-P12	10 meters
R8LL-P32 to R8SS-P18	3 meters
R8LL-P33 to R6SS-P12	5 meters
R8LL-P34 to R8SS-P22	3 meters
R8LL-P35 to R5SS-P29	10 meters
R8LL-P36 to R4SS-P29	10 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-42 Eight-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R8LL-P5	10 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R7UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R6LL-P5	10 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R8UL-P5	10 meters
R1SS-P16	R6UL-P7	10 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R7LL-P5	10 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R8UL-P3	10 meters
R1SS-P22	R2LL-P3	5 meters
R1SS-P23	R1UL-P3	3 meters
R1SS-P24	R6LL-P3	10 meters
R1SS-P25	R3UL-P3	5 meters
R1SS-P26	R7LL-P3	10 meters
R1SS-P27	R4LL-P3	10 meters
R1SS-P28	R5UL-P3	10 meters
R1SS-P29	R8LL-P3	10 meters
R1SS-P30	R7UL-P3	10 meters
R1SS-P31	R4UL-P3	10 meters
R1SS-P32	R3LL-P2	5 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R7UL-P7	10 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R8LL-P7	10 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R6LL-P4	10 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R8UL-P7	10 meters
R2SS-P16	R6UL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters

Table 22-42 (Cont.) Eight-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R2SS-P19	R7LL-P7	10 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R8UL-P1	10 meters
R2SS-P22	R5LL-P3	10 meters
R2SS-P23	R1UL-P34	5 meters
R2SS-P24	R1LL-P1	5 meters
R2SS-P25	R3UL-P2	5 meters
R2SS-P26	R7LL-P1	10 meters
R2SS-P27	R4LL-P2	5 meters
R2SS-P28	R5UL-P2	10 meters
R2SS-P29	R8LL-P1	10 meters
R2SS-P30	R7UL-P1	10 meters
R2SS-P31	R4UL-P2	5 meters
R2SS-P32	R3LL-P3	5 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R6LL-P31	10 meters
R3SS-P8	R5UL-P6	5 meters
R3SS-P9	R7UL-P4	10 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R8LL-P4	10 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R7LL-P4	10 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R8UL-P4	10 meters
R3SS-P19	R6UL-P6	10 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R8UL-P2	10 meters
R3SS-P22	R5LL-P34	5 meters
R3SS-P23	R1UL-P1	5 meters
R3SS-P24	R1LL-P2	5 meters
R3SS-P25	R2UL-P2	5 meters
R3SS-P26	R7LL-P2	10 meters
R3SS-P27	R4LL-P34	5 meters
R3SS-P28	R5UL-P1	5 meters
R3SS-P29	R8LL-P2	10 meters
R3SS-P30	R6UL-P2	10 meters
R3SS-P31	R4UL-P35	5 meters
R3SS-P32	R3LL-P34	3 meters

Table 22-42 (Cont.) Eight-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R8LL-P6	10 meters
R4SS-P8	R7UL-P6	10 meters
R4SS-P9	R6LL-P33	5 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R5UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R6UL-P31	5 meters
R4SS-P16	R8UL-P6	10 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R7LL-P6	10 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R8UL-P36	10 meters
R4SS-P22	R5LL-P35	5 meters
R4SS-P23	R1UL-P2	10 meters
R4SS-P24	R1LL-P34	10 meters
R4SS-P25	R2UL-P3	5 meters
R4SS-P26	R6LL-P34	5 meters
R4SS-P27	R4LL-P35	3 meters
R4SS-P28	R5UL-P34	5 meters
R4SS-P29	R8LL-P36	10 meters
R4SS-P30	R6UL-P1	5 meters
R4SS-P31	R4UL-P34	3 meters
R4SS-P32	R2LL-P2	5 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R6LL-P30	5 meters
R5SS-P9	R7UL-P31	5 meters
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R8LL-P31	10 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R7LL-P31	5 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R6UL-P33	5 meters

Table 22-42 (Cont.) Eight-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R5SS-P19	R8UL-P31	10 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R7UL-P2	5 meters
R5SS-P22	R5LL-P36	3 meters
R5SS-P23	R1UL-P35	10 meters
R5SS-P24	R1LL-P3	10 meters
R5SS-P25	R2UL-P1	10 meters
R5SS-P26	R6LL-P35	5 meters
R5SS-P27	R4LL-P1	5 meters
R5SS-P28	R5UL-P35	3 meters
R5SS-P29	R8LL-P35	10 meters
R5SS-P30	R6UL-P34	5 meters
R5SS-P31	R3UL-P34	5 meters
R5SS-P32	R2LL-P35	10 meters
R6SS-P5	R1UL-P32	10 meters
R6SS-P6	R2UL-P33	10 meters
R6SS-P7	R5UL-P7	5 meters
R6SS-P8	R6LL-P32	3 meters
R6SS-P9	R7UL-P33	5 meters
R6SS-P10	R4UL-P30	5 meters
R6SS-P11	R3UL-P7	10 meters
R6SS-P12	R8LL-P33	5 meters
R6SS-P13	R1LL-P32	10 meters
R6SS-P14	R2LL-P7	10 meters
R6SS-P15	R5LL-P4	5 meters
R6SS-P16	R6UL-P30	3 meters
R6SS-P17	R4LL-P6	5 meters
R6SS-P18	R7LL-P33	5 meters
R6SS-P19	R8UL-P33	5 meters
R6SS-P20	R3LL-P30	10 meters
R6SS-P21	R7UL-P35	5 meters
R6SS-P22	R5LL-P2	5 meters
R6SS-P23	R1UL-P36	10 meters
R6SS-P24	R1LL-P35	10 meters
R6SS-P25	R2UL-P35	10 meters
R6SS-P26	R6LL-P36	3 meters
R6SS-P27	R3LL-P35	10 meters
R6SS-P28	R5UL-P36	5 meters
R6SS-P29	R7LL-P34	5 meters
R6SS-P30	R6UL-P35	3 meters
R6SS-P31	R3UL-P1	10 meters
R6SS-P32	R2LL-P1	10 meters

Table 22-42 (Cont.) Eight-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R7SS-P5	R1UL-P33	10 meters
R7SS-P6	R2UL-P32	10 meters
R7SS-P7	R5UL-P32	5 meters
R7SS-P8	R6LL-P6	5 meters
R7SS-P9	R7UL-P30	3 meters
R7SS-P10	R4UL-P31	10 meters
R7SS-P11	R3UL-P31	10 meters
R7SS-P12	R8LL-P30	5 meters
R7SS-P13	R1LL-P6	10 meters
R7SS-P14	R2LL-P32	10 meters
R7SS-P15	R5LL-P5	5 meters
R7SS-P16	R6UL-P5	5 meters
R7SS-P17	R4LL-P32	10 meters
R7SS-P18	R7LL-P30	3 meters
R7SS-P19	R8UL-P30	5 meters
R7SS-P20	R3LL-P7	10 meters
R7SS-P21	R7UL-P36	3 meters
R7SS-P22	R5LL-P1	5 meters
R7SS-P23	R4UL-P36	10 meters
R7SS-P24	R1LL-P36	10 meters
R7SS-P25	R2UL-P34	10 meters
R7SS-P26	R6LL-P2	5 meters
R7SS-P27	R3LL-P1	10 meters
R7SS-P28	R8UL-P34	5 meters
R7SS-P29	R7LL-P35	3 meters
R7SS-P30	R6UL-P3	5 meters
R7SS-P31	R3UL-P35	10 meters
R7SS-P32	R2LL-P34	10 meters
R8SS-P5	R1UL-P7	10 meters
R8SS-P6	R2UL-P30	10 meters
R8SS-P7	R5UL-P31	10 meters
R8SS-P8	R6LL-P7	5 meters
R8SS-P9	R7UL-P32	5 meters
R8SS-P10	R4UL-P4	10 meters
R8SS-P11	R3UL-P33	10 meters
R8SS-P12	R7LL-P32	5 meters
R8SS-P13	R1LL-P33	10 meters
R8SS-P14	R2LL-P30	10 meters
R8SS-P15	R5LL-P32	10 meters
R8SS-P16	R6UL-P32	5 meters
R8SS-P17	R4LL-P33	10 meters
R8SS-P18	R8LL-P32	3 meters

Table 22-42 (Cont.) Eight-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R8SS-P19	R8UL-P32	3 meters
R8SS-P20	R3LL-P6	10 meters
R8SS-P21	R7UL-P34	5 meters
R8SS-P22	R8LL-P34	3 meters
R8SS-P23	R4UL-P1	10 meters
R8SS-P24	R4LL-P36	10 meters
R8SS-P25	R2UL-P36	10 meters
R8SS-P26	R6LL-P1	5 meters
R8SS-P27	R3LL-P36	10 meters
R8SS-P28	R8UL-P35	3 meters
R8SS-P29	R7LL-P36	5 meters
R8SS-P30	R6UL-P36	5 meters
R8SS-P31	R3UL-P36	10 meters
R8SS-P32	R2LL-P36	10 meters

Nine-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect nine racks using RoCE Network Fabric.

 Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, for these systems only, ignore the connections to port number 30 on every leaf switch in the following tables. This adjustment leaves only 13 inter-switch links on each leaf switch and only applies to X9M-8 systems with three database servers and 11 storage servers.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a 9 rack system.

Table 22-43 Leaf Switch Connections for the First Rack in a 9 Rack System

Connection	Cable Length
R1UL-P1 to R3SS-P23	5 meters
R1UL-P2 to R4SS-P23	10 meters
R1UL-P3 to R1SS-P23	3 meters
R1UL-P4 to R4SS-P5	10 meters
R1UL-P5 to R1SS-P5	3 meters
R1UL-P6 to R3SS-P5	5 meters
R1UL-P7 to R8SS-P5	10 meters
R1UL-P30 to R5SS-P5	10 meters
R1UL-P31 to R2SS-P5	5 meters
R1UL-P32 to R6SS-P5	10 meters
R1UL-P33 to R7SS-P5	10 meters

Table 22-43 (Cont.) Leaf Switch Connections for the First Rack in a 9 Rack System

Connection	Cable Length
R1UL-P34 to R2SS-P23	5 meters
R1UL-P35 to R5SS-P23	10 meters
R1UL-P36 to R9SS-P5	10 meters
R1LL-P1 to R2SS-P25	5 meters
R1LL-P2 to R3SS-P25	5 meters
R1LL-P3 to R5SS-P25	10 meters
R1LL-P4 to R3SS-P13	5 meters
R1LL-P5 to R1SS-P13	3 meters
R1LL-P6 to R7SS-P13	10 meters
R1LL-P7 to R5SS-P13	10 meters
R1LL-P30 to R4SS-P13	10 meters
R1LL-P31 to R2SS-P13	5 meters
R1LL-P32 to R6SS-P13	10 meters
R1LL-P33 to R8SS-P13	10 meters
R1LL-P34 to R4SS-P25	10 meters
R1LL-P35 to R6SS-P25	10 meters
R1LL-P36 to R9SS-P13	10 meters

Table 22-44 Leaf Switch Connections for the Second Rack in a 9 Rack System

Connection	Cable Length
R2UL-P1 to R5SS-P31	10 meters
R2UL-P2 to R3SS-P31	5 meters
R2UL-P3 to R4SS-P31	5 meters
R2UL-P4 to R4SS-P6	5 meters
R2UL-P5 to R1SS-P6	5 meters
R2UL-P6 to R3SS-P6	5 meters
R2UL-P7 to R5SS-P6	10 meters
R2UL-P30 to R8SS-P6	10 meters
R2UL-P31 to R2SS-P6	3 meters
R2UL-P32 to R7SS-P6	10 meters
R2UL-P33 to R6SS-P6	10 meters
R2UL-P34 to R7SS-P31	10 meters
R2UL-P35 to R6SS-P31	10 meters
R2UL-P36 to R9SS-P6	10 meters
R2LL-P1 to R6SS-P27	10 meters
R2LL-P2 to R4SS-P27	5 meters
R2LL-P3 to R9SS-P14	10 meters
R2LL-P4 to R4SS-P14	5 meters
R2LL-P5 to R1SS-P14	5 meters
R2LL-P6 to R3SS-P14	5 meters
R2LL-P7 to R6SS-P14	10 meters
R2LL-P30 to R8SS-P14	10 meters

Table 22-44 (Cont.) Leaf Switch Connections for the Second Rack in a 9 Rack System

Connection	Cable Length
R2LL-P31 to R2SS-P14	3 meters
R2LL-P32 to R7SS-P14	10 meters
R2LL-P33 to R5SS-P14	10 meters
R2LL-P34 to R7SS-P27	10 meters
R2LL-P35 to R5SS-P27	10 meters
R2LL-P36 to R8SS-P27	10 meters

Table 22-45 Leaf Switch Connections for the Third Rack in a 9 Rack System

Connection	Cable Length
R3UL-P1 to R6SS-P28	10 meters
R3UL-P2 to R9SS-P11	10 meters
R3UL-P3 to R9SS-P28	10 meters
R3UL-P4 to R4SS-P11	5 meters
R3UL-P5 to R1SS-P11	5 meters
R3UL-P6 to R2SS-P11	5 meters
R3UL-P7 to R6SS-P11	10 meters
R3UL-P30 to R3SS-P11	3 meters
R3UL-P31 to R7SS-P11	10 meters
R3UL-P32 to R5SS-P11	5 meters
R3UL-P33 to R8SS-P11	10 meters
R3UL-P34 to R5SS-P28	5 meters
R3UL-P35 to R7SS-P28	10 meters
R3UL-P36 to R8SS-P28	10 meters
R3LL-P1 to R7SS-P23	10 meters
R3LL-P2 to R1SS-P32	5 meters
R3LL-P3 to R9SS-P20	10 meters
R3LL-P4 to R1SS-P20	5 meters
R3LL-P5 to R5SS-P20	5 meters
R3LL-P6 to R8SS-P20	10 meters
R3LL-P7 to R7SS-P20	10 meters
R3LL-P30 to R6SS-P20	10 meters
R3LL-P31 to R2SS-P20	5 meters
R3LL-P32 to R3SS-P20	3 meters
R3LL-P33 to R4SS-P20	5 meters
R3LL-P34 to R9SS-P23	10 meters
R3LL-P35 to R6SS-P23	10 meters
R3LL-P36 to R8SS-P23	10 meters

Table 22-46 Leaf Switch Connections for the Fourth Rack in a 9 Rack System

Connection	Cable Length
R4UL-P1 to R8SS-P25	10 meters

Table 22-46 (Cont.) Leaf Switch Connections for the Fourth Rack in a 9 Rack System

Connection	Cable Length
R4UL-P2 to R2SS-P31	5 meters
R4UL-P3 to R1SS-P31	10 meters
R4UL-P4 to R8SS-P10	10 meters
R4UL-P5 to R5SS-P10	5 meters
R4UL-P6 to R2SS-P10	5 meters
R4UL-P7 to R1SS-P10	10 meters
R4UL-P30 to R6SS-P10	5 meters
R4UL-P31 to R7SS-P10	10 meters
R4UL-P32 to R4SS-P10	3 meters
R4UL-P33 to R3SS-P10	5 meters
R4UL-P34 to R9SS-P10	10 meters
R4UL-P35 to R9SS-P25	10 meters
R4UL-P36 to R7SS-P25	10 meters
R4LL-P1 to R9SS-P17	10 meters
R4LL-P2 to R2SS-P27	5 meters
R4LL-P3 to R1SS-P27	10 meters
R4LL-P4 to R2SS-P17	5 meters
R4LL-P5 to R1SS-P17	10 meters
R4LL-P6 to R6SS-P17	5 meters
R4LL-P7 to R5SS-P17	5 meters
R4LL-P30 to R4SS-P17	3 meters
R4LL-P31 to R3SS-P17	5 meters
R4LL-P32 to R7SS-P17	10 meters
R4LL-P33 to R8SS-P17	10 meters
R4LL-P34 to R3SS-P27	5 meters
R4LL-P35 to R9SS-P31	10 meters
R4LL-P36 to R8SS-P31	10 meters

Table 22-47 Leaf Switch Connections for the Fifth Rack in a 9 Rack System

Connection	Cable Length
R5UL-P1 to R3SS-P28	5 meters
R5UL-P2 to R2SS-P28	10 meters
R5UL-P3 to R1SS-P28	10 meters
R5UL-P4 to R2SS-P8	10 meters
R5UL-P5 to R1SS-P8	10 meters
R5UL-P6 to R3SS-P8	5 meters
R5UL-P7 to R6SS-P7	5 meters
R5UL-P30 to R5SS-P7	3 meters
R5UL-P31 to R8SS-P7	10 meters
R5UL-P32 to R7SS-P7	5 meters
R5UL-P33 to R4SS-P12	5 meters
R5UL-P34 to R4SS-P28	5 meters

Table 22-47 (Cont.) Leaf Switch Connections for the Fifth Rack in a 9 Rack System

Connection	Cable Length
R5UL-P35 to R9SS-P7	10 meters
R5UL-P36 to R9SS-P27	10 meters
R5LL-P1 to R9SS-P15	10 meters
R5LL-P2 to R1SS-P22	10 meters
R5LL-P3 to R2SS-P22	10 meters
R5LL-P4 to R6SS-P15	5 meters
R5LL-P5 to R7SS-P15	5 meters
R5LL-P6 to R2SS-P18	10 meters
R5LL-P7 to R1SS-P18	10 meters
R5LL-P30 to R5SS-P15	3 meters
R5LL-P31 to R3SS-P16	5 meters
R5LL-P32 to R8SS-P15	10 meters
R5LL-P33 to R4SS-P19	5 meters
R5LL-P34 to R3SS-P22	5 meters
R5LL-P35 to R4SS-P22	5 meters
R5LL-P36 to R5SS-P22	3 meters

Table 22-48 Leaf Switch Connections for the Sixth Rack in a 9 Rack System

Connection	Cable Length
R6UL-P1 to R4SS-P32	5 meters
R6UL-P2 to R3SS-P32	10 meters
R6UL-P3 to R9SS-P16	10 meters
R6UL-P4 to R2SS-P16	10 meters
R6UL-P5 to R7SS-P16	5 meters
R6UL-P6 to R3SS-P19	10 meters
R6UL-P7 to R1SS-P16	10 meters
R6UL-P30 to R6SS-P16	3 meters
R6UL-P31 to R4SS-P15	5 meters
R6UL-P32 to R8SS-P16	5 meters
R6UL-P33 to R5SS-P18	5 meters
R6UL-P34 to R5SS-P32	5 meters
R6UL-P35 to R6SS-P32	3 meters
R6UL-P36 to R2SS-P32	10 meters
R6LL-P1 to R9SS-P8	10 meters
R6LL-P2 to R7SS-P21	5 meters
R6LL-P3 to R3SS-P21	10 meters
R6LL-P4 to R2SS-P12	10 meters
R6LL-P5 to R1SS-P12	10 meters
R6LL-P6 to R7SS-P8	5 meters
R6LL-P7 to R8SS-P8	5 meters
R6LL-P30 to R5SS-P8	5 meters
R6LL-P31 to R3SS-P7	10 meters

Table 22-48 (Cont.) Leaf Switch Connections for the Sixth Rack in a 9 Rack System

Connection	Cable Length
R6LL-P32 to R6SS-P8	3 meters
R6LL-P33 to R4SS-P9	5 meters
R6LL-P34 to R4SS-P21	5 meters
R6LL-P35 to R5SS-P21	5 meters
R6LL-P36 to R6SS-P21	3 meters

Table 22-49 Leaf Switch Connections for the Seventh Rack in a 9 Rack System

Connection	Cable Length
R7UL-P1 to R5SS-P26	5 meters
R7UL-P2 to R9SS-P9	5 meters
R7UL-P3 to R4SS-P26	10 meters
R7UL-P4 to R3SS-P9	10 meters
R7UL-P5 to R1SS-P9	10 meters
R7UL-P6 to R4SS-P8	10 meters
R7UL-P7 to R2SS-P7	10 meters
R7UL-P30 to R7SS-P9	3 meters
R7UL-P31 to R5SS-P9	5 meters
R7UL-P32 to R8SS-P9	5 meters
R7UL-P33 to R6SS-P9	5 meters
R7UL-P34 to R8SS-P26	5 meters
R7UL-P35 to R6SS-P26	5 meters
R7UL-P36 to R7SS-P26	3 meters
R7LL-P1 to R9SS-P12	5 meters
R7LL-P2 to R9SS-P30	5 meters
R7LL-P3 to R5SS-P30	5 meters
R7LL-P4 to R3SS-P15	10 meters
R7LL-P5 to R1SS-P19	10 meters
R7LL-P6 to R4SS-P18	10 meters
R7LL-P7 to R2SS-P19	10 meters
R7LL-P30 to R7SS-P18	3 meters
R7LL-P31 to R5SS-P16	5 meters
R7LL-P32 to R8SS-P12	5 meters
R7LL-P33 to R6SS-P18	5 meters
R7LL-P34 to R6SS-P30	5 meters
R7LL-P35 to R7SS-P30	3 meters
R7LL-P36 to R8SS-P30	5 meters

Table 22-50 Leaf Switch Connections for the Eighth Rack in a 9 Rack System

Connection	Cable Length
R8UL-P1 to R9SS-P19	5 meters
R8UL-P2 to R9SS-P22	5 meters

Table 22-50 (Cont.) Leaf Switch Connections for the Eighth Rack in a 9 Rack System

Connection	Cable Length
R8UL-P3 to R1SS-P25	10 meters
R8UL-P4 to R3SS-P18	10 meters
R8UL-P5 to R1SS-P15	10 meters
R8UL-P6 to R4SS-P16	10 meters
R8UL-P7 to R2SS-P15	10 meters
R8UL-P30 to R7SS-P19	5 meters
R8UL-P31 to R5SS-P19	10 meters
R8UL-P32 to R8SS-P19	3 meters
R8UL-P33 to R6SS-P19	5 meters
R8UL-P34 to R7SS-P22	5 meters
R8UL-P35 to R8SS-P22	3 meters
R8UL-P36 to R6SS-P22	5 meters
R8LL-P1 to R2SS-P21	10 meters
R8LL-P2 to R9SS-P18	5 meters
R8LL-P3 to R1SS-P21	10 meters
R8LL-P4 to R3SS-P12	10 meters
R8LL-P5 to R1SS-P7	10 meters
R8LL-P6 to R4SS-P7	10 meters
R8LL-P7 to R2SS-P9	10 meters
R8LL-P30 to R7SS-P12	5 meters
R8LL-P31 to R5SS-P12	10 meters
R8LL-P32 to R8SS-P18	3 meters
R8LL-P33 to R6SS-P12	5 meters
R8LL-P34 to R8SS-P32	3 meters
R8LL-P35 to R9SS-P32	5 meters
R8LL-P36 to R7SS-P32	5 meters

Table 22-51 Leaf Switch Connections for the Ninth Rack in a 9 Rack System

Connection	Cable Length
R9UL-P1 to R1SS-P24	10 meters
R9UL-P2 to R1SS-P26	10 meters
R9UL-P3 to R2SS-P24	10 meters
R9UL-P4 to R2SS-P26	10 meters
R9UL-P5 to R3SS-P24	10 meters
R9UL-P6 to R3SS-P26	10 meters
R9UL-P7 to R4SS-P24	10 meters
R9UL-P30 to R5SS-P24	10 meters
R9UL-P31 to R6SS-P24	10 meters
R9UL-P32 to R7SS-P24	5 meters
R9UL-P33 to R8SS-P21	5 meters
R9UL-P34 to R8SS-P24	5 meters
R9UL-P35 to R9SS-P21	3 meters

Table 22-51 (Cont.) Leaf Switch Connections for the Ninth Rack in a 9 Rack System

Connection	Cable Length
R9UL-P36 to R9SS-P24	3 meters
R9LL-P1 to R1SS-P29	10 meters
R9LL-P2 to R1SS-P30	10 meters
R9LL-P3 to R2SS-P29	10 meters
R9LL-P4 to R2SS-P30	10 meters
R9LL-P5 to R3SS-P29	10 meters
R9LL-P6 to R3SS-P30	10 meters
R9LL-P7 to R4SS-P29	10 meters
R9LL-P30 to R4SS-P30	10 meters
R9LL-P31 to R5SS-P29	10 meters
R9LL-P32 to R6SS-P29	10 meters
R9LL-P33 to R7SS-P29	5 meters
R9LL-P34 to R8SS-P29	5 meters
R9LL-P35 to R9SS-P26	3 meters
R9LL-P36 to R9SS-P29	3 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-52 Nine-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R8LL-P5	10 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R7UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R6LL-P5	10 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R8UL-P5	10 meters
R1SS-P16	R6UL-P7	10 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R7LL-P5	10 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R8LL-P3	10 meters
R1SS-P22	R5LL-P2	10 meters
R1SS-P23	R1UL-P3	3 meters
R1SS-P24	R9UL-P1	10 meters
R1SS-P25	R8UL-P3	10 meters

Table 22-52 (Cont.) Nine-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P26	R9UL-P2	10 meters
R1SS-P27	R4LL-P3	10 meters
R1SS-P28	R5UL-P3	10 meters
R1SS-P29	R9LL-P1	10 meters
R1SS-P30	R9LL-P2	10 meters
R1SS-P31	R4UL-P3	10 meters
R1SS-P32	R3LL-P2	5 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R7UL-P7	10 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R8LL-P7	10 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R6LL-P4	10 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R8UL-P7	10 meters
R2SS-P16	R6UL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters
R2SS-P19	R7LL-P7	10 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R8LL-P1	10 meters
R2SS-P22	R5LL-P3	10 meters
R2SS-P23	R1UL-P34	5 meters
R2SS-P24	R9UL-P3	10 meters
R2SS-P25	R1LL-P1	5 meters
R2SS-P26	R9UL-P4	10 meters
R2SS-P27	R4LL-P2	5 meters
R2SS-P28	R5UL-P2	10 meters
R2SS-P29	R9LL-P3	10 meters
R2SS-P30	R9LL-P4	10 meters
R2SS-P31	R4UL-P2	5 meters
R2SS-P32	R6UL-P36	10 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R6LL-P31	10 meters
R3SS-P8	R5UL-P6	5 meters
R3SS-P9	R7UL-P4	10 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters

Table 22-52 (Cont.) Nine-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R3SS-P12	R8LL-P4	10 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R7LL-P4	10 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R8UL-P4	10 meters
R3SS-P19	R6UL-P6	10 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R6LL-P3	10 meters
R3SS-P22	R5LL-P34	5 meters
R3SS-P23	R1UL-P1	5 meters
R3SS-P24	R9UL-P5	10 meters
R3SS-P25	R1LL-P2	5 meters
R3SS-P26	R9UL-P6	10 meters
R3SS-P27	R4LL-P34	5 meters
R3SS-P28	R5UL-P1	5 meters
R3SS-P29	R9LL-P5	10 meters
R3SS-P30	R9LL-P6	10 meters
R3SS-P31	R2UL-P2	5 meters
R3SS-P32	R6UL-P2	10 meters
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R8LL-P6	10 meters
R4SS-P8	R7UL-P6	10 meters
R4SS-P9	R6LL-P33	5 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R5UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R6UL-P31	5 meters
R4SS-P16	R8UL-P6	10 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R7LL-P6	10 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R6LL-P34	5 meters
R4SS-P22	R5LL-P35	5 meters
R4SS-P23	R1UL-P2	10 meters
R4SS-P24	R9UL-P7	10 meters
R4SS-P25	R1LL-P34	10 meters

Table 22-52 (Cont.) Nine-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R4SS-P26	R7UL-P3	10 meters
R4SS-P27	R2LL-P2	5 meters
R4SS-P28	R5UL-P34	5 meters
R4SS-P29	R9LL-P7	10 meters
R4SS-P30	R9LL-P30	10 meters
R4SS-P31	R2UL-P3	5 meters
R4SS-P32	R6UL-P1	5 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R6LL-P30	5 meters
R5SS-P9	R7UL-P31	5 meters
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R8LL-P31	10 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R7LL-P31	5 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R6UL-P33	5 meters
R5SS-P19	R8UL-P31	10 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R6LL-P35	5 meters
R5SS-P22	R5LL-P36	3 meters
R5SS-P23	R1UL-P35	10 meters
R5SS-P24	R9UL-P30	10 meters
R5SS-P25	R1LL-P3	10 meters
R5SS-P26	R7UL-P1	5 meters
R5SS-P27	R2LL-P35	10 meters
R5SS-P28	R3UL-P34	5 meters
R5SS-P29	R9LL-P31	10 meters
R5SS-P30	R7LL-P3	5 meters
R5SS-P31	R2UL-P1	10 meters
R5SS-P32	R6UL-P34	5 meters
R6SS-P5	R1UL-P32	10 meters
R6SS-P6	R2UL-P33	10 meters
R6SS-P7	R5UL-P7	5 meters
R6SS-P8	R6LL-P32	3 meters
R6SS-P9	R7UL-P33	5 meters
R6SS-P10	R4UL-P30	5 meters
R6SS-P11	R3UL-P7	10 meters

Table 22-52 (Cont.) Nine-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R6SS-P12	R8LL-P33	5 meters
R6SS-P13	R1LL-P32	10 meters
R6SS-P14	R2LL-P7	10 meters
R6SS-P15	R5LL-P4	5 meters
R6SS-P16	R6UL-P30	3 meters
R6SS-P17	R4LL-P6	5 meters
R6SS-P18	R7LL-P33	5 meters
R6SS-P19	R8UL-P33	5 meters
R6SS-P20	R3LL-P30	10 meters
R6SS-P21	R6LL-P36	3 meters
R6SS-P22	R8UL-P36	5 meters
R6SS-P23	R3LL-P35	10 meters
R6SS-P24	R9UL-P31	10 meters
R6SS-P25	R1LL-P35	10 meters
R6SS-P26	R7UL-P35	5 meters
R6SS-P27	R2LL-P1	10 meters
R6SS-P28	R3UL-P1	10 meters
R6SS-P29	R9LL-P32	10 meters
R6SS-P30	R7LL-P34	5 meters
R6SS-P31	R2UL-P35	10 meters
R6SS-P32	R6UL-P35	3 meters
R7SS-P5	R1UL-P33	10 meters
R7SS-P6	R2UL-P32	10 meters
R7SS-P7	R5UL-P32	5 meters
R7SS-P8	R6LL-P6	5 meters
R7SS-P9	R7UL-P30	3 meters
R7SS-P10	R4UL-P31	10 meters
R7SS-P11	R3UL-P31	10 meters
R7SS-P12	R8LL-P30	5 meters
R7SS-P13	R1LL-P6	10 meters
R7SS-P14	R2LL-P32	10 meters
R7SS-P15	R5LL-P5	5 meters
R7SS-P16	R6UL-P5	5 meters
R7SS-P17	R4LL-P32	10 meters
R7SS-P18	R7LL-P30	3 meters
R7SS-P19	R8UL-P30	5 meters
R7SS-P20	R3LL-P7	10 meters
R7SS-P21	R6LL-P2	5 meters
R7SS-P22	R8UL-P34	5 meters
R7SS-P23	R3LL-P1	10 meters
R7SS-P24	R9UL-P32	5 meters
R7SS-P25	R4UL-P36	10 meters

Table 22-52 (Cont.) Nine-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R7SS-P26	R7UL-P36	3 meters
R7SS-P27	R2LL-P34	10 meters
R7SS-P28	R3UL-P35	10 meters
R7SS-P29	R9LL-P33	5 meters
R7SS-P30	R7LL-P35	3 meters
R7SS-P31	R2UL-P34	10 meters
R7SS-P32	R8LL-P36	5 meters
R8SS-P5	R1UL-P7	10 meters
R8SS-P6	R2UL-P30	10 meters
R8SS-P7	R5UL-P31	10 meters
R8SS-P8	R6LL-P7	5 meters
R8SS-P9	R7UL-P32	5 meters
R8SS-P10	R4UL-P4	10 meters
R8SS-P11	R3UL-P33	10 meters
R8SS-P12	R7LL-P32	5 meters
R8SS-P13	R1LL-P33	10 meters
R8SS-P14	R2LL-P30	10 meters
R8SS-P15	R5LL-P32	10 meters
R8SS-P16	R6UL-P32	5 meters
R8SS-P17	R4LL-P33	10 meters
R8SS-P18	R8LL-P32	3 meters
R8SS-P19	R8UL-P32	3 meters
R8SS-P20	R3LL-P6	10 meters
R8SS-P21	R9UL-P33	5 meters
R8SS-P22	R8UL-P35	3 meters
R8SS-P23	R3LL-P36	10 meters
R8SS-P24	R9UL-P34	5 meters
R8SS-P25	R4UL-P1	10 meters
R8SS-P26	R7UL-P34	5 meters
R8SS-P27	R2LL-P36	10 meters
R8SS-P28	R3UL-P36	10 meters
R8SS-P29	R9LL-P34	5 meters
R8SS-P30	R7LL-P36	5 meters
R8SS-P31	R4LL-P36	10 meters
R8SS-P32	R8LL-P34	3 meters
R9SS-P5	R1UL-P36	10 meters
R9SS-P6	R2UL-P36	10 meters
R9SS-P7	R5UL-P35	10 meters
R9SS-P8	R6LL-P1	10 meters
R9SS-P9	R7UL-P2	5 meters
R9SS-P10	R4UL-P34	10 meters
R9SS-P11	R3UL-P2	10 meters

Table 22-52 (Cont.) Nine-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R9SS-P12	R7LL-P1	5 meters
R9SS-P13	R1LL-P36	10 meters
R9SS-P14	R2LL-P3	10 meters
R9SS-P15	R5LL-P1	10 meters
R9SS-P16	R6UL-P3	10 meters
R9SS-P17	R4LL-P1	10 meters
R9SS-P18	R8LL-P2	5 meters
R9SS-P19	R8UL-P1	5 meters
R9SS-P20	R3LL-P3	10 meters
R9SS-P21	R9UL-P35	3 meters
R9SS-P22	R8UL-P2	5 meters
R9SS-P23	R3LL-P34	10 meters
R9SS-P24	R9UL-P36	3 meters
R9SS-P25	R4UL-P35	10 meters
R9SS-P26	R9LL-P35	3 meters
R9SS-P27	R5UL-P36	10 meters
R9SS-P28	R3UL-P3	10 meters
R9SS-P29	R9LL-P36	3 meters
R9SS-P30	R7LL-P2	5 meters
R9SS-P31	R4LL-P35	10 meters
R9SS-P32	R8LL-P35	5 meters

Ten-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect ten racks using RoCE Network Fabric.

 Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, interconnecting 10 or more of these racks requires modification to the following cabling tables. Contact Oracle for further details.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a 10 rack system.

Table 22-53 Leaf Switch Connections for the First Rack in a 10 Rack System

Connection	Cable Length
R1UL-P1 to R3SS-P23	5 meters
R1UL-P2 to R4SS-P23	10 meters
R1UL-P3 to R1SS-P23	3 meters
R1UL-P4 to R4SS-P5	10 meters
R1UL-P5 to R1SS-P5	3 meters
R1UL-P6 to R3SS-P5	5 meters
R1UL-P7 to R8SS-P5	10 meters
R1UL-P30 to R5SS-P5	10 meters
R1UL-P31 to R2SS-P5	5 meters
R1UL-P32 to R6SS-P5	10 meters
R1UL-P33 to R7SS-P5	10 meters
R1UL-P34 to R2SS-P23	5 meters

Table 22-53 (Cont.) Leaf Switch Connections for the First Rack in a 10 Rack System

Connection	Cable Length
R1UL-P35 to R10SS-P5	15 meters
R1UL-P36 to R9SS-P5	10 meters
R1LL-P1 to R2SS-P27	5 meters
R1LL-P2 to R3SS-P27	5 meters
R1LL-P3 to R5SS-P27	10 meters
R1LL-P4 to R3SS-P13	5 meters
R1LL-P5 to R1SS-P13	3 meters
R1LL-P6 to R7SS-P13	10 meters
R1LL-P7 to R5SS-P13	10 meters
R1LL-P30 to R4SS-P13	10 meters
R1LL-P31 to R2SS-P13	5 meters
R1LL-P32 to R6SS-P13	10 meters
R1LL-P33 to R8SS-P13	10 meters
R1LL-P34 to R4SS-P27	10 meters
R1LL-P35 to R10SS-P13	15 meters
R1LL-P36 to R9SS-P13	10 meters

Table 22-54 Leaf Switch Connections for the Second Rack in a 10 Rack System

Connection	Cable Length
R2UL-P1 to R5SS-P28	10 meters
R2UL-P2 to R3SS-P28	5 meters
R2UL-P3 to R4SS-P28	5 meters
R2UL-P4 to R4SS-P6	5 meters
R2UL-P5 to R1SS-P6	5 meters
R2UL-P6 to R3SS-P6	5 meters
R2UL-P7 to R5SS-P6	10 meters
R2UL-P30 to R8SS-P6	10 meters
R2UL-P31 to R2SS-P6	3 meters
R2UL-P32 to R7SS-P6	10 meters
R2UL-P33 to R6SS-P6	10 meters
R2UL-P34 to R10SS-P6	10 meters
R2UL-P35 to R6SS-P28	10 meters
R2UL-P36 to R9SS-P6	10 meters
R2LL-P1 to R6SS-P22	10 meters
R2LL-P2 to R4SS-P22	5 meters
R2LL-P3 to R9SS-P14	10 meters
R2LL-P4 to R4SS-P14	5 meters
R2LL-P5 to R1SS-P14	5 meters
R2LL-P6 to R3SS-P14	5 meters
R2LL-P7 to R6SS-P14	10 meters
R2LL-P30 to R8SS-P14	10 meters
R2LL-P31 to R2SS-P14	3 meters

Table 22-54 (Cont.) Leaf Switch Connections for the Second Rack in a 10 Rack System

Connection	Cable Length
R2LL-P32 to R7SS-P14	10 meters
R2LL-P33 to R5SS-P14	10 meters
R2LL-P34 to R7SS-P22	10 meters
R2LL-P35 to R5SS-P22	10 meters
R2LL-P36 to R10SS-P14	10 meters

Table 22-55 Leaf Switch Connections for the Third Rack in a 10 Rack System

Connection	Cable Length
R3UL-P1 to R6SS-P23	10 meters
R3UL-P2 to R9SS-P11	10 meters
R3UL-P3 to R10SS-P11	10 meters
R3UL-P4 to R4SS-P11	5 meters
R3UL-P5 to R1SS-P11	5 meters
R3UL-P6 to R2SS-P11	5 meters
R3UL-P7 to R6SS-P11	10 meters
R3UL-P30 to R3SS-P11	3 meters
R3UL-P31 to R7SS-P11	10 meters
R3UL-P32 to R5SS-P11	5 meters
R3UL-P33 to R8SS-P11	10 meters
R3UL-P34 to R5SS-P23	5 meters
R3UL-P35 to R7SS-P23	10 meters
R3UL-P36 to R8SS-P23	10 meters
R3LL-P1 to R7SS-P27	10 meters
R3LL-P2 to R10SS-P20	10 meters
R3LL-P3 to R9SS-P20	10 meters
R3LL-P4 to R1SS-P20	5 meters
R3LL-P5 to R5SS-P20	5 meters
R3LL-P6 to R8SS-P20	10 meters
R3LL-P7 to R7SS-P20	10 meters
R3LL-P30 to R6SS-P20	10 meters
R3LL-P31 to R2SS-P20	5 meters
R3LL-P32 to R3SS-P20	3 meters
R3LL-P33 to R4SS-P20	5 meters
R3LL-P34 to R9SS-P27	10 meters
R3LL-P35 to R6SS-P27	10 meters
R3LL-P36 to R8SS-P27	10 meters

Table 22-56 Leaf Switch Connections for the Fourth Rack in a 10 Rack System

Connection	Cable Length
R4UL-P1 to R8SS-P28	10 meters
R4UL-P2 to R10SS-P28	10 meters

Table 22-56 (Cont.) Leaf Switch Connections for the Fourth Rack in a 10 Rack System

Connection	Cable Length
R4UL-P3 to R10SS-P10	10 meters
R4UL-P4 to R8SS-P10	10 meters
R4UL-P5 to R5SS-P10	5 meters
R4UL-P6 to R2SS-P10	5 meters
R4UL-P7 to R1SS-P10	10 meters
R4UL-P30 to R6SS-P10	5 meters
R4UL-P31 to R7SS-P10	10 meters
R4UL-P32 to R4SS-P10	3 meters
R4UL-P33 to R3SS-P10	5 meters
R4UL-P34 to R9SS-P10	10 meters
R4UL-P35 to R9SS-P28	10 meters
R4UL-P36 to R7SS-P28	10 meters
R4LL-P1 to R9SS-P17	10 meters
R4LL-P2 to R10SS-P17	10 meters
R4LL-P3 to R1SS-P27	10 meters
R4LL-P4 to R2SS-P17	5 meters
R4LL-P5 to R1SS-P17	10 meters
R4LL-P6 to R6SS-P17	5 meters
R4LL-P7 to R5SS-P17	5 meters
R4LL-P30 to R4SS-P17	3 meters
R4LL-P31 to R3SS-P17	5 meters
R4LL-P32 to R7SS-P17	10 meters
R4LL-P33 to R8SS-P17	10 meters
R4LL-P34 to R10SS-P22	10 meters
R4LL-P35 to R9SS-P22	10 meters
R4LL-P36 to R8SS-P22	10 meters

Table 22-57 Leaf Switch Connections for the Fifth Rack in a 10 Rack System

Connection	Cable Length
R5UL-P1 to R10SS-P7	10 meters
R5UL-P2 to R2SS-P28	10 meters
R5UL-P3 to R1SS-P28	10 meters
R5UL-P4 to R2SS-P8	10 meters
R5UL-P5 to R1SS-P8	10 meters
R5UL-P6 to R3SS-P8	5 meters
R5UL-P7 to R6SS-P7	5 meters
R5UL-P30 to R5SS-P7	3 meters
R5UL-P31 to R8SS-P7	10 meters
R5UL-P32 to R7SS-P7	5 meters
R5UL-P33 to R4SS-P12	5 meters
R5UL-P34 to R10SS-P23	10 meters
R5UL-P35 to R9SS-P7	10 meters

Table 22-57 (Cont.) Leaf Switch Connections for the Fifth Rack in a 10 Rack System

Connection	Cable Length
R5UL-P36 to R9SS-P23	10 meters
R5LL-P1 to R9SS-P15	10 meters
R5LL-P2 to R1SS-P22	10 meters
R5LL-P3 to R2SS-P22	10 meters
R5LL-P4 to R6SS-P15	5 meters
R5LL-P5 to R7SS-P15	5 meters
R5LL-P6 to R2SS-P18	10 meters
R5LL-P7 to R1SS-P18	10 meters
R5LL-P30 to R5SS-P15	3 meters
R5LL-P31 to R3SS-P16	5 meters
R5LL-P32 to R8SS-P15	10 meters
R5LL-P33 to R4SS-P19	5 meters
R5LL-P34 to R3SS-P22	5 meters
R5LL-P35 to R10SS-P15	10 meters
R5LL-P36 to R10SS-P27	10 meters

Table 22-58 Leaf Switch Connections for the Sixth Rack in a 10 Rack System

Connection	Cable Length
R6UL-P1 to R4SS-P26	5 meters
R6UL-P2 to R3SS-P26	10 meters
R6UL-P3 to R9SS-P16	10 meters
R6UL-P4 to R2SS-P16	10 meters
R6UL-P5 to R7SS-P16	5 meters
R6UL-P6 to R3SS-P19	10 meters
R6UL-P7 to R1SS-P16	10 meters
R6UL-P30 to R6SS-P16	3 meters
R6UL-P31 to R4SS-P15	5 meters
R6UL-P32 to R8SS-P16	5 meters
R6UL-P33 to R5SS-P18	5 meters
R6UL-P34 to R1SS-P26	10 meters
R6UL-P35 to R10SS-P16	10 meters
R6UL-P36 to R2SS-P26	10 meters
R6LL-P1 to R9SS-P8	10 meters
R6LL-P2 to R10SS-P8	10 meters
R6LL-P3 to R3SS-P29	10 meters
R6LL-P4 to R2SS-P12	10 meters
R6LL-P5 to R1SS-P12	10 meters
R6LL-P6 to R7SS-P8	5 meters
R6LL-P7 to R8SS-P8	5 meters
R6LL-P30 to R5SS-P8	5 meters
R6LL-P31 to R3SS-P7	10 meters
R6LL-P32 to R6SS-P8	3 meters

Table 22-58 (Cont.) Leaf Switch Connections for the Sixth Rack in a 10 Rack System

Connection	Cable Length
R6LL-P33 to R4SS-P9	5 meters
R6LL-P34 to R4SS-P29	5 meters
R6LL-P35 to R5SS-P29	5 meters
R6LL-P36 to R2SS-P29	10 meters

Table 22-59 Leaf Switch Connections for the Seventh Rack in a 10 Rack System

Connection	Cable Length
R7UL-P1 to R5SS-P25	5 meters
R7UL-P2 to R9SS-P9	5 meters
R7UL-P3 to R4SS-P25	10 meters
R7UL-P4 to R3SS-P9	10 meters
R7UL-P5 to R1SS-P9	10 meters
R7UL-P6 to R4SS-P8	10 meters
R7UL-P7 to R2SS-P7	10 meters
R7UL-P30 to R7SS-P9	3 meters
R7UL-P31 to R5SS-P9	5 meters
R7UL-P32 to R8SS-P9	5 meters
R7UL-P33 to R6SS-P9	5 meters
R7UL-P34 to R3SS-P25	10 meters
R7UL-P35 to R6SS-P25	5 meters
R7UL-P36 to R10SS-P9	10 meters
R7LL-P1 to R9SS-P12	5 meters
R7LL-P2 to R10SS-P12	10 meters
R7LL-P3 to R5SS-P31	5 meters
R7LL-P4 to R3SS-P15	10 meters
R7LL-P5 to R1SS-P19	10 meters
R7LL-P6 to R4SS-P18	10 meters
R7LL-P7 to R2SS-P19	10 meters
R7LL-P30 to R7SS-P18	3 meters
R7LL-P31 to R5SS-P16	5 meters
R7LL-P32 to R8SS-P12	5 meters
R7LL-P33 to R6SS-P18	5 meters
R7LL-P34 to R6SS-P31	5 meters
R7LL-P35 to R7SS-P31	3 meters
R7LL-P36 to R4SS-P31	10 meters

Table 22-60 Leaf Switch Connections for the Eighth Rack in a 10 Rack System

Connection	Cable Length
R8UL-P1 to R9SS-P19	5 meters
R8UL-P2 to R10SS-P19	5 meters
R8UL-P3 to R5SS-P26	10 meters

Table 22-60 (Cont.) Leaf Switch Connections for the Eighth Rack in a 10 Rack System

Connection	Cable Length
R8UL-P4 to R3SS-P18	10 meters
R8UL-P5 to R1SS-P15	10 meters
R8UL-P6 to R4SS-P16	10 meters
R8UL-P7 to R2SS-P15	10 meters
R8UL-P30 to R7SS-P19	5 meters
R8UL-P31 to R5SS-P19	10 meters
R8UL-P32 to R8SS-P19	3 meters
R8UL-P33 to R6SS-P19	5 meters
R8UL-P34 to R7SS-P26	5 meters
R8UL-P35 to R8SS-P26	3 meters
R8UL-P36 to R6SS-P26	5 meters
R8LL-P1 to R7SS-P29	5 meters
R8LL-P2 to R9SS-P18	5 meters
R8LL-P3 to R6SS-P29	5 meters
R8LL-P4 to R3SS-P12	10 meters
R8LL-P5 to R1SS-P7	10 meters
R8LL-P6 to R4SS-P7	10 meters
R8LL-P7 to R2SS-P9	10 meters
R8LL-P30 to R7SS-P12	5 meters
R8LL-P31 to R5SS-P12	10 meters
R8LL-P32 to R8SS-P18	3 meters
R8LL-P33 to R6SS-P12	5 meters
R8LL-P34 to R8SS-P29	3 meters
R8LL-P35 to R9SS-P29	5 meters
R8LL-P36 to R10SS-P18	5 meters

Table 22-61 Leaf Switch Connections for the Ninth Rack in a 10 Rack System

Connection	Cable Length
R9UL-P1 to R1SS-P24	10 meters
R9UL-P2 to R7SS-P25	5 meters
R9UL-P3 to R2SS-P24	10 meters
R9UL-P4 to R10SS-P24	5 meters
R9UL-P5 to R3SS-P24	10 meters
R9UL-P6 to R10SS-P25	5 meters
R9UL-P7 to R4SS-P24	10 meters
R9UL-P30 to R5SS-P24	10 meters
R9UL-P31 to R6SS-P24	10 meters
R9UL-P32 to R7SS-P24	5 meters
R9UL-P33 to R8SS-P24	5 meters
R9UL-P34 to R8SS-P25	5 meters
R9UL-P35 to R9SS-P24	3 meters
R9UL-P36 to R9SS-P25	3 meters

Table 22-61 (Cont.) Leaf Switch Connections for the Ninth Rack in a 10 Rack System

Connection	Cable Length
R9LL-P1 to R1SS-P29	10 meters
R9LL-P2 to R1SS-P30	10 meters
R9LL-P3 to R2SS-P30	10 meters
R9LL-P4 to R8SS-P31	5 meters
R9LL-P5 to R3SS-P30	10 meters
R9LL-P6 to R10SS-P30	5 meters
R9LL-P7 to R4SS-P30	10 meters
R9LL-P30 to R10SS-P31	5 meters
R9LL-P31 to R5SS-P30	10 meters
R9LL-P32 to R6SS-P30	10 meters
R9LL-P33 to R7SS-P30	5 meters
R9LL-P34 to R8SS-P30	5 meters
R9LL-P35 to R9SS-P30	3 meters
R9LL-P36 to R9SS-P31	3 meters

Table 22-62 Leaf Switch Connections for the Tenth Rack in a 10 Rack System

Connection	Cable Length
R10UL-P1 to R1SS-P21	15 meters
R10UL-P2 to R1SS-P25	15 meters
R10UL-P3 to R2SS-P21	10 meters
R10UL-P4 to R10SS-P21	3 meters
R10UL-P5 to R3SS-P21	10 meters
R10UL-P6 to R10SS-P26	3 meters
R10UL-P7 to R4SS-P21	10 meters
R10UL-P30 to R5SS-P21	10 meters
R10UL-P31 to R6SS-P21	10 meters
R10UL-P32 to R7SS-P21	10 meters
R10UL-P33 to R8SS-P21	5 meters
R10UL-P34 to R2SS-P25	10 meters
R10UL-P35 to R9SS-P21	5 meters
R10UL-P36 to R9SS-P26	5 meters
R10LL-P1 to R1SS-P31	15 meters
R10LL-P2 to R1SS-P32	15 meters
R10LL-P3 to R2SS-P31	10 meters
R10LL-P4 to R2SS-P32	10 meters
R10LL-P5 to R3SS-P31	10 meters
R10LL-P6 to R10SS-P29	3 meters
R10LL-P7 to R4SS-P32	10 meters
R10LL-P30 to R10SS-P32	3 meters
R10LL-P31 to R5SS-P32	10 meters
R10LL-P32 to R6SS-P32	10 meters
R10LL-P33 to R7SS-P32	10 meters

Table 22-62 (Cont.) Leaf Switch Connections for the Tenth Rack in a 10 Rack System

Connection	Cable Length
R10LL-P34 to R8SS-P32	5 meters
R10LL-P35 to R9SS-P32	5 meters
R10LL-P36 to R3SS-P32	10 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-63 Ten-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R8LL-P5	10 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R7UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R6LL-P5	10 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R8UL-P5	10 meters
R1SS-P16	R6UL-P7	10 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R7LL-P5	10 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R10UL-P1	15 meters
R1SS-P22	R5LL-P2	10 meters
R1SS-P23	R1UL-P3	3 meters
R1SS-P24	R9UL-P1	10 meters
R1SS-P25	R10UL-P2	15 meters
R1SS-P26	R6UL-P34	10 meters
R1SS-P27	R4LL-P3	10 meters
R1SS-P28	R5UL-P3	10 meters
R1SS-P29	R9LL-P1	10 meters
R1SS-P30	R9LL-P2	10 meters
R1SS-P31	R10LL-P1	15 meters
R1SS-P32	R10LL-P2	15 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R7UL-P7	10 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R8LL-P7	10 meters

Table 22-63 (Cont.) Ten-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R6LL-P4	10 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R8UL-P7	10 meters
R2SS-P16	R6UL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters
R2SS-P19	R7LL-P7	10 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R10UL-P3	10 meters
R2SS-P22	R5LL-P3	10 meters
R2SS-P23	R1UL-P34	5 meters
R2SS-P24	R9UL-P3	10 meters
R2SS-P25	R10UL-P34	10 meters
R2SS-P26	R6UL-P36	10 meters
R2SS-P27	R1LL-P1	5 meters
R2SS-P28	R5UL-P2	10 meters
R2SS-P29	R6LL-P36	10 meters
R2SS-P30	R9LL-P3	10 meters
R2SS-P31	R10LL-P3	10 meters
R2SS-P32	R10LL-P4	10 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R6LL-P31	10 meters
R3SS-P8	R5UL-P6	5 meters
R3SS-P9	R7UL-P4	10 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R8LL-P4	10 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R7LL-P4	10 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R8UL-P4	10 meters
R3SS-P19	R6UL-P6	10 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R10UL-P5	10 meters
R3SS-P22	R5LL-P34	5 meters
R3SS-P23	R1UL-P1	5 meters

Table 22-63 (Cont.) Ten-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R3SS-P24	R9UL-P5	10 meters
R3SS-P25	R7UL-P34	10 meters
R3SS-P26	R6UL-P2	10 meters
R3SS-P27	R1LL-P2	5 meters
R3SS-P28	R2UL-P2	5 meters
R3SS-P29	R6LL-P3	10 meters
R3SS-P30	R9LL-P5	10 meters
R3SS-P31	R10LL-P5	10 meters
R3SS-P32	R10LL-P36	10 meters
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R8LL-P6	10 meters
R4SS-P8	R7UL-P6	10 meters
R4SS-P9	R6LL-P33	5 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R5UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R6UL-P31	5 meters
R4SS-P16	R8UL-P6	10 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R7LL-P6	10 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R10UL-P7	10 meters
R4SS-P22	R2LL-P2	5 meters
R4SS-P23	R1UL-P2	10 meters
R4SS-P24	R9UL-P7	10 meters
R4SS-P25	R7UL-P3	10 meters
R4SS-P26	R6UL-P1	5 meters
R4SS-P27	R1LL-P34	10 meters
R4SS-P28	R2UL-P3	5 meters
R4SS-P29	R6LL-P34	5 meters
R4SS-P30	R9LL-P7	10 meters
R4SS-P31	R7LL-P36	10 meters
R4SS-P32	R10LL-P7	10 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R6LL-P30	5 meters
R5SS-P9	R7UL-P31	5 meters

Table 22-63 (Cont.) Ten-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R8LL-P31	10 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R7LL-P31	5 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R6UL-P33	5 meters
R5SS-P19	R8UL-P31	10 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R10UL-P30	10 meters
R5SS-P22	R2LL-P35	10 meters
R5SS-P23	R3UL-P34	5 meters
R5SS-P24	R9UL-P30	10 meters
R5SS-P25	R7UL-P1	5 meters
R5SS-P26	R8UL-P3	10 meters
R5SS-P27	R1LL-P3	10 meters
R5SS-P28	R2UL-P1	10 meters
R5SS-P29	R6LL-P35	5 meters
R5SS-P30	R9LL-P31	10 meters
R5SS-P31	R7LL-P3	5 meters
R5SS-P32	R10LL-P31	10 meters
R6SS-P5	R1UL-P32	10 meters
R6SS-P6	R2UL-P33	10 meters
R6SS-P7	R5UL-P7	5 meters
R6SS-P8	R6LL-P32	3 meters
R6SS-P9	R7UL-P33	5 meters
R6SS-P10	R4UL-P30	5 meters
R6SS-P11	R3UL-P7	10 meters
R6SS-P12	R8LL-P33	5 meters
R6SS-P13	R1LL-P32	10 meters
R6SS-P14	R2LL-P7	10 meters
R6SS-P15	R5LL-P4	5 meters
R6SS-P16	R6UL-P30	3 meters
R6SS-P17	R4LL-P6	5 meters
R6SS-P18	R7LL-P33	5 meters
R6SS-P19	R8UL-P33	5 meters
R6SS-P20	R3LL-P30	10 meters
R6SS-P21	R10UL-P31	10 meters
R6SS-P22	R2LL-P1	10 meters
R6SS-P23	R3UL-P1	10 meters

Table 22-63 (Cont.) Ten-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R6SS-P24	R9UL-P31	10 meters
R6SS-P25	R7UL-P35	5 meters
R6SS-P26	R8UL-P36	5 meters
R6SS-P27	R3LL-P35	10 meters
R6SS-P28	R2UL-P35	10 meters
R6SS-P29	R8LL-P3	5 meters
R6SS-P30	R9LL-P32	10 meters
R6SS-P31	R7LL-P34	5 meters
R6SS-P32	R10LL-P32	10 meters
R7SS-P5	R1UL-P33	10 meters
R7SS-P6	R2UL-P32	10 meters
R7SS-P7	R5UL-P32	5 meters
R7SS-P8	R6LL-P6	5 meters
R7SS-P9	R7UL-P30	3 meters
R7SS-P10	R4UL-P31	10 meters
R7SS-P11	R3UL-P31	10 meters
R7SS-P12	R8LL-P30	5 meters
R7SS-P13	R1LL-P6	10 meters
R7SS-P14	R2LL-P32	10 meters
R7SS-P15	R5LL-P5	5 meters
R7SS-P16	R6UL-P5	5 meters
R7SS-P17	R4LL-P32	10 meters
R7SS-P18	R7LL-P30	3 meters
R7SS-P19	R8UL-P30	5 meters
R7SS-P20	R3LL-P7	10 meters
R7SS-P21	R10UL-P32	10 meters
R7SS-P22	R2LL-P34	10 meters
R7SS-P23	R3UL-P35	10 meters
R7SS-P24	R9UL-P32	5 meters
R7SS-P25	R9UL-P2	5 meters
R7SS-P26	R8UL-P34	5 meters
R7SS-P27	R3LL-P1	10 meters
R7SS-P28	R4UL-P36	10 meters
R7SS-P29	R8LL-P1	5 meters
R7SS-P30	R9LL-P33	5 meters
R7SS-P31	R7LL-P35	3 meters
R7SS-P32	R10LL-P33	10 meters
R8SS-P5	R1UL-P7	10 meters
R8SS-P6	R2UL-P30	10 meters
R8SS-P7	R5UL-P31	10 meters
R8SS-P8	R6LL-P7	5 meters
R8SS-P9	R7UL-P32	5 meters

Table 22-63 (Cont.) Ten-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R8SS-P10	R4UL-P4	10 meters
R8SS-P11	R3UL-P33	10 meters
R8SS-P12	R7LL-P32	5 meters
R8SS-P13	R1LL-P33	10 meters
R8SS-P14	R2LL-P30	10 meters
R8SS-P15	R5LL-P32	10 meters
R8SS-P16	R6UL-P32	5 meters
R8SS-P17	R4LL-P33	10 meters
R8SS-P18	R8LL-P32	3 meters
R8SS-P19	R8UL-P32	3 meters
R8SS-P20	R3LL-P6	10 meters
R8SS-P21	R10UL-P33	5 meters
R8SS-P22	R4LL-P36	10 meters
R8SS-P23	R3UL-P36	10 meters
R8SS-P24	R9UL-P33	5 meters
R8SS-P25	R9UL-P34	5 meters
R8SS-P26	R8UL-P35	3 meters
R8SS-P27	R3LL-P36	10 meters
R8SS-P28	R4UL-P1	10 meters
R8SS-P29	R8LL-P34	3 meters
R8SS-P30	R9LL-P34	5 meters
R8SS-P31	R9LL-P4	5 meters
R8SS-P32	R10LL-P34	5 meters
R9SS-P5	R1UL-P36	10 meters
R9SS-P6	R2UL-P36	10 meters
R9SS-P7	R5UL-P35	10 meters
R9SS-P8	R6LL-P1	10 meters
R9SS-P9	R7UL-P2	5 meters
R9SS-P10	R4UL-P34	10 meters
R9SS-P11	R3UL-P2	10 meters
R9SS-P12	R7LL-P1	5 meters
R9SS-P13	R1LL-P36	10 meters
R9SS-P14	R2LL-P3	10 meters
R9SS-P15	R5LL-P1	10 meters
R9SS-P16	R6UL-P3	10 meters
R9SS-P17	R4LL-P1	10 meters
R9SS-P18	R8LL-P2	5 meters
R9SS-P19	R8UL-P1	5 meters
R9SS-P20	R3LL-P3	10 meters
R9SS-P21	R10UL-P35	5 meters
R9SS-P22	R4LL-P35	10 meters
R9SS-P23	R5UL-P36	10 meters

Table 22-63 (Cont.) Ten-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R9SS-P24	R9UL-P35	3 meters
R9SS-P25	R9UL-P36	3 meters
R9SS-P26	R10UL-P36	5 meters
R9SS-P27	R3LL-P34	10 meters
R9SS-P28	R4UL-P35	10 meters
R9SS-P29	R8LL-P35	5 meters
R9SS-P30	R9LL-P35	3 meters
R9SS-P31	R9LL-P36	3 meters
R9SS-P32	R10LL-P35	5 meters
R10SS-P5	R1UL-P35	15 meters
R10SS-P6	R2UL-P34	10 meters
R10SS-P7	R5UL-P1	10 meters
R10SS-P8	R6LL-P2	10 meters
R10SS-P9	R7UL-P36	10 meters
R10SS-P10	R4UL-P3	10 meters
R10SS-P11	R3UL-P3	10 meters
R10SS-P12	R7LL-P2	10 meters
R10SS-P13	R1LL-P35	15 meters
R10SS-P14	R2LL-P36	10 meters
R10SS-P15	R5LL-P35	10 meters
R10SS-P16	R6UL-P35	10 meters
R10SS-P17	R4LL-P2	10 meters
R10SS-P18	R8LL-P36	5 meters
R10SS-P19	R8UL-P2	5 meters
R10SS-P20	R3LL-P2	10 meters
R10SS-P21	R10UL-P4	3 meters
R10SS-P22	R4LL-P34	10 meters
R10SS-P23	R5UL-P34	10 meters
R10SS-P24	R9UL-P4	5 meters
R10SS-P25	R9UL-P6	5 meters
R10SS-P26	R10UL-P6	3 meters
R10SS-P27	R5LL-P36	10 meters
R10SS-P28	R4UL-P2	10 meters
R10SS-P29	R10LL-P6	3 meters
R10SS-P30	R9LL-P6	5 meters
R10SS-P31	R9LL-P30	5 meters
R10SS-P32	R10LL-P30	3 meters

Eleven-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect eleven racks using RoCE Network Fabric.

 Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, interconnecting 10 or more of these racks requires modification to the following cabling tables. Contact Oracle for further details.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a 11 rack system.

Table 22-64 Leaf Switch Connections for the First Rack in a 11 Rack System

Connection	Cable Length
R1UL-P1 to R3SS-P23	5 meters
R1UL-P2 to R11SS-P5	15 meters
R1UL-P3 to R1SS-P23	3 meters
R1UL-P4 to R4SS-P5	10 meters
R1UL-P5 to R1SS-P5	3 meters
R1UL-P6 to R3SS-P5	5 meters
R1UL-P7 to R8SS-P5	10 meters
R1UL-P30 to R5SS-P5	10 meters
R1UL-P31 to R2SS-P5	5 meters
R1UL-P32 to R6SS-P5	10 meters
R1UL-P33 to R7SS-P5	10 meters
R1UL-P34 to R2SS-P23	5 meters

Table 22-64 (Cont.) Leaf Switch Connections for the First Rack in a 11 Rack System

Connection	Cable Length
R1UL-P35 to R10SS-P5	15 meters
R1UL-P36 to R9SS-P5	10 meters
R1LL-P1 to R2SS-P22	5 meters
R1LL-P2 to R3SS-P22	5 meters
R1LL-P3 to R11SS-P13	15 meters
R1LL-P4 to R3SS-P13	5 meters
R1LL-P5 to R1SS-P13	3 meters
R1LL-P6 to R7SS-P13	10 meters
R1LL-P7 to R5SS-P13	10 meters
R1LL-P30 to R4SS-P13	10 meters
R1LL-P31 to R2SS-P13	5 meters
R1LL-P32 to R6SS-P13	10 meters
R1LL-P33 to R8SS-P13	10 meters
R1LL-P34 to R4SS-P22	10 meters
R1LL-P35 to R10SS-P13	15 meters
R1LL-P36 to R9SS-P13	10 meters

Table 22-65 Leaf Switch Connections for the Second Rack in a 11 Rack System

Connection	Cable Length
R2UL-P1 to R5SS-P26	10 meters
R2UL-P2 to R3SS-P26	5 meters
R2UL-P3 to R4SS-P26	5 meters
R2UL-P4 to R4SS-P6	5 meters
R2UL-P5 to R1SS-P6	5 meters
R2UL-P6 to R3SS-P6	5 meters
R2UL-P7 to R5SS-P6	10 meters
R2UL-P30 to R8SS-P6	10 meters
R2UL-P31 to R2SS-P6	3 meters
R2UL-P32 to R7SS-P6	10 meters
R2UL-P33 to R6SS-P6	10 meters
R2UL-P34 to R10SS-P6	10 meters
R2UL-P35 to R11SS-P6	15 meters
R2UL-P36 to R9SS-P6	10 meters
R2LL-P1 to R6SS-P23	10 meters
R2LL-P2 to R4SS-P23	5 meters
R2LL-P3 to R9SS-P14	10 meters
R2LL-P4 to R4SS-P14	5 meters
R2LL-P5 to R1SS-P14	5 meters
R2LL-P6 to R3SS-P14	5 meters
R2LL-P7 to R6SS-P14	10 meters
R2LL-P30 to R8SS-P14	10 meters
R2LL-P31 to R2SS-P14	3 meters

Table 22-65 (Cont.) Leaf Switch Connections for the Second Rack in a 11 Rack System

Connection	Cable Length
R2LL-P32 to R7SS-P14	10 meters
R2LL-P33 to R5SS-P14	10 meters
R2LL-P34 to R11SS-P14	15 meters
R2LL-P35 to R5SS-P23	10 meters
R2LL-P36 to R10SS-P14	10 meters

Table 22-66 Leaf Switch Connections for the Third Rack in a 11 Rack System

Connection	Cable Length
R3UL-P1 to R6SS-P22	10 meters
R3UL-P2 to R9SS-P11	10 meters
R3UL-P3 to R10SS-P11	10 meters
R3UL-P4 to R4SS-P11	5 meters
R3UL-P5 to R1SS-P11	5 meters
R3UL-P6 to R2SS-P11	5 meters
R3UL-P7 to R6SS-P11	10 meters
R3UL-P30 to R3SS-P11	3 meters
R3UL-P31 to R7SS-P11	10 meters
R3UL-P32 to R5SS-P11	5 meters
R3UL-P33 to R8SS-P11	10 meters
R3UL-P34 to R5SS-P22	5 meters
R3UL-P35 to R7SS-P22	10 meters
R3UL-P36 to R11SS-P11	10 meters
R3LL-P1 to R7SS-P26	10 meters
R3LL-P2 to R10SS-P20	10 meters
R3LL-P3 to R9SS-P20	10 meters
R3LL-P4 to R1SS-P20	5 meters
R3LL-P5 to R5SS-P20	5 meters
R3LL-P6 to R8SS-P20	10 meters
R3LL-P7 to R7SS-P20	10 meters
R3LL-P30 to R6SS-P20	10 meters
R3LL-P31 to R2SS-P20	5 meters
R3LL-P32 to R3SS-P20	3 meters
R3LL-P33 to R4SS-P20	5 meters
R3LL-P34 to R11SS-P20	10 meters
R3LL-P35 to R6SS-P26	10 meters
R3LL-P36 to R8SS-P26	10 meters

Table 22-67 Leaf Switch Connections for the Fourth Rack in a 11 Rack System

Connection	Cable Length
R4UL-P1 to R8SS-P23	10 meters
R4UL-P2 to R11SS-P10	10 meters

Table 22-67 (Cont.) Leaf Switch Connections for the Fourth Rack in a 11 Rack System

Connection	Cable Length
R4UL-P3 to R10SS-P10	10 meters
R4UL-P4 to R8SS-P10	10 meters
R4UL-P5 to R5SS-P10	5 meters
R4UL-P6 to R2SS-P10	5 meters
R4UL-P7 to R1SS-P10	10 meters
R4UL-P30 to R6SS-P10	5 meters
R4UL-P31 to R7SS-P10	10 meters
R4UL-P32 to R4SS-P10	3 meters
R4UL-P33 to R3SS-P10	5 meters
R4UL-P34 to R9SS-P10	10 meters
R4UL-P35 to R9SS-P23	10 meters
R4UL-P36 to R7SS-P23	10 meters
R4LL-P1 to R9SS-P17	10 meters
R4LL-P2 to R10SS-P17	10 meters
R4LL-P3 to R11SS-P17	10 meters
R4LL-P4 to R2SS-P17	5 meters
R4LL-P5 to R1SS-P17	10 meters
R4LL-P6 to R6SS-P17	5 meters
R4LL-P7 to R5SS-P17	5 meters
R4LL-P30 to R4SS-P17	3 meters
R4LL-P31 to R3SS-P17	5 meters
R4LL-P32 to R7SS-P17	10 meters
R4LL-P33 to R8SS-P17	10 meters
R4LL-P34 to R10SS-P22	10 meters
R4LL-P35 to R9SS-P22	10 meters
R4LL-P36 to R8SS-P22	10 meters

Table 22-68 Leaf Switch Connections for the Fifth Rack in a 11 Rack System

Connection	Cable Length
R5UL-P1 to R10SS-P7	10 meters
R5UL-P2 to R11SS-P26	10 meters
R5UL-P3 to R11SS-P7	10 meters
R5UL-P4 to R2SS-P8	10 meters
R5UL-P5 to R1SS-P8	10 meters
R5UL-P6 to R3SS-P8	5 meters
R5UL-P7 to R6SS-P7	5 meters
R5UL-P30 to R5SS-P7	3 meters
R5UL-P31 to R8SS-P7	10 meters
R5UL-P32 to R7SS-P7	5 meters
R5UL-P33 to R4SS-P12	5 meters
R5UL-P34 to R10SS-P26	10 meters
R5UL-P35 to R9SS-P7	10 meters

Table 22-68 (Cont.) Leaf Switch Connections for the Fifth Rack in a 11 Rack System

Connection	Cable Length
R5UL-P36 to R9SS-P26	10 meters
R5LL-P1 to R9SS-P15	10 meters
R5LL-P2 to R1SS-P22	10 meters
R5LL-P3 to R11SS-P15	10 meters
R5LL-P4 to R6SS-P15	5 meters
R5LL-P5 to R7SS-P15	5 meters
R5LL-P6 to R2SS-P18	10 meters
R5LL-P7 to R1SS-P18	10 meters
R5LL-P30 to R5SS-P15	3 meters
R5LL-P31 to R3SS-P16	5 meters
R5LL-P32 to R8SS-P15	10 meters
R5LL-P33 to R4SS-P19	5 meters
R5LL-P34 to R11SS-P23	10 meters
R5LL-P35 to R10SS-P15	10 meters
R5LL-P36 to R10SS-P23	10 meters

Table 22-69 Leaf Switch Connections for the Sixth Rack in a 11 Rack System

Connection	Cable Length
R6UL-P1 to R11SS-P22	10 meters
R6UL-P2 to R11SS-P16	10 meters
R6UL-P3 to R9SS-P16	10 meters
R6UL-P4 to R2SS-P16	10 meters
R6UL-P5 to R7SS-P16	5 meters
R6UL-P6 to R3SS-P19	10 meters
R6UL-P7 to R1SS-P16	10 meters
R6UL-P30 to R6SS-P16	3 meters
R6UL-P31 to R4SS-P15	5 meters
R6UL-P32 to R8SS-P16	5 meters
R6UL-P33 to R5SS-P18	5 meters
R6UL-P34 to R1SS-P26	10 meters
R6UL-P35 to R10SS-P16	10 meters
R6UL-P36 to R2SS-P26	10 meters
R6LL-P1 to R9SS-P8	10 meters
R6LL-P2 to R10SS-P8	10 meters
R6LL-P3 to R3SS-P31	10 meters
R6LL-P4 to R2SS-P12	10 meters
R6LL-P5 to R1SS-P12	10 meters
R6LL-P6 to R7SS-P8	5 meters
R6LL-P7 to R8SS-P8	5 meters
R6LL-P30 to R5SS-P8	5 meters
R6LL-P31 to R3SS-P7	10 meters
R6LL-P32 to R6SS-P8	3 meters

Table 22-69 (Cont.) Leaf Switch Connections for the Sixth Rack in a 11 Rack System

Connection	Cable Length
R6LL-P33 to R4SS-P9	5 meters
R6LL-P34 to R1SS-P31	10 meters
R6LL-P35 to R11SS-P8	10 meters
R6LL-P36 to R2SS-P31	10 meters

Table 22-70 Leaf Switch Connections for the Seventh Rack in a 11 Rack System

Connection	Cable Length
R7UL-P1 to R2SS-P27	10 meters
R7UL-P2 to R9SS-P9	5 meters
R7UL-P3 to R4SS-P27	10 meters
R7UL-P4 to R3SS-P9	10 meters
R7UL-P5 to R1SS-P9	10 meters
R7UL-P6 to R4SS-P8	10 meters
R7UL-P7 to R2SS-P7	10 meters
R7UL-P30 to R7SS-P9	3 meters
R7UL-P31 to R5SS-P9	5 meters
R7UL-P32 to R8SS-P9	5 meters
R7UL-P33 to R6SS-P9	5 meters
R7UL-P34 to R3SS-P27	10 meters
R7UL-P35 to R11SS-P9	10 meters
R7UL-P36 to R10SS-P9	10 meters
R7LL-P1 to R9SS-P12	5 meters
R7LL-P2 to R10SS-P12	10 meters
R7LL-P3 to R5SS-P30	5 meters
R7LL-P4 to R3SS-P15	10 meters
R7LL-P5 to R1SS-P19	10 meters
R7LL-P6 to R4SS-P18	10 meters
R7LL-P7 to R2SS-P19	10 meters
R7LL-P30 to R7SS-P18	3 meters
R7LL-P31 to R5SS-P16	5 meters
R7LL-P32 to R8SS-P12	5 meters
R7LL-P33 to R6SS-P18	5 meters
R7LL-P34 to R3SS-P30	10 meters
R7LL-P35 to R11SS-P12	10 meters
R7LL-P36 to R4SS-P30	10 meters

Table 22-71 Leaf Switch Connections for the Eighth Rack in a 11 Rack System

Connection	Cable Length
R8UL-P1 to R9SS-P19	5 meters
R8UL-P2 to R10SS-P19	5 meters
R8UL-P3 to R5SS-P31	10 meters

Table 22-71 (Cont.) Leaf Switch Connections for the Eighth Rack in a 11 Rack System

Connection	Cable Length
R8UL-P4 to R3SS-P18	10 meters
R8UL-P5 to R1SS-P15	10 meters
R8UL-P6 to R4SS-P16	10 meters
R8UL-P7 to R2SS-P15	10 meters
R8UL-P30 to R7SS-P19	5 meters
R8UL-P31 to R5SS-P19	10 meters
R8UL-P32 to R8SS-P19	3 meters
R8UL-P33 to R6SS-P19	5 meters
R8UL-P34 to R4SS-P31	10 meters
R8UL-P35 to R11SS-P19	10 meters
R8UL-P36 to R6SS-P31	5 meters
R8LL-P1 to R7SS-P27	5 meters
R8LL-P2 to R9SS-P18	5 meters
R8LL-P3 to R6SS-P27	5 meters
R8LL-P4 to R3SS-P12	10 meters
R8LL-P5 to R1SS-P7	10 meters
R8LL-P6 to R4SS-P7	10 meters
R8LL-P7 to R2SS-P9	10 meters
R8LL-P30 to R7SS-P12	5 meters
R8LL-P31 to R5SS-P12	10 meters
R8LL-P32 to R8SS-P18	3 meters
R8LL-P33 to R6SS-P12	5 meters
R8LL-P34 to R5SS-P27	10 meters
R8LL-P35 to R11SS-P18	10 meters
R8LL-P36 to R10SS-P18	5 meters

Table 22-72 Leaf Switch Connections for the Ninth Rack in a 11 Rack System

Connection	Cable Length
R9UL-P1 to R1SS-P24	10 meters
R9UL-P2 to R7SS-P30	5 meters
R9UL-P3 to R2SS-P24	10 meters
R9UL-P4 to R10SS-P24	5 meters
R9UL-P5 to R3SS-P24	10 meters
R9UL-P6 to R11SS-P24	5 meters
R9UL-P7 to R4SS-P24	10 meters
R9UL-P30 to R5SS-P24	10 meters
R9UL-P31 to R6SS-P24	10 meters
R9UL-P32 to R7SS-P24	5 meters
R9UL-P33 to R8SS-P24	5 meters
R9UL-P34 to R8SS-P30	5 meters
R9UL-P35 to R9SS-P24	3 meters
R9UL-P36 to R6SS-P30	10 meters

Table 22-72 (Cont.) Leaf Switch Connections for the Ninth Rack in a 11 Rack System

Connection	Cable Length
R9LL-P1 to R1SS-P29	10 meters
R9LL-P2 to R11SS-P29	5 meters
R9LL-P3 to R2SS-P29	10 meters
R9LL-P4 to R8SS-P31	5 meters
R9LL-P5 to R3SS-P29	10 meters
R9LL-P6 to R10SS-P29	5 meters
R9LL-P7 to R4SS-P29	10 meters
R9LL-P30 to R7SS-P31	5 meters
R9LL-P31 to R5SS-P29	10 meters
R9LL-P32 to R6SS-P29	10 meters
R9LL-P33 to R7SS-P29	5 meters
R9LL-P34 to R8SS-P29	5 meters
R9LL-P35 to R9SS-P29	3 meters
R9LL-P36 to R9SS-P31	3 meters

Table 22-73 Leaf Switch Connections for the Tenth Rack in a 11 Rack System

Connection	Cable Length
R10UL-P1 to R1SS-P21	15 meters
R10UL-P2 to R11SS-P21	5 meters
R10UL-P3 to R2SS-P21	10 meters
R10UL-P4 to R10SS-P21	3 meters
R10UL-P5 to R3SS-P21	10 meters
R10UL-P6 to R10SS-P27	3 meters
R10UL-P7 to R4SS-P21	10 meters
R10UL-P30 to R5SS-P21	10 meters
R10UL-P31 to R6SS-P21	10 meters
R10UL-P32 to R7SS-P21	10 meters
R10UL-P33 to R8SS-P21	5 meters
R10UL-P34 to R8SS-P27	5 meters
R10UL-P35 to R9SS-P21	5 meters
R10UL-P36 to R9SS-P27	5 meters
R10LL-P1 to R1SS-P32	15 meters
R10LL-P2 to R9SS-P30	5 meters
R10LL-P3 to R2SS-P32	10 meters
R10LL-P4 to R11SS-P30	5 meters
R10LL-P5 to R3SS-P32	10 meters
R10LL-P6 to R10SS-P30	3 meters
R10LL-P7 to R4SS-P32	10 meters
R10LL-P30 to R10SS-P32	3 meters
R10LL-P31 to R5SS-P32	10 meters
R10LL-P32 to R6SS-P32	10 meters
R10LL-P33 to R7SS-P32	10 meters

Table 22-73 (Cont.) Leaf Switch Connections for the Tenth Rack in a 11 Rack System

Connection	Cable Length
R10LL-P34 to R8SS-P32	5 meters
R10LL-P35 to R9SS-P32	5 meters
R10LL-P36 to R11SS-P32	5 meters

Table 22-74 Leaf Switch Connections for the Eleventh Rack in a 11 Rack System

Connection	Cable Length
R11UL-P1 to R1SS-P25	15 meters
R11UL-P2 to R11SS-P25	3 meters
R11UL-P3 to R2SS-P25	15 meters
R11UL-P4 to R10SS-P25	5 meters
R11UL-P5 to R3SS-P25	10 meters
R11UL-P6 to R10SS-P31	5 meters
R11UL-P7 to R4SS-P25	10 meters
R11UL-P30 to R5SS-P25	10 meters
R11UL-P31 to R6SS-P25	10 meters
R11UL-P32 to R7SS-P25	10 meters
R11UL-P33 to R8SS-P25	10 meters
R11UL-P34 to R11SS-P31	3 meters
R11UL-P35 to R9SS-P25	5 meters
R11UL-P36 to R1SS-P27	15 meters
R11LL-P1 to R1SS-P28	15 meters
R11LL-P2 to R1SS-P30	15 meters
R11LL-P3 to R2SS-P28	15 meters
R11LL-P4 to R11SS-P28	3 meters
R11LL-P5 to R3SS-P28	10 meters
R11LL-P6 to R10SS-P28	5 meters
R11LL-P7 to R4SS-P28	10 meters
R11LL-P30 to R2SS-P30	15 meters
R11LL-P31 to R5SS-P28	10 meters
R11LL-P32 to R6SS-P28	10 meters
R11LL-P33 to R7SS-P28	10 meters
R11LL-P34 to R8SS-P28	10 meters
R11LL-P35 to R9SS-P28	5 meters
R11LL-P36 to R11SS-P27	3 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-75 Eleven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R8LL-P5	10 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R7UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R6LL-P5	10 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R8UL-P5	10 meters
R1SS-P16	R6UL-P7	10 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R7LL-P5	10 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R10UL-P1	15 meters
R1SS-P22	R5LL-P2	10 meters
R1SS-P23	R1UL-P3	3 meters
R1SS-P24	R9UL-P1	10 meters
R1SS-P25	R11UL-P1	15 meters
R1SS-P26	R6UL-P34	10 meters
R1SS-P27	R11UL-P36	15 meters
R1SS-P28	R11LL-P1	15 meters
R1SS-P29	R9LL-P1	10 meters
R1SS-P30	R11LL-P2	15 meters
R1SS-P31	R6LL-P34	10 meters
R1SS-P32	R10LL-P1	15 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R7UL-P7	10 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R8LL-P7	10 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R6LL-P4	10 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R8UL-P7	10 meters
R2SS-P16	R6UL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters

Table 22-75 (Cont.) Eleven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R2SS-P19	R7LL-P7	10 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R10UL-P3	10 meters
R2SS-P22	R1LL-P1	5 meters
R2SS-P23	R1UL-P34	5 meters
R2SS-P24	R9UL-P3	10 meters
R2SS-P25	R11UL-P3	15 meters
R2SS-P26	R6UL-P36	10 meters
R2SS-P27	R7UL-P1	10 meters
R2SS-P28	R11LL-P3	15 meters
R2SS-P29	R9LL-P3	10 meters
R2SS-P30	R11LL-P30	15 meters
R2SS-P31	R6LL-P36	10 meters
R2SS-P32	R10LL-P3	10 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R6LL-P31	10 meters
R3SS-P8	R5UL-P6	5 meters
R3SS-P9	R7UL-P4	10 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R8LL-P4	10 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R7LL-P4	10 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R8UL-P4	10 meters
R3SS-P19	R6UL-P6	10 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R10UL-P5	10 meters
R3SS-P22	R1LL-P2	5 meters
R3SS-P23	R1UL-P1	5 meters
R3SS-P24	R9UL-P5	10 meters
R3SS-P25	R11UL-P5	10 meters
R3SS-P26	R2UL-P2	5 meters
R3SS-P27	R7UL-P34	10 meters
R3SS-P28	R11LL-P5	10 meters
R3SS-P29	R9LL-P5	10 meters
R3SS-P30	R7LL-P34	10 meters
R3SS-P31	R6LL-P3	10 meters
R3SS-P32	R10LL-P5	10 meters

Table 22-75 (Cont.) Eleven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R8LL-P6	10 meters
R4SS-P8	R7UL-P6	10 meters
R4SS-P9	R6LL-P33	5 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R5UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R6UL-P31	5 meters
R4SS-P16	R8UL-P6	10 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R7LL-P6	10 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R10UL-P7	10 meters
R4SS-P22	R1LL-P34	10 meters
R4SS-P23	R2LL-P2	5 meters
R4SS-P24	R9UL-P7	10 meters
R4SS-P25	R11UL-P7	10 meters
R4SS-P26	R2UL-P3	5 meters
R4SS-P27	R7UL-P3	10 meters
R4SS-P28	R11LL-P7	10 meters
R4SS-P29	R9LL-P7	10 meters
R4SS-P30	R7LL-P36	10 meters
R4SS-P31	R8UL-P34	10 meters
R4SS-P32	R10LL-P7	10 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R6LL-P30	5 meters
R5SS-P9	R7UL-P31	5 meters
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R8LL-P31	10 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R7LL-P31	5 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R6UL-P33	5 meters

Table 22-75 (Cont.) Eleven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R5SS-P19	R8UL-P31	10 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R10UL-P30	10 meters
R5SS-P22	R3UL-P34	5 meters
R5SS-P23	R2LL-P35	10 meters
R5SS-P24	R9UL-P30	10 meters
R5SS-P25	R11UL-P30	10 meters
R5SS-P26	R2UL-P1	10 meters
R5SS-P27	R8LL-P34	10 meters
R5SS-P28	R11LL-P31	10 meters
R5SS-P29	R9LL-P31	10 meters
R5SS-P30	R7LL-P3	5 meters
R5SS-P31	R8UL-P3	10 meters
R5SS-P32	R10LL-P31	10 meters
R6SS-P5	R1UL-P32	10 meters
R6SS-P6	R2UL-P33	10 meters
R6SS-P7	R5UL-P7	5 meters
R6SS-P8	R6LL-P32	3 meters
R6SS-P9	R7UL-P33	5 meters
R6SS-P10	R4UL-P30	5 meters
R6SS-P11	R3UL-P7	10 meters
R6SS-P12	R8LL-P33	5 meters
R6SS-P13	R1LL-P32	10 meters
R6SS-P14	R2LL-P7	10 meters
R6SS-P15	R5LL-P4	5 meters
R6SS-P16	R6UL-P30	3 meters
R6SS-P17	R4LL-P6	5 meters
R6SS-P18	R7LL-P33	5 meters
R6SS-P19	R8UL-P33	5 meters
R6SS-P20	R3LL-P30	10 meters
R6SS-P21	R10UL-P31	10 meters
R6SS-P22	R3UL-P1	10 meters
R6SS-P23	R2LL-P1	10 meters
R6SS-P24	R9UL-P31	10 meters
R6SS-P25	R11UL-P31	10 meters
R6SS-P26	R3LL-P35	10 meters
R6SS-P27	R8LL-P3	5 meters
R6SS-P28	R11LL-P32	10 meters
R6SS-P29	R9LL-P32	10 meters
R6SS-P30	R9UL-P36	10 meters
R6SS-P31	R8UL-P36	5 meters
R6SS-P32	R10LL-P32	10 meters

Table 22-75 (Cont.) Eleven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R7SS-P5	R1UL-P33	10 meters
R7SS-P6	R2UL-P32	10 meters
R7SS-P7	R5UL-P32	5 meters
R7SS-P8	R6LL-P6	5 meters
R7SS-P9	R7UL-P30	3 meters
R7SS-P10	R4UL-P31	10 meters
R7SS-P11	R3UL-P31	10 meters
R7SS-P12	R8LL-P30	5 meters
R7SS-P13	R1LL-P6	10 meters
R7SS-P14	R2LL-P32	10 meters
R7SS-P15	R5LL-P5	5 meters
R7SS-P16	R6UL-P5	5 meters
R7SS-P17	R4LL-P32	10 meters
R7SS-P18	R7LL-P30	3 meters
R7SS-P19	R8UL-P30	5 meters
R7SS-P20	R3LL-P7	10 meters
R7SS-P21	R10UL-P32	10 meters
R7SS-P22	R3UL-P35	10 meters
R7SS-P23	R4UL-P36	10 meters
R7SS-P24	R9UL-P32	5 meters
R7SS-P25	R11UL-P32	10 meters
R7SS-P26	R3LL-P1	10 meters
R7SS-P27	R8LL-P1	5 meters
R7SS-P28	R11LL-P33	10 meters
R7SS-P29	R9LL-P33	5 meters
R7SS-P30	R9UL-P2	5 meters
R7SS-P31	R9LL-P30	5 meters
R7SS-P32	R10LL-P33	10 meters
R8SS-P5	R1UL-P7	10 meters
R8SS-P6	R2UL-P30	10 meters
R8SS-P7	R5UL-P31	10 meters
R8SS-P8	R6LL-P7	5 meters
R8SS-P9	R7UL-P32	5 meters
R8SS-P10	R4UL-P4	10 meters
R8SS-P11	R3UL-P33	10 meters
R8SS-P12	R7LL-P32	5 meters
R8SS-P13	R1LL-P33	10 meters
R8SS-P14	R2LL-P30	10 meters
R8SS-P15	R5LL-P32	10 meters
R8SS-P16	R6UL-P32	5 meters
R8SS-P17	R4LL-P33	10 meters
R8SS-P18	R8LL-P32	3 meters

Table 22-75 (Cont.) Eleven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R8SS-P19	R8UL-P32	3 meters
R8SS-P20	R3LL-P6	10 meters
R8SS-P21	R10UL-P33	5 meters
R8SS-P22	R4LL-P36	10 meters
R8SS-P23	R4UL-P1	10 meters
R8SS-P24	R9UL-P33	5 meters
R8SS-P25	R11UL-P33	10 meters
R8SS-P26	R3LL-P36	10 meters
R8SS-P27	R10UL-P34	5 meters
R8SS-P28	R11LL-P34	10 meters
R8SS-P29	R9LL-P34	5 meters
R8SS-P30	R9UL-P34	5 meters
R8SS-P31	R9LL-P4	5 meters
R8SS-P32	R10LL-P34	5 meters
R9SS-P5	R1UL-P36	10 meters
R9SS-P6	R2UL-P36	10 meters
R9SS-P7	R5UL-P35	10 meters
R9SS-P8	R6LL-P1	10 meters
R9SS-P9	R7UL-P2	5 meters
R9SS-P10	R4UL-P34	10 meters
R9SS-P11	R3UL-P2	10 meters
R9SS-P12	R7LL-P1	5 meters
R9SS-P13	R1LL-P36	10 meters
R9SS-P14	R2LL-P3	10 meters
R9SS-P15	R5LL-P1	10 meters
R9SS-P16	R6UL-P3	10 meters
R9SS-P17	R4LL-P1	10 meters
R9SS-P18	R8LL-P2	5 meters
R9SS-P19	R8UL-P1	5 meters
R9SS-P20	R3LL-P3	10 meters
R9SS-P21	R10UL-P35	5 meters
R9SS-P22	R4LL-P35	10 meters
R9SS-P23	R4UL-P35	10 meters
R9SS-P24	R9UL-P35	3 meters
R9SS-P25	R11UL-P35	5 meters
R9SS-P26	R5UL-P36	10 meters
R9SS-P27	R10UL-P36	5 meters
R9SS-P28	R11LL-P35	5 meters
R9SS-P29	R9LL-P35	3 meters
R9SS-P30	R10LL-P2	5 meters
R9SS-P31	R9LL-P36	3 meters
R9SS-P32	R10LL-P35	5 meters

Table 22-75 (Cont.) Eleven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R10SS-P5	R1UL-P35	15 meters
R10SS-P6	R2UL-P34	10 meters
R10SS-P7	R5UL-P1	10 meters
R10SS-P8	R6LL-P2	10 meters
R10SS-P9	R7UL-P36	10 meters
R10SS-P10	R4UL-P3	10 meters
R10SS-P11	R3UL-P3	10 meters
R10SS-P12	R7LL-P2	10 meters
R10SS-P13	R1LL-P35	15 meters
R10SS-P14	R2LL-P36	10 meters
R10SS-P15	R5LL-P35	10 meters
R10SS-P16	R6UL-P35	10 meters
R10SS-P17	R4LL-P2	10 meters
R10SS-P18	R8LL-P36	5 meters
R10SS-P19	R8UL-P2	5 meters
R10SS-P20	R3LL-P2	10 meters
R10SS-P21	R10UL-P4	3 meters
R10SS-P22	R4LL-P34	10 meters
R10SS-P23	R5LL-P36	10 meters
R10SS-P24	R9UL-P4	5 meters
R10SS-P25	R11UL-P4	5 meters
R10SS-P26	R5UL-P34	10 meters
R10SS-P27	R10UL-P6	3 meters
R10SS-P28	R11LL-P6	5 meters
R10SS-P29	R9LL-P6	5 meters
R10SS-P30	R10LL-P6	3 meters
R10SS-P31	R11UL-P6	5 meters
R10SS-P32	R10LL-P30	3 meters
R11SS-P5	R1UL-P2	15 meters
R11SS-P6	R2UL-P35	15 meters
R11SS-P7	R5UL-P3	10 meters
R11SS-P8	R6LL-P35	10 meters
R11SS-P9	R7UL-P35	10 meters
R11SS-P10	R4UL-P2	10 meters
R11SS-P11	R3UL-P36	10 meters
R11SS-P12	R7LL-P35	10 meters
R11SS-P13	R1LL-P3	15 meters
R11SS-P14	R2LL-P34	15 meters
R11SS-P15	R5LL-P3	10 meters
R11SS-P16	R6UL-P2	10 meters
R11SS-P17	R4LL-P3	10 meters
R11SS-P18	R8LL-P35	10 meters

Table 22-75 (Cont.) Eleven-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R11SS-P19	R8UL-P35	10 meters
R11SS-P20	R3LL-P34	10 meters
R11SS-P21	R10UL-P2	5 meters
R11SS-P22	R6UL-P1	10 meters
R11SS-P23	R5LL-P34	10 meters
R11SS-P24	R9UL-P6	5 meters
R11SS-P25	R11UL-P2	3 meters
R11SS-P26	R5UL-P2	10 meters
R11SS-P27	R11LL-P36	3 meters
R11SS-P28	R11LL-P4	3 meters
R11SS-P29	R9LL-P2	5 meters
R11SS-P30	R10LL-P4	5 meters
R11SS-P31	R11UL-P34	3 meters
R11SS-P32	R10LL-P36	5 meters

Twelve-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect twelve racks using RoCE Network Fabric.

 Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, interconnecting 10 or more of these racks requires modification to the following cabling tables. Contact Oracle for further details.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a 12 rack system.

Table 22-76 Leaf Switch Connections for the First Rack in a 12 Rack System

Connection	Cable Length
R1UL-P1 to R12SS-P5	15 meters
R1UL-P2 to R11SS-P5	15 meters
R1UL-P3 to R1SS-P23	3 meters
R1UL-P4 to R4SS-P5	10 meters
R1UL-P5 to R1SS-P5	3 meters
R1UL-P6 to R3SS-P5	5 meters
R1UL-P7 to R8SS-P5	10 meters
R1UL-P30 to R5SS-P5	10 meters
R1UL-P31 to R2SS-P5	5 meters
R1UL-P32 to R6SS-P5	10 meters
R1UL-P33 to R7SS-P5	10 meters
R1UL-P34 to R2SS-P23	5 meters

Table 22-76 (Cont.) Leaf Switch Connections for the First Rack in a 12 Rack System

Connection	Cable Length
R1UL-P35 to R10SS-P5	15 meters
R1UL-P36 to R9SS-P5	10 meters
R1LL-P1 to R2SS-P22	5 meters
R1LL-P2 to R3SS-P22	5 meters
R1LL-P3 to R11SS-P13	15 meters
R1LL-P4 to R3SS-P13	5 meters
R1LL-P5 to R1SS-P13	3 meters
R1LL-P6 to R7SS-P13	10 meters
R1LL-P7 to R5SS-P13	10 meters
R1LL-P30 to R4SS-P13	10 meters
R1LL-P31 to R2SS-P13	5 meters
R1LL-P32 to R6SS-P13	10 meters
R1LL-P33 to R8SS-P13	10 meters
R1LL-P34 to R12SS-P13	15 meters
R1LL-P35 to R10SS-P13	15 meters
R1LL-P36 to R9SS-P13	10 meters

Table 22-77 Leaf Switch Connections for the Second Rack in a 12 Rack System

Connection	Cable Length
R2UL-P1 to R12SS-P6	15 meters
R2UL-P2 to R3SS-P23	5 meters
R2UL-P3 to R4SS-P23	5 meters
R2UL-P4 to R4SS-P6	5 meters
R2UL-P5 to R1SS-P6	5 meters
R2UL-P6 to R3SS-P6	5 meters
R2UL-P7 to R5SS-P6	10 meters
R2UL-P30 to R8SS-P6	10 meters
R2UL-P31 to R2SS-P6	3 meters
R2UL-P32 to R7SS-P6	10 meters
R2UL-P33 to R6SS-P6	10 meters
R2UL-P34 to R10SS-P6	10 meters
R2UL-P35 to R11SS-P6	15 meters
R2UL-P36 to R9SS-P6	10 meters
R2LL-P1 to R12SS-P14	15 meters
R2LL-P2 to R4SS-P22	5 meters
R2LL-P3 to R9SS-P14	10 meters
R2LL-P4 to R4SS-P14	5 meters
R2LL-P5 to R1SS-P14	5 meters
R2LL-P6 to R3SS-P14	5 meters
R2LL-P7 to R6SS-P14	10 meters
R2LL-P30 to R8SS-P14	10 meters
R2LL-P31 to R2SS-P14	3 meters

Table 22-77 (Cont.) Leaf Switch Connections for the Second Rack in a 12 Rack System

Connection	Cable Length
R2LL-P32 to R7SS-P14	10 meters
R2LL-P33 to R5SS-P14	10 meters
R2LL-P34 to R11SS-P14	15 meters
R2LL-P35 to R5SS-P22	10 meters
R2LL-P36 to R10SS-P14	10 meters

Table 22-78 Leaf Switch Connections for the Third Rack in a 12 Rack System

Connection	Cable Length
R3UL-P1 to R6SS-P23	10 meters
R3UL-P2 to R9SS-P11	10 meters
R3UL-P3 to R10SS-P11	10 meters
R3UL-P4 to R4SS-P11	5 meters
R3UL-P5 to R1SS-P11	5 meters
R3UL-P6 to R2SS-P11	5 meters
R3UL-P7 to R6SS-P11	10 meters
R3UL-P30 to R3SS-P11	3 meters
R3UL-P31 to R7SS-P11	10 meters
R3UL-P32 to R5SS-P11	5 meters
R3UL-P33 to R8SS-P11	10 meters
R3UL-P34 to R5SS-P23	5 meters
R3UL-P35 to R12SS-P11	15 meters
R3UL-P36 to R11SS-P11	10 meters
R3LL-P1 to R7SS-P22	10 meters
R3LL-P2 to R10SS-P20	10 meters
R3LL-P3 to R9SS-P20	10 meters
R3LL-P4 to R1SS-P20	5 meters
R3LL-P5 to R5SS-P20	5 meters
R3LL-P6 to R8SS-P20	10 meters
R3LL-P7 to R7SS-P20	10 meters
R3LL-P30 to R6SS-P20	10 meters
R3LL-P31 to R2SS-P20	5 meters
R3LL-P32 to R3SS-P20	3 meters
R3LL-P33 to R4SS-P20	5 meters
R3LL-P34 to R11SS-P20	10 meters
R3LL-P35 to R6SS-P22	10 meters
R3LL-P36 to R12SS-P20	15 meters

Table 22-79 Leaf Switch Connections for the Fourth Rack in a 12 Rack System

Connection	Cable Length
R4UL-P1 to R8SS-P23	10 meters
R4UL-P2 to R11SS-P10	10 meters

Table 22-79 (Cont.) Leaf Switch Connections for the Fourth Rack in a 12 Rack System

Connection	Cable Length
R4UL-P3 to R10SS-P10	10 meters
R4UL-P4 to R8SS-P10	10 meters
R4UL-P5 to R5SS-P10	5 meters
R4UL-P6 to R2SS-P10	5 meters
R4UL-P7 to R1SS-P10	10 meters
R4UL-P30 to R6SS-P10	5 meters
R4UL-P31 to R7SS-P10	10 meters
R4UL-P32 to R4SS-P10	3 meters
R4UL-P33 to R3SS-P10	5 meters
R4UL-P34 to R9SS-P10	10 meters
R4UL-P35 to R12SS-P10	10 meters
R4UL-P36 to R7SS-P23	10 meters
R4LL-P1 to R9SS-P17	10 meters
R4LL-P2 to R10SS-P17	10 meters
R4LL-P3 to R11SS-P17	10 meters
R4LL-P4 to R2SS-P17	5 meters
R4LL-P5 to R1SS-P17	10 meters
R4LL-P6 to R6SS-P17	5 meters
R4LL-P7 to R5SS-P17	5 meters
R4LL-P30 to R4SS-P17	3 meters
R4LL-P31 to R3SS-P17	5 meters
R4LL-P32 to R7SS-P17	10 meters
R4LL-P33 to R8SS-P17	10 meters
R4LL-P34 to R12SS-P17	10 meters
R4LL-P35 to R9SS-P22	10 meters
R4LL-P36 to R8SS-P22	10 meters

Table 22-80 Leaf Switch Connections for the Fifth Rack in a 12 Rack System

Connection	Cable Length
R5UL-P1 to R10SS-P7	10 meters
R5UL-P2 to R12SS-P7	10 meters
R5UL-P3 to R11SS-P7	10 meters
R5UL-P4 to R2SS-P8	10 meters
R5UL-P5 to R1SS-P8	10 meters
R5UL-P6 to R3SS-P8	5 meters
R5UL-P7 to R6SS-P7	5 meters
R5UL-P30 to R5SS-P7	3 meters
R5UL-P31 to R8SS-P7	10 meters
R5UL-P32 to R7SS-P7	5 meters
R5UL-P33 to R4SS-P12	5 meters
R5UL-P34 to R10SS-P23	10 meters
R5UL-P35 to R9SS-P7	10 meters

Table 22-80 (Cont.) Leaf Switch Connections for the Fifth Rack in a 12 Rack System

Connection	Cable Length
R5UL-P36 to R9SS-P23	10 meters
R5LL-P1 to R9SS-P15	10 meters
R5LL-P2 to R12SS-P15	10 meters
R5LL-P3 to R11SS-P15	10 meters
R5LL-P4 to R6SS-P15	5 meters
R5LL-P5 to R7SS-P15	5 meters
R5LL-P6 to R2SS-P18	10 meters
R5LL-P7 to R1SS-P18	10 meters
R5LL-P30 to R5SS-P15	3 meters
R5LL-P31 to R3SS-P16	5 meters
R5LL-P32 to R8SS-P15	10 meters
R5LL-P33 to R4SS-P19	5 meters
R5LL-P34 to R11SS-P22	10 meters
R5LL-P35 to R10SS-P15	10 meters
R5LL-P36 to R10SS-P22	10 meters

Table 22-81 Leaf Switch Connections for the Sixth Rack in a 12 Rack System

Connection	Cable Length
R6UL-P1 to R11SS-P23	10 meters
R6UL-P2 to R11SS-P16	10 meters
R6UL-P3 to R9SS-P16	10 meters
R6UL-P4 to R2SS-P16	10 meters
R6UL-P5 to R7SS-P16	5 meters
R6UL-P6 to R3SS-P19	10 meters
R6UL-P7 to R1SS-P16	10 meters
R6UL-P30 to R6SS-P16	3 meters
R6UL-P31 to R4SS-P15	5 meters
R6UL-P32 to R8SS-P16	5 meters
R6UL-P33 to R5SS-P18	5 meters
R6UL-P34 to R12SS-P16	10 meters
R6UL-P35 to R10SS-P16	10 meters
R6UL-P36 to R12SS-P23	10 meters
R6LL-P1 to R9SS-P8	10 meters
R6LL-P2 to R10SS-P8	10 meters
R6LL-P3 to R1SS-P22	10 meters
R6LL-P4 to R2SS-P12	10 meters
R6LL-P5 to R1SS-P12	10 meters
R6LL-P6 to R7SS-P8	5 meters
R6LL-P7 to R8SS-P8	5 meters
R6LL-P30 to R5SS-P8	5 meters
R6LL-P31 to R3SS-P7	10 meters
R6LL-P32 to R6SS-P8	3 meters

Table 22-81 (Cont.) Leaf Switch Connections for the Sixth Rack in a 12 Rack System

Connection	Cable Length
R6LL-P33 to R4SS-P9	5 meters
R6LL-P34 to R12SS-P22	10 meters
R6LL-P35 to R11SS-P8	10 meters
R6LL-P36 to R12SS-P8	10 meters

Table 22-82 Leaf Switch Connections for the Seventh Rack in a 12 Rack System

Connection	Cable Length
R7UL-P1 to R2SS-P26	10 meters
R7UL-P2 to R9SS-P9	5 meters
R7UL-P3 to R1SS-P26	10 meters
R7UL-P4 to R3SS-P9	10 meters
R7UL-P5 to R1SS-P9	10 meters
R7UL-P6 to R4SS-P8	10 meters
R7UL-P7 to R2SS-P7	10 meters
R7UL-P30 to R7SS-P9	3 meters
R7UL-P31 to R5SS-P9	5 meters
R7UL-P32 to R8SS-P9	5 meters
R7UL-P33 to R6SS-P9	5 meters
R7UL-P34 to R12SS-P9	10 meters
R7UL-P35 to R11SS-P9	10 meters
R7UL-P36 to R10SS-P9	10 meters
R7LL-P1 to R9SS-P12	5 meters
R7LL-P2 to R10SS-P12	10 meters
R7LL-P3 to R3SS-P30	10 meters
R7LL-P4 to R3SS-P15	10 meters
R7LL-P5 to R1SS-P19	10 meters
R7LL-P6 to R4SS-P18	10 meters
R7LL-P7 to R2SS-P19	10 meters
R7LL-P30 to R7SS-P18	3 meters
R7LL-P31 to R5SS-P16	5 meters
R7LL-P32 to R8SS-P12	5 meters
R7LL-P33 to R6SS-P18	5 meters
R7LL-P34 to R12SS-P12	10 meters
R7LL-P35 to R11SS-P12	10 meters
R7LL-P36 to R2SS-P30	10 meters

Table 22-83 Leaf Switch Connections for the Eighth Rack in a 12 Rack System

Connection	Cable Length
R8UL-P1 to R9SS-P19	5 meters
R8UL-P2 to R10SS-P19	5 meters
R8UL-P3 to R3SS-P26	10 meters

Table 22-83 (Cont.) Leaf Switch Connections for the Eighth Rack in a 12 Rack System

Connection	Cable Length
R8UL-P4 to R3SS-P18	10 meters
R8UL-P5 to R1SS-P15	10 meters
R8UL-P6 to R4SS-P16	10 meters
R8UL-P7 to R2SS-P15	10 meters
R8UL-P30 to R7SS-P19	5 meters
R8UL-P31 to R5SS-P19	10 meters
R8UL-P32 to R8SS-P19	3 meters
R8UL-P33 to R6SS-P19	5 meters
R8UL-P34 to R12SS-P19	10 meters
R8UL-P35 to R11SS-P19	10 meters
R8UL-P36 to R4SS-P26	10 meters
R8LL-P1 to R5SS-P30	10 meters
R8LL-P2 to R9SS-P18	5 meters
R8LL-P3 to R4SS-P30	10 meters
R8LL-P4 to R3SS-P12	10 meters
R8LL-P5 to R1SS-P7	10 meters
R8LL-P6 to R4SS-P7	10 meters
R8LL-P7 to R2SS-P9	10 meters
R8LL-P30 to R7SS-P12	5 meters
R8LL-P31 to R5SS-P12	10 meters
R8LL-P32 to R8SS-P18	3 meters
R8LL-P33 to R6SS-P12	5 meters
R8LL-P34 to R12SS-P18	10 meters
R8LL-P35 to R11SS-P18	10 meters
R8LL-P36 to R10SS-P18	5 meters

Table 22-84 Leaf Switch Connections for the Ninth Rack in a 12 Rack System

Connection	Cable Length
R9UL-P1 to R1SS-P24	10 meters
R9UL-P2 to R5SS-P26	10 meters
R9UL-P3 to R2SS-P24	10 meters
R9UL-P4 to R10SS-P24	5 meters
R9UL-P5 to R3SS-P24	10 meters
R9UL-P6 to R11SS-P24	5 meters
R9UL-P7 to R4SS-P24	10 meters
R9UL-P30 to R5SS-P24	10 meters
R9UL-P31 to R6SS-P24	10 meters
R9UL-P32 to R7SS-P24	5 meters
R9UL-P33 to R8SS-P24	5 meters
R9UL-P34 to R6SS-P26	10 meters
R9UL-P35 to R9SS-P24	3 meters
R9UL-P36 to R12SS-P24	10 meters

Table 22-84 (Cont.) Leaf Switch Connections for the Ninth Rack in a 12 Rack System

Connection	Cable Length
R9LL-P1 to R1SS-P29	10 meters
R9LL-P2 to R11SS-P29	5 meters
R9LL-P3 to R2SS-P29	10 meters
R9LL-P4 to R6SS-P30	10 meters
R9LL-P5 to R3SS-P29	10 meters
R9LL-P6 to R10SS-P29	5 meters
R9LL-P7 to R4SS-P29	10 meters
R9LL-P30 to R12SS-P29	10 meters
R9LL-P31 to R5SS-P29	10 meters
R9LL-P32 to R6SS-P29	10 meters
R9LL-P33 to R7SS-P29	5 meters
R9LL-P34 to R8SS-P29	5 meters
R9LL-P35 to R9SS-P29	3 meters
R9LL-P36 to R7SS-P30	5 meters

Table 22-85 Leaf Switch Connections for the Tenth Rack in a 12 Rack System

Connection	Cable Length
R10UL-P1 to R1SS-P21	15 meters
R10UL-P2 to R11SS-P21	5 meters
R10UL-P3 to R2SS-P21	10 meters
R10UL-P4 to R10SS-P21	3 meters
R10UL-P5 to R3SS-P21	10 meters
R10UL-P6 to R8SS-P26	5 meters
R10UL-P7 to R4SS-P21	10 meters
R10UL-P30 to R5SS-P21	10 meters
R10UL-P31 to R6SS-P21	10 meters
R10UL-P32 to R7SS-P21	10 meters
R10UL-P33 to R8SS-P21	5 meters
R10UL-P34 to R12SS-P21	5 meters
R10UL-P35 to R9SS-P21	5 meters
R10UL-P36 to R7SS-P26	10 meters
R10LL-P1 to R1SS-P31	15 meters
R10LL-P2 to R12SS-P31	5 meters
R10LL-P3 to R2SS-P31	10 meters
R10LL-P4 to R11SS-P31	5 meters
R10LL-P5 to R3SS-P31	10 meters
R10LL-P6 to R10SS-P31	3 meters
R10LL-P7 to R4SS-P31	10 meters
R10LL-P30 to R8SS-P31	5 meters
R10LL-P31 to R5SS-P31	10 meters
R10LL-P32 to R6SS-P31	10 meters
R10LL-P33 to R7SS-P31	10 meters

Table 22-85 (Cont.) Leaf Switch Connections for the Tenth Rack in a 12 Rack System

Connection	Cable Length
R10LL-P34 to R8SS-P30	5 meters
R10LL-P35 to R9SS-P30	5 meters
R10LL-P36 to R9SS-P31	5 meters

Table 22-86 Leaf Switch Connections for the Eleventh Rack in a 12 Rack System

Connection	Cable Length
R11UL-P1 to R1SS-P25	15 meters
R11UL-P2 to R11SS-P25	3 meters
R11UL-P3 to R2SS-P25	15 meters
R11UL-P4 to R10SS-P25	5 meters
R11UL-P5 to R3SS-P25	10 meters
R11UL-P6 to R10SS-P26	5 meters
R11UL-P7 to R4SS-P25	10 meters
R11UL-P30 to R5SS-P25	10 meters
R11UL-P31 to R6SS-P25	10 meters
R11UL-P32 to R7SS-P25	10 meters
R11UL-P33 to R8SS-P25	10 meters
R11UL-P34 to R12SS-P25	5 meters
R11UL-P35 to R9SS-P25	5 meters
R11UL-P36 to R9SS-P26	5 meters
R11LL-P1 to R1SS-P28	15 meters
R11LL-P2 to R10SS-P30	5 meters
R11LL-P3 to R2SS-P28	15 meters
R11LL-P4 to R11SS-P28	3 meters
R11LL-P5 to R3SS-P28	10 meters
R11LL-P6 to R10SS-P28	5 meters
R11LL-P7 to R4SS-P28	10 meters
R11LL-P30 to R11SS-P30	3 meters
R11LL-P31 to R5SS-P28	10 meters
R11LL-P32 to R6SS-P28	10 meters
R11LL-P33 to R7SS-P28	10 meters
R11LL-P34 to R8SS-P28	10 meters
R11LL-P35 to R9SS-P28	5 meters
R11LL-P36 to R12SS-P28	5 meters

Table 22-87 Leaf Switch Connections for the Twelfth Rack in a 12 Rack System

Connection	Cable Length
R12UL-P1 to R1SS-P27	15 meters
R12UL-P2 to R11SS-P27	5 meters
R12UL-P3 to R2SS-P27	15 meters
R12UL-P4 to R10SS-P27	5 meters

Table 22-87 (Cont.) Leaf Switch Connections for the Twelfth Rack in a 12 Rack System

Connection	Cable Length
R12UL-P5 to R3SS-P27	15 meters
R12UL-P6 to R11SS-P26	5 meters
R12UL-P7 to R4SS-P27	10 meters
R12UL-P30 to R5SS-P27	10 meters
R12UL-P31 to R6SS-P27	10 meters
R12UL-P32 to R7SS-P27	10 meters
R12UL-P33 to R8SS-P27	10 meters
R12UL-P34 to R12SS-P27	3 meters
R12UL-P35 to R9SS-P27	10 meters
R12UL-P36 to R12SS-P26	3 meters
R12LL-P1 to R1SS-P32	15 meters
R12LL-P2 to R1SS-P30	15 meters
R12LL-P3 to R2SS-P32	15 meters
R12LL-P4 to R11SS-P32	5 meters
R12LL-P5 to R3SS-P32	15 meters
R12LL-P6 to R10SS-P32	5 meters
R12LL-P7 to R4SS-P32	10 meters
R12LL-P30 to R12SS-P30	3 meters
R12LL-P31 to R5SS-P32	10 meters
R12LL-P32 to R6SS-P32	10 meters
R12LL-P33 to R7SS-P32	10 meters
R12LL-P34 to R8SS-P32	10 meters
R12LL-P35 to R9SS-P32	10 meters
R12LL-P36 to R12SS-P32	3 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-88 Twelve-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R8LL-P5	10 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R7UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R6LL-P5	10 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R8UL-P5	10 meters

Table 22-88 (Cont.) Twelve-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P16	R6UL-P7	10 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R7LL-P5	10 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R10UL-P1	15 meters
R1SS-P22	R6LL-P3	10 meters
R1SS-P23	R1UL-P3	3 meters
R1SS-P24	R9UL-P1	10 meters
R1SS-P25	R11UL-P1	15 meters
R1SS-P26	R7UL-P3	10 meters
R1SS-P27	R12UL-P1	15 meters
R1SS-P28	R11LL-P1	15 meters
R1SS-P29	R9LL-P1	10 meters
R1SS-P30	R12LL-P2	15 meters
R1SS-P31	R10LL-P1	15 meters
R1SS-P32	R12LL-P1	15 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R7UL-P7	10 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R8LL-P7	10 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R6LL-P4	10 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R8UL-P7	10 meters
R2SS-P16	R6UL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters
R2SS-P19	R7LL-P7	10 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R10UL-P3	10 meters
R2SS-P22	R1LL-P1	5 meters
R2SS-P23	R1UL-P34	5 meters
R2SS-P24	R9UL-P3	10 meters
R2SS-P25	R11UL-P3	15 meters
R2SS-P26	R7UL-P1	10 meters
R2SS-P27	R12UL-P3	15 meters
R2SS-P28	R11LL-P3	15 meters
R2SS-P29	R9LL-P3	10 meters

Table 22-88 (Cont.) Twelve-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R2SS-P30	R7LL-P36	10 meters
R2SS-P31	R10LL-P3	10 meters
R2SS-P32	R12LL-P3	15 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R6LL-P31	10 meters
R3SS-P8	R5UL-P6	5 meters
R3SS-P9	R7UL-P4	10 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R8LL-P4	10 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R7LL-P4	10 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R8UL-P4	10 meters
R3SS-P19	R6UL-P6	10 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R10UL-P5	10 meters
R3SS-P22	R1LL-P2	5 meters
R3SS-P23	R2UL-P2	5 meters
R3SS-P24	R9UL-P5	10 meters
R3SS-P25	R11UL-P5	10 meters
R3SS-P26	R8UL-P3	10 meters
R3SS-P27	R12UL-P5	15 meters
R3SS-P28	R11LL-P5	10 meters
R3SS-P29	R9LL-P5	10 meters
R3SS-P30	R7LL-P3	10 meters
R3SS-P31	R10LL-P5	10 meters
R3SS-P32	R12LL-P5	15 meters
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R8LL-P6	10 meters
R4SS-P8	R7UL-P6	10 meters
R4SS-P9	R6LL-P33	5 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R5UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R6UL-P31	5 meters

Table 22-88 (Cont.) Twelve-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R4SS-P16	R8UL-P6	10 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R7LL-P6	10 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R10UL-P7	10 meters
R4SS-P22	R2LL-P2	5 meters
R4SS-P23	R2UL-P3	5 meters
R4SS-P24	R9UL-P7	10 meters
R4SS-P25	R11UL-P7	10 meters
R4SS-P26	R8UL-P36	10 meters
R4SS-P27	R12UL-P7	10 meters
R4SS-P28	R11LL-P7	10 meters
R4SS-P29	R9LL-P7	10 meters
R4SS-P30	R8LL-P3	10 meters
R4SS-P31	R10LL-P7	10 meters
R4SS-P32	R12LL-P7	10 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R6LL-P30	5 meters
R5SS-P9	R7UL-P31	5 meters
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R8LL-P31	10 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R7LL-P31	5 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R6UL-P33	5 meters
R5SS-P19	R8UL-P31	10 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R10UL-P30	10 meters
R5SS-P22	R2LL-P35	10 meters
R5SS-P23	R3UL-P34	5 meters
R5SS-P24	R9UL-P30	10 meters
R5SS-P25	R11UL-P30	10 meters
R5SS-P26	R9UL-P2	10 meters
R5SS-P27	R12UL-P30	10 meters
R5SS-P28	R11LL-P31	10 meters
R5SS-P29	R9LL-P31	10 meters

Table 22-88 (Cont.) Twelve-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R5SS-P30	R8LL-P1	10 meters
R5SS-P31	R10LL-P31	10 meters
R5SS-P32	R12LL-P31	10 meters
R6SS-P5	R1UL-P32	10 meters
R6SS-P6	R2UL-P33	10 meters
R6SS-P7	R5UL-P7	5 meters
R6SS-P8	R6LL-P32	3 meters
R6SS-P9	R7UL-P33	5 meters
R6SS-P10	R4UL-P30	5 meters
R6SS-P11	R3UL-P7	10 meters
R6SS-P12	R8LL-P33	5 meters
R6SS-P13	R1LL-P32	10 meters
R6SS-P14	R2LL-P7	10 meters
R6SS-P15	R5LL-P4	5 meters
R6SS-P16	R6UL-P30	3 meters
R6SS-P17	R4LL-P6	5 meters
R6SS-P18	R7LL-P33	5 meters
R6SS-P19	R8UL-P33	5 meters
R6SS-P20	R3LL-P30	10 meters
R6SS-P21	R10UL-P31	10 meters
R6SS-P22	R3LL-P35	10 meters
R6SS-P23	R3UL-P1	10 meters
R6SS-P24	R9UL-P31	10 meters
R6SS-P25	R11UL-P31	10 meters
R6SS-P26	R9UL-P34	10 meters
R6SS-P27	R12UL-P31	10 meters
R6SS-P28	R11LL-P32	10 meters
R6SS-P29	R9LL-P32	10 meters
R6SS-P30	R9LL-P4	10 meters
R6SS-P31	R10LL-P32	10 meters
R6SS-P32	R12LL-P32	10 meters
R7SS-P5	R1UL-P33	10 meters
R7SS-P6	R2UL-P32	10 meters
R7SS-P7	R5UL-P32	5 meters
R7SS-P8	R6LL-P6	5 meters
R7SS-P9	R7UL-P30	3 meters
R7SS-P10	R4UL-P31	10 meters
R7SS-P11	R3UL-P31	10 meters
R7SS-P12	R8LL-P30	5 meters
R7SS-P13	R1LL-P6	10 meters
R7SS-P14	R2LL-P32	10 meters
R7SS-P15	R5LL-P5	5 meters

Table 22-88 (Cont.) Twelve-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R7SS-P16	R6UL-P5	5 meters
R7SS-P17	R4LL-P32	10 meters
R7SS-P18	R7LL-P30	3 meters
R7SS-P19	R8UL-P30	5 meters
R7SS-P20	R3LL-P7	10 meters
R7SS-P21	R10UL-P32	10 meters
R7SS-P22	R3LL-P1	10 meters
R7SS-P23	R4UL-P36	10 meters
R7SS-P24	R9UL-P32	5 meters
R7SS-P25	R11UL-P32	10 meters
R7SS-P26	R10UL-P36	10 meters
R7SS-P27	R12UL-P32	10 meters
R7SS-P28	R11LL-P33	10 meters
R7SS-P29	R9LL-P33	5 meters
R7SS-P30	R9LL-P36	5 meters
R7SS-P31	R10LL-P33	10 meters
R7SS-P32	R12LL-P33	10 meters
R8SS-P5	R1UL-P7	10 meters
R8SS-P6	R2UL-P30	10 meters
R8SS-P7	R5UL-P31	10 meters
R8SS-P8	R6LL-P7	5 meters
R8SS-P9	R7UL-P32	5 meters
R8SS-P10	R4UL-P4	10 meters
R8SS-P11	R3UL-P33	10 meters
R8SS-P12	R7LL-P32	5 meters
R8SS-P13	R1LL-P33	10 meters
R8SS-P14	R2LL-P30	10 meters
R8SS-P15	R5LL-P32	10 meters
R8SS-P16	R6UL-P32	5 meters
R8SS-P17	R4LL-P33	10 meters
R8SS-P18	R8LL-P32	3 meters
R8SS-P19	R8UL-P32	3 meters
R8SS-P20	R3LL-P6	10 meters
R8SS-P21	R10UL-P33	5 meters
R8SS-P22	R4LL-P36	10 meters
R8SS-P23	R4UL-P1	10 meters
R8SS-P24	R9UL-P33	5 meters
R8SS-P25	R11UL-P33	10 meters
R8SS-P26	R10UL-P6	5 meters
R8SS-P27	R12UL-P33	10 meters
R8SS-P28	R11LL-P34	10 meters
R8SS-P29	R9LL-P34	5 meters

Table 22-88 (Cont.) Twelve-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R8SS-P30	R10LL-P34	5 meters
R8SS-P31	R10LL-P30	5 meters
R8SS-P32	R12LL-P34	10 meters
R9SS-P5	R1UL-P36	10 meters
R9SS-P6	R2UL-P36	10 meters
R9SS-P7	R5UL-P35	10 meters
R9SS-P8	R6LL-P1	10 meters
R9SS-P9	R7UL-P2	5 meters
R9SS-P10	R4UL-P34	10 meters
R9SS-P11	R3UL-P2	10 meters
R9SS-P12	R7LL-P1	5 meters
R9SS-P13	R1LL-P36	10 meters
R9SS-P14	R2LL-P3	10 meters
R9SS-P15	R5LL-P1	10 meters
R9SS-P16	R6UL-P3	10 meters
R9SS-P17	R4LL-P1	10 meters
R9SS-P18	R8LL-P2	5 meters
R9SS-P19	R8UL-P1	5 meters
R9SS-P20	R3LL-P3	10 meters
R9SS-P21	R10UL-P35	5 meters
R9SS-P22	R4LL-P35	10 meters
R9SS-P23	R5UL-P36	10 meters
R9SS-P24	R9UL-P35	3 meters
R9SS-P25	R11UL-P35	5 meters
R9SS-P26	R11UL-P36	5 meters
R9SS-P27	R12UL-P35	10 meters
R9SS-P28	R11LL-P35	5 meters
R9SS-P29	R9LL-P35	3 meters
R9SS-P30	R10LL-P35	5 meters
R9SS-P31	R10LL-P36	5 meters
R9SS-P32	R12LL-P35	10 meters
R10SS-P5	R1UL-P35	15 meters
R10SS-P6	R2UL-P34	10 meters
R10SS-P7	R5UL-P1	10 meters
R10SS-P8	R6LL-P2	10 meters
R10SS-P9	R7UL-P36	10 meters
R10SS-P10	R4UL-P3	10 meters
R10SS-P11	R3UL-P3	10 meters
R10SS-P12	R7LL-P2	10 meters
R10SS-P13	R1LL-P35	15 meters
R10SS-P14	R2LL-P36	10 meters
R10SS-P15	R5LL-P35	10 meters

Table 22-88 (Cont.) Twelve-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R10SS-P16	R6UL-P35	10 meters
R10SS-P17	R4LL-P2	10 meters
R10SS-P18	R8LL-P36	5 meters
R10SS-P19	R8UL-P2	5 meters
R10SS-P20	R3LL-P2	10 meters
R10SS-P21	R10UL-P4	3 meters
R10SS-P22	R5LL-P36	10 meters
R10SS-P23	R5UL-P34	10 meters
R10SS-P24	R9UL-P4	5 meters
R10SS-P25	R11UL-P4	5 meters
R10SS-P26	R11UL-P6	5 meters
R10SS-P27	R12UL-P4	5 meters
R10SS-P28	R11LL-P6	5 meters
R10SS-P29	R9LL-P6	5 meters
R10SS-P30	R11LL-P2	5 meters
R10SS-P31	R10LL-P6	3 meters
R10SS-P32	R12LL-P6	5 meters
R11SS-P5	R1UL-P2	15 meters
R11SS-P6	R2UL-P35	15 meters
R11SS-P7	R5UL-P3	10 meters
R11SS-P8	R6LL-P35	10 meters
R11SS-P9	R7UL-P35	10 meters
R11SS-P10	R4UL-P2	10 meters
R11SS-P11	R3UL-P36	10 meters
R11SS-P12	R7LL-P35	10 meters
R11SS-P13	R1LL-P3	15 meters
R11SS-P14	R2LL-P34	15 meters
R11SS-P15	R5LL-P3	10 meters
R11SS-P16	R6UL-P2	10 meters
R11SS-P17	R4LL-P3	10 meters
R11SS-P18	R8LL-P35	10 meters
R11SS-P19	R8UL-P35	10 meters
R11SS-P20	R3LL-P34	10 meters
R11SS-P21	R10UL-P2	5 meters
R11SS-P22	R5LL-P34	10 meters
R11SS-P23	R6UL-P1	10 meters
R11SS-P24	R9UL-P6	5 meters
R11SS-P25	R11UL-P2	3 meters
R11SS-P26	R12UL-P6	5 meters
R11SS-P27	R12UL-P2	5 meters
R11SS-P28	R11LL-P4	3 meters
R11SS-P29	R9LL-P2	5 meters

Table 22-88 (Cont.) Twelve-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R11SS-P30	R11LL-P30	3 meters
R11SS-P31	R10LL-P4	5 meters
R11SS-P32	R12LL-P4	5 meters
R12SS-P5	R1UL-P1	15 meters
R12SS-P6	R2UL-P1	15 meters
R12SS-P7	R5UL-P2	10 meters
R12SS-P8	R6LL-P36	10 meters
R12SS-P9	R7UL-P34	10 meters
R12SS-P10	R4UL-P35	10 meters
R12SS-P11	R3UL-P35	15 meters
R12SS-P12	R7LL-P34	10 meters
R12SS-P13	R1LL-P34	15 meters
R12SS-P14	R2LL-P1	15 meters
R12SS-P15	R5LL-P2	10 meters
R12SS-P16	R6UL-P34	10 meters
R12SS-P17	R4LL-P34	10 meters
R12SS-P18	R8LL-P34	10 meters
R12SS-P19	R8UL-P34	10 meters
R12SS-P20	R3LL-P36	15 meters
R12SS-P21	R10UL-P34	5 meters
R12SS-P22	R6LL-P34	10 meters
R12SS-P23	R6UL-P36	10 meters
R12SS-P24	R9UL-P36	10 meters
R12SS-P25	R11UL-P34	5 meters
R12SS-P26	R12UL-P36	3 meters
R12SS-P27	R12UL-P34	3 meters
R12SS-P28	R11LL-P36	5 meters
R12SS-P29	R9LL-P30	10 meters
R12SS-P30	R12LL-P30	3 meters
R12SS-P31	R10LL-P2	5 meters
R12SS-P32	R12LL-P36	3 meters

Thirteen-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect thirteen racks using RoCE Network Fabric.



Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, interconnecting 10 or more of these racks requires modification to the following cabling tables. Contact Oracle for further details.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a 13 rack system.

Table 22-89 Leaf Switch Connections for the First Rack in a 13 Rack System

Connection	Cable Length
R1UL-P1 to R12SS-P5	15 meters
R1UL-P2 to R11SS-P5	15 meters
R1UL-P3 to R1SS-P23	3 meters
R1UL-P4 to R4SS-P5	10 meters
R1UL-P5 to R1SS-P5	3 meters
R1UL-P6 to R3SS-P5	5 meters
R1UL-P7 to R8SS-P5	10 meters
R1UL-P30 to R5SS-P5	10 meters
R1UL-P31 to R2SS-P5	5 meters
R1UL-P32 to R6SS-P5	10 meters
R1UL-P33 to R7SS-P5	10 meters
R1UL-P34 to R13SS-P5	15 meters

Table 22-89 (Cont.) Leaf Switch Connections for the First Rack in a 13 Rack System

Connection	Cable Length
R1UL-P35 to R10SS-P5	15 meters
R1UL-P36 to R9SS-P5	10 meters
R1LL-P1 to R2SS-P23	5 meters
R1LL-P2 to R13SS-P13	15 meters
R1LL-P3 to R11SS-P13	15 meters
R1LL-P4 to R3SS-P13	5 meters
R1LL-P5 to R1SS-P13	3 meters
R1LL-P6 to R7SS-P13	10 meters
R1LL-P7 to R5SS-P13	10 meters
R1LL-P30 to R4SS-P13	10 meters
R1LL-P31 to R2SS-P13	5 meters
R1LL-P32 to R6SS-P13	10 meters
R1LL-P33 to R8SS-P13	10 meters
R1LL-P34 to R12SS-P13	15 meters
R1LL-P35 to R10SS-P13	15 meters
R1LL-P36 to R9SS-P13	10 meters

Table 22-90 Leaf Switch Connections for the Second Rack in a 13 Rack System

Connection	Cable Length
R2UL-P1 to R12SS-P6	15 meters
R2UL-P2 to R3SS-P23	5 meters
R2UL-P3 to R13SS-P6	15 meters
R2UL-P4 to R4SS-P6	5 meters
R2UL-P5 to R1SS-P6	5 meters
R2UL-P6 to R3SS-P6	5 meters
R2UL-P7 to R5SS-P6	10 meters
R2UL-P30 to R8SS-P6	10 meters
R2UL-P31 to R2SS-P6	3 meters
R2UL-P32 to R7SS-P6	10 meters
R2UL-P33 to R6SS-P6	10 meters
R2UL-P34 to R10SS-P6	10 meters
R2UL-P35 to R11SS-P6	15 meters
R2UL-P36 to R9SS-P6	10 meters
R2LL-P1 to R12SS-P14	15 meters
R2LL-P2 to R4SS-P23	5 meters
R2LL-P3 to R9SS-P14	10 meters
R2LL-P4 to R4SS-P14	5 meters
R2LL-P5 to R1SS-P14	5 meters
R2LL-P6 to R3SS-P14	5 meters
R2LL-P7 to R6SS-P14	10 meters
R2LL-P30 to R8SS-P14	10 meters
R2LL-P31 to R2SS-P14	3 meters

Table 22-90 (Cont.) Leaf Switch Connections for the Second Rack in a 13 Rack System

Connection	Cable Length
R2LL-P32 to R7SS-P14	10 meters
R2LL-P33 to R5SS-P14	10 meters
R2LL-P34 to R11SS-P14	15 meters
R2LL-P35 to R13SS-P14	15 meters
R2LL-P36 to R10SS-P14	10 meters

Table 22-91 Leaf Switch Connections for the Third Rack in a 13 Rack System

Connection	Cable Length
R3UL-P1 to R13SS-P11	15 meters
R3UL-P2 to R9SS-P11	10 meters
R3UL-P3 to R10SS-P11	10 meters
R3UL-P4 to R4SS-P11	5 meters
R3UL-P5 to R1SS-P11	5 meters
R3UL-P6 to R2SS-P11	5 meters
R3UL-P7 to R6SS-P11	10 meters
R3UL-P30 to R3SS-P11	3 meters
R3UL-P31 to R7SS-P11	10 meters
R3UL-P32 to R5SS-P11	5 meters
R3UL-P33 to R8SS-P11	10 meters
R3UL-P34 to R5SS-P23	5 meters
R3UL-P35 to R12SS-P11	15 meters
R3UL-P36 to R11SS-P11	10 meters
R3LL-P1 to R13SS-P20	15 meters
R3LL-P2 to R10SS-P20	10 meters
R3LL-P3 to R9SS-P20	10 meters
R3LL-P4 to R1SS-P20	5 meters
R3LL-P5 to R5SS-P20	5 meters
R3LL-P6 to R8SS-P20	10 meters
R3LL-P7 to R7SS-P20	10 meters
R3LL-P30 to R6SS-P20	10 meters
R3LL-P31 to R2SS-P20	5 meters
R3LL-P32 to R3SS-P20	3 meters
R3LL-P33 to R4SS-P20	5 meters
R3LL-P34 to R11SS-P20	10 meters
R3LL-P35 to R6SS-P23	10 meters
R3LL-P36 to R12SS-P20	15 meters

Table 22-92 Leaf Switch Connections for the Fourth Rack in a 13 Rack System

Connection	Cable Length
R4UL-P1 to R13SS-P10	15 meters
R4UL-P2 to R11SS-P10	10 meters

Table 22-92 (Cont.) Leaf Switch Connections for the Fourth Rack in a 13 Rack System

Connection	Cable Length
R4UL-P3 to R10SS-P10	10 meters
R4UL-P4 to R8SS-P10	10 meters
R4UL-P5 to R5SS-P10	5 meters
R4UL-P6 to R2SS-P10	5 meters
R4UL-P7 to R1SS-P10	10 meters
R4UL-P30 to R6SS-P10	5 meters
R4UL-P31 to R7SS-P10	10 meters
R4UL-P32 to R4SS-P10	3 meters
R4UL-P33 to R3SS-P10	5 meters
R4UL-P34 to R9SS-P10	10 meters
R4UL-P35 to R12SS-P10	10 meters
R4UL-P36 to R7SS-P23	10 meters
R4LL-P1 to R9SS-P17	10 meters
R4LL-P2 to R10SS-P17	10 meters
R4LL-P3 to R11SS-P17	10 meters
R4LL-P4 to R2SS-P17	5 meters
R4LL-P5 to R1SS-P17	10 meters
R4LL-P6 to R6SS-P17	5 meters
R4LL-P7 to R5SS-P17	5 meters
R4LL-P30 to R4SS-P17	3 meters
R4LL-P31 to R3SS-P17	5 meters
R4LL-P32 to R7SS-P17	10 meters
R4LL-P33 to R8SS-P17	10 meters
R4LL-P34 to R12SS-P17	10 meters
R4LL-P35 to R13SS-P17	15 meters
R4LL-P36 to R8SS-P23	10 meters

Table 22-93 Leaf Switch Connections for the Fifth Rack in a 13 Rack System

Connection	Cable Length
R5UL-P1 to R10SS-P7	10 meters
R5UL-P2 to R12SS-P7	10 meters
R5UL-P3 to R11SS-P7	10 meters
R5UL-P4 to R2SS-P8	10 meters
R5UL-P5 to R1SS-P8	10 meters
R5UL-P6 to R3SS-P8	5 meters
R5UL-P7 to R6SS-P7	5 meters
R5UL-P30 to R5SS-P7	3 meters
R5UL-P31 to R8SS-P7	10 meters
R5UL-P32 to R7SS-P7	5 meters
R5UL-P33 to R4SS-P12	5 meters
R5UL-P34 to R13SS-P7	10 meters
R5UL-P35 to R9SS-P7	10 meters

Table 22-93 (Cont.) Leaf Switch Connections for the Fifth Rack in a 13 Rack System

Connection	Cable Length
R5UL-P36 to R9SS-P23	10 meters
R5LL-P1 to R9SS-P15	10 meters
R5LL-P2 to R12SS-P15	10 meters
R5LL-P3 to R11SS-P15	10 meters
R5LL-P4 to R6SS-P15	5 meters
R5LL-P5 to R7SS-P15	5 meters
R5LL-P6 to R2SS-P18	10 meters
R5LL-P7 to R1SS-P18	10 meters
R5LL-P30 to R5SS-P15	3 meters
R5LL-P31 to R3SS-P16	5 meters
R5LL-P32 to R8SS-P15	10 meters
R5LL-P33 to R4SS-P19	5 meters
R5LL-P34 to R13SS-P15	10 meters
R5LL-P35 to R10SS-P15	10 meters
R5LL-P36 to R10SS-P23	10 meters

Table 22-94 Leaf Switch Connections for the Sixth Rack in a 13 Rack System

Connection	Cable Length
R6UL-P1 to R11SS-P23	10 meters
R6UL-P2 to R11SS-P16	10 meters
R6UL-P3 to R9SS-P16	10 meters
R6UL-P4 to R2SS-P16	10 meters
R6UL-P5 to R7SS-P16	5 meters
R6UL-P6 to R3SS-P19	10 meters
R6UL-P7 to R1SS-P16	10 meters
R6UL-P30 to R6SS-P16	3 meters
R6UL-P31 to R4SS-P15	5 meters
R6UL-P32 to R8SS-P16	5 meters
R6UL-P33 to R5SS-P18	5 meters
R6UL-P34 to R12SS-P16	10 meters
R6UL-P35 to R10SS-P16	10 meters
R6UL-P36 to R13SS-P16	10 meters
R6LL-P1 to R9SS-P8	10 meters
R6LL-P2 to R10SS-P8	10 meters
R6LL-P3 to R13SS-P8	10 meters
R6LL-P4 to R2SS-P12	10 meters
R6LL-P5 to R1SS-P12	10 meters
R6LL-P6 to R7SS-P8	5 meters
R6LL-P7 to R8SS-P8	5 meters
R6LL-P30 to R5SS-P8	5 meters
R6LL-P31 to R3SS-P7	10 meters
R6LL-P32 to R6SS-P8	3 meters

Table 22-94 (Cont.) Leaf Switch Connections for the Sixth Rack in a 13 Rack System

Connection	Cable Length
R6LL-P33 to R4SS-P9	5 meters
R6LL-P34 to R12SS-P23	10 meters
R6LL-P35 to R11SS-P8	10 meters
R6LL-P36 to R12SS-P8	10 meters

Table 22-95 Leaf Switch Connections for the Seventh Rack in a 13 Rack System

Connection	Cable Length
R7UL-P1 to R13SS-P23	10 meters
R7UL-P2 to R9SS-P9	5 meters
R7UL-P3 to R13SS-P9	10 meters
R7UL-P4 to R3SS-P9	10 meters
R7UL-P5 to R1SS-P9	10 meters
R7UL-P6 to R4SS-P8	10 meters
R7UL-P7 to R2SS-P7	10 meters
R7UL-P30 to R7SS-P9	3 meters
R7UL-P31 to R5SS-P9	5 meters
R7UL-P32 to R8SS-P9	5 meters
R7UL-P33 to R6SS-P9	5 meters
R7UL-P34 to R12SS-P9	10 meters
R7UL-P35 to R11SS-P9	10 meters
R7UL-P36 to R10SS-P9	10 meters
R7LL-P1 to R9SS-P12	5 meters
R7LL-P2 to R10SS-P12	10 meters
R7LL-P3 to R1SS-P26	10 meters
R7LL-P4 to R3SS-P15	10 meters
R7LL-P5 to R1SS-P19	10 meters
R7LL-P6 to R4SS-P18	10 meters
R7LL-P7 to R2SS-P19	10 meters
R7LL-P30 to R7SS-P18	3 meters
R7LL-P31 to R5SS-P16	5 meters
R7LL-P32 to R8SS-P12	5 meters
R7LL-P33 to R6SS-P18	5 meters
R7LL-P34 to R12SS-P12	10 meters
R7LL-P35 to R11SS-P12	10 meters
R7LL-P36 to R13SS-P12	10 meters

Table 22-96 Leaf Switch Connections for the Eighth Rack in a 13 Rack System

Connection	Cable Length
R8UL-P1 to R9SS-P19	5 meters
R8UL-P2 to R10SS-P19	5 meters
R8UL-P3 to R2SS-P26	10 meters

Table 22-96 (Cont.) Leaf Switch Connections for the Eighth Rack in a 13 Rack System

Connection	Cable Length
R8UL-P4 to R3SS-P18	10 meters
R8UL-P5 to R1SS-P15	10 meters
R8UL-P6 to R4SS-P16	10 meters
R8UL-P7 to R2SS-P15	10 meters
R8UL-P30 to R7SS-P19	5 meters
R8UL-P31 to R5SS-P19	10 meters
R8UL-P32 to R8SS-P19	3 meters
R8UL-P33 to R6SS-P19	5 meters
R8UL-P34 to R12SS-P19	10 meters
R8UL-P35 to R11SS-P19	10 meters
R8UL-P36 to R13SS-P19	10 meters
R8LL-P1 to R13SS-P18	10 meters
R8LL-P2 to R9SS-P18	5 meters
R8LL-P3 to R3SS-P26	10 meters
R8LL-P4 to R3SS-P12	10 meters
R8LL-P5 to R1SS-P7	10 meters
R8LL-P6 to R4SS-P7	10 meters
R8LL-P7 to R2SS-P9	10 meters
R8LL-P30 to R7SS-P12	5 meters
R8LL-P31 to R5SS-P12	10 meters
R8LL-P32 to R8SS-P18	3 meters
R8LL-P33 to R6SS-P12	5 meters
R8LL-P34 to R12SS-P18	10 meters
R8LL-P35 to R11SS-P18	10 meters
R8LL-P36 to R10SS-P18	5 meters

Table 22-97 Leaf Switch Connections for the Ninth Rack in a 13 Rack System

Connection	Cable Length
R9UL-P1 to R1SS-P24	10 meters
R9UL-P2 to R4SS-P26	10 meters
R9UL-P3 to R2SS-P24	10 meters
R9UL-P4 to R10SS-P24	5 meters
R9UL-P5 to R3SS-P24	10 meters
R9UL-P6 to R11SS-P24	5 meters
R9UL-P7 to R4SS-P24	10 meters
R9UL-P30 to R5SS-P24	10 meters
R9UL-P31 to R6SS-P24	10 meters
R9UL-P32 to R7SS-P24	5 meters
R9UL-P33 to R8SS-P24	5 meters
R9UL-P34 to R13SS-P24	10 meters
R9UL-P35 to R9SS-P24	3 meters
R9UL-P36 to R12SS-P24	10 meters

Table 22-97 (Cont.) Leaf Switch Connections for the Ninth Rack in a 13 Rack System

Connection	Cable Length
R9LL-P1 to R1SS-P29	10 meters
R9LL-P2 to R11SS-P29	5 meters
R9LL-P3 to R2SS-P29	10 meters
R9LL-P4 to R13SS-P29	10 meters
R9LL-P5 to R3SS-P29	10 meters
R9LL-P6 to R10SS-P29	5 meters
R9LL-P7 to R4SS-P29	10 meters
R9LL-P30 to R12SS-P29	10 meters
R9LL-P31 to R5SS-P29	10 meters
R9LL-P32 to R6SS-P29	10 meters
R9LL-P33 to R7SS-P29	5 meters
R9LL-P34 to R8SS-P29	5 meters
R9LL-P35 to R9SS-P29	3 meters
R9LL-P36 to R5SS-P26	10 meters

Table 22-98 Leaf Switch Connections for the Tenth Rack in a 13 Rack System

Connection	Cable Length
R10UL-P1 to R1SS-P21	15 meters
R10UL-P2 to R11SS-P21	5 meters
R10UL-P3 to R2SS-P21	10 meters
R10UL-P4 to R10SS-P21	3 meters
R10UL-P5 to R3SS-P21	10 meters
R10UL-P6 to R13SS-P21	10 meters
R10UL-P7 to R4SS-P21	10 meters
R10UL-P30 to R5SS-P21	10 meters
R10UL-P31 to R6SS-P21	10 meters
R10UL-P32 to R7SS-P21	10 meters
R10UL-P33 to R8SS-P21	5 meters
R10UL-P34 to R12SS-P21	5 meters
R10UL-P35 to R9SS-P21	5 meters
R10UL-P36 to R6SS-P26	10 meters
R10LL-P1 to R1SS-P31	15 meters
R10LL-P2 to R12SS-P31	5 meters
R10LL-P3 to R2SS-P31	10 meters
R10LL-P4 to R11SS-P31	5 meters
R10LL-P5 to R3SS-P31	10 meters
R10LL-P6 to R10SS-P31	3 meters
R10LL-P7 to R4SS-P31	10 meters
R10LL-P30 to R7SS-P26	10 meters
R10LL-P31 to R5SS-P31	10 meters
R10LL-P32 to R6SS-P31	10 meters
R10LL-P33 to R7SS-P31	10 meters

Table 22-98 (Cont.) Leaf Switch Connections for the Tenth Rack in a 13 Rack System

Connection	Cable Length
R10LL-P34 to R8SS-P31	5 meters
R10LL-P35 to R9SS-P31	5 meters
R10LL-P36 to R13SS-P31	10 meters

Table 22-99 Leaf Switch Connections for the Eleventh Rack in a 13 Rack System

Connection	Cable Length
R11UL-P1 to R1SS-P25	15 meters
R11UL-P2 to R11SS-P25	3 meters
R11UL-P3 to R2SS-P25	15 meters
R11UL-P4 to R10SS-P25	5 meters
R11UL-P5 to R3SS-P25	10 meters
R11UL-P6 to R13SS-P25	5 meters
R11UL-P7 to R4SS-P25	10 meters
R11UL-P30 to R5SS-P25	10 meters
R11UL-P31 to R6SS-P25	10 meters
R11UL-P32 to R7SS-P25	10 meters
R11UL-P33 to R8SS-P25	10 meters
R11UL-P34 to R12SS-P25	5 meters
R11UL-P35 to R9SS-P25	5 meters
R11UL-P36 to R8SS-P26	10 meters
R11LL-P1 to R1SS-P28	15 meters
R11LL-P2 to R13SS-P28	5 meters
R11LL-P3 to R2SS-P28	15 meters
R11LL-P4 to R11SS-P28	3 meters
R11LL-P5 to R3SS-P28	10 meters
R11LL-P6 to R10SS-P28	5 meters
R11LL-P7 to R4SS-P28	10 meters
R11LL-P30 to R9SS-P26	5 meters
R11LL-P31 to R5SS-P28	10 meters
R11LL-P32 to R6SS-P28	10 meters
R11LL-P33 to R7SS-P28	10 meters
R11LL-P34 to R8SS-P28	10 meters
R11LL-P35 to R9SS-P28	5 meters
R11LL-P36 to R12SS-P28	5 meters

Table 22-100 Leaf Switch Connections for the Twelfth Rack in a 13 Rack System

Connection	Cable Length
R12UL-P1 to R1SS-P22	15 meters
R12UL-P2 to R11SS-P22	5 meters
R12UL-P3 to R2SS-P22	15 meters
R12UL-P4 to R10SS-P22	5 meters

Table 22-100 (Cont.) Leaf Switch Connections for the Twelfth Rack in a 13 Rack System

Connection	Cable Length
R12UL-P5 to R3SS-P22	15 meters
R12UL-P6 to R10SS-P26	5 meters
R12UL-P7 to R4SS-P22	10 meters
R12UL-P30 to R5SS-P22	10 meters
R12UL-P31 to R6SS-P22	10 meters
R12UL-P32 to R7SS-P22	10 meters
R12UL-P33 to R8SS-P22	10 meters
R12UL-P34 to R12SS-P22	3 meters
R12UL-P35 to R9SS-P22	10 meters
R12UL-P36 to R13SS-P22	5 meters
R12LL-P1 to R1SS-P32	15 meters
R12LL-P2 to R11SS-P26	5 meters
R12LL-P3 to R2SS-P32	15 meters
R12LL-P4 to R11SS-P32	5 meters
R12LL-P5 to R3SS-P32	15 meters
R12LL-P6 to R10SS-P32	5 meters
R12LL-P7 to R4SS-P32	10 meters
R12LL-P30 to R13SS-P32	5 meters
R12LL-P31 to R5SS-P32	10 meters
R12LL-P32 to R6SS-P32	10 meters
R12LL-P33 to R7SS-P32	10 meters
R12LL-P34 to R8SS-P32	10 meters
R12LL-P35 to R9SS-P32	10 meters
R12LL-P36 to R12SS-P32	3 meters

Table 22-101 Leaf Switch Connections for the Thirteenth Rack in a 13 Rack System

Connection	Cable Length
R13UL-P1 to R1SS-P27	15 meters
R13UL-P2 to R11SS-P27	5 meters
R13UL-P3 to R2SS-P27	15 meters
R13UL-P4 to R10SS-P27	10 meters
R13UL-P5 to R3SS-P27	15 meters
R13UL-P6 to R12SS-P26	5 meters
R13UL-P7 to R4SS-P27	15 meters
R13UL-P30 to R5SS-P27	10 meters
R13UL-P31 to R6SS-P27	10 meters
R13UL-P32 to R7SS-P27	10 meters
R13UL-P33 to R8SS-P27	10 meters
R13UL-P34 to R12SS-P27	5 meters
R13UL-P35 to R9SS-P27	10 meters
R13UL-P36 to R13SS-P27	3 meters

Table 22-101 (Cont.) Leaf Switch Connections for the Thirteenth Rack in a 13 Rack System

Connection	Cable Length
R13LL-P1 to R1SS-P30	15 meters
R13LL-P2 to R13SS-P26	3 meters
R13LL-P3 to R2SS-P30	15 meters
R13LL-P4 to R11SS-P30	5 meters
R13LL-P5 to R3SS-P30	15 meters
R13LL-P6 to R10SS-P30	10 meters
R13LL-P7 to R4SS-P30	15 meters
R13LL-P30 to R13SS-P30	3 meters
R13LL-P31 to R5SS-P30	10 meters
R13LL-P32 to R6SS-P30	10 meters
R13LL-P33 to R7SS-P30	10 meters
R13LL-P34 to R8SS-P30	10 meters
R13LL-P35 to R9SS-P30	10 meters
R13LL-P36 to R12SS-P30	5 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-102 Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R8LL-P5	10 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R7UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R6LL-P5	10 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R8UL-P5	10 meters
R1SS-P16	R6UL-P7	10 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R7LL-P5	10 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R10UL-P1	15 meters
R1SS-P22	R12UL-P1	15 meters
R1SS-P23	R1UL-P3	3 meters
R1SS-P24	R9UL-P1	10 meters
R1SS-P25	R11UL-P1	15 meters

Table 22-102 (Cont.) Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P26	R7LL-P3	10 meters
R1SS-P27	R13UL-P1	15 meters
R1SS-P28	R11LL-P1	15 meters
R1SS-P29	R9LL-P1	10 meters
R1SS-P30	R13LL-P1	15 meters
R1SS-P31	R10LL-P1	15 meters
R1SS-P32	R12LL-P1	15 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R7UL-P7	10 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R8LL-P7	10 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R6LL-P4	10 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R8UL-P7	10 meters
R2SS-P16	R6UL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters
R2SS-P19	R7LL-P7	10 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R10UL-P3	10 meters
R2SS-P22	R12UL-P3	15 meters
R2SS-P23	R1LL-P1	5 meters
R2SS-P24	R9UL-P3	10 meters
R2SS-P25	R11UL-P3	15 meters
R2SS-P26	R8UL-P3	10 meters
R2SS-P27	R13UL-P3	15 meters
R2SS-P28	R11LL-P3	15 meters
R2SS-P29	R9LL-P3	10 meters
R2SS-P30	R13LL-P3	15 meters
R2SS-P31	R10LL-P3	10 meters
R2SS-P32	R12LL-P3	15 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R6LL-P31	10 meters
R3SS-P8	R5UL-P6	5 meters
R3SS-P9	R7UL-P4	10 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters

Table 22-102 (Cont.) Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R3SS-P12	R8LL-P4	10 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R7LL-P4	10 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R8UL-P4	10 meters
R3SS-P19	R6UL-P6	10 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R10UL-P5	10 meters
R3SS-P22	R12UL-P5	15 meters
R3SS-P23	R2UL-P2	5 meters
R3SS-P24	R9UL-P5	10 meters
R3SS-P25	R11UL-P5	10 meters
R3SS-P26	R8LL-P3	10 meters
R3SS-P27	R13UL-P5	15 meters
R3SS-P28	R11LL-P5	10 meters
R3SS-P29	R9LL-P5	10 meters
R3SS-P30	R13LL-P5	15 meters
R3SS-P31	R10LL-P5	10 meters
R3SS-P32	R12LL-P5	15 meters
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R8LL-P6	10 meters
R4SS-P8	R7UL-P6	10 meters
R4SS-P9	R6LL-P33	5 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R5UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R6UL-P31	5 meters
R4SS-P16	R8UL-P6	10 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R7LL-P6	10 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R10UL-P7	10 meters
R4SS-P22	R12UL-P7	10 meters
R4SS-P23	R2LL-P2	5 meters
R4SS-P24	R9UL-P7	10 meters
R4SS-P25	R11UL-P7	10 meters

Table 22-102 (Cont.) Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R4SS-P26	R9UL-P2	10 meters
R4SS-P27	R13UL-P7	15 meters
R4SS-P28	R11LL-P7	10 meters
R4SS-P29	R9LL-P7	10 meters
R4SS-P30	R13LL-P7	15 meters
R4SS-P31	R10LL-P7	10 meters
R4SS-P32	R12LL-P7	10 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R6LL-P30	5 meters
R5SS-P9	R7UL-P31	5 meters
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R8LL-P31	10 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R7LL-P31	5 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R6UL-P33	5 meters
R5SS-P19	R8UL-P31	10 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R10UL-P30	10 meters
R5SS-P22	R12UL-P30	10 meters
R5SS-P23	R3UL-P34	5 meters
R5SS-P24	R9UL-P30	10 meters
R5SS-P25	R11UL-P30	10 meters
R5SS-P26	R9LL-P36	10 meters
R5SS-P27	R13UL-P30	10 meters
R5SS-P28	R11LL-P31	10 meters
R5SS-P29	R9LL-P31	10 meters
R5SS-P30	R13LL-P31	10 meters
R5SS-P31	R10LL-P31	10 meters
R5SS-P32	R12LL-P31	10 meters
R6SS-P5	R1UL-P32	10 meters
R6SS-P6	R2UL-P33	10 meters
R6SS-P7	R5UL-P7	5 meters
R6SS-P8	R6LL-P32	3 meters
R6SS-P9	R7UL-P33	5 meters
R6SS-P10	R4UL-P30	5 meters
R6SS-P11	R3UL-P7	10 meters

Table 22-102 (Cont.) Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R6SS-P12	R8LL-P33	5 meters
R6SS-P13	R1LL-P32	10 meters
R6SS-P14	R2LL-P7	10 meters
R6SS-P15	R5LL-P4	5 meters
R6SS-P16	R6UL-P30	3 meters
R6SS-P17	R4LL-P6	5 meters
R6SS-P18	R7LL-P33	5 meters
R6SS-P19	R8UL-P33	5 meters
R6SS-P20	R3LL-P30	10 meters
R6SS-P21	R10UL-P31	10 meters
R6SS-P22	R12UL-P31	10 meters
R6SS-P23	R3LL-P35	10 meters
R6SS-P24	R9UL-P31	10 meters
R6SS-P25	R11UL-P31	10 meters
R6SS-P26	R10UL-P36	10 meters
R6SS-P27	R13UL-P31	10 meters
R6SS-P28	R11LL-P32	10 meters
R6SS-P29	R9LL-P32	10 meters
R6SS-P30	R13LL-P32	10 meters
R6SS-P31	R10LL-P32	10 meters
R6SS-P32	R12LL-P32	10 meters
R7SS-P5	R1UL-P33	10 meters
R7SS-P6	R2UL-P32	10 meters
R7SS-P7	R5UL-P32	5 meters
R7SS-P8	R6LL-P6	5 meters
R7SS-P9	R7UL-P30	3 meters
R7SS-P10	R4UL-P31	10 meters
R7SS-P11	R3UL-P31	10 meters
R7SS-P12	R8LL-P30	5 meters
R7SS-P13	R1LL-P6	10 meters
R7SS-P14	R2LL-P32	10 meters
R7SS-P15	R5LL-P5	5 meters
R7SS-P16	R6UL-P5	5 meters
R7SS-P17	R4LL-P32	10 meters
R7SS-P18	R7LL-P30	3 meters
R7SS-P19	R8UL-P30	5 meters
R7SS-P20	R3LL-P7	10 meters
R7SS-P21	R10UL-P32	10 meters
R7SS-P22	R12UL-P32	10 meters
R7SS-P23	R4UL-P36	10 meters
R7SS-P24	R9UL-P32	5 meters
R7SS-P25	R11UL-P32	10 meters

Table 22-102 (Cont.) Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R7SS-P26	R10LL-P30	10 meters
R7SS-P27	R13UL-P32	10 meters
R7SS-P28	R11LL-P33	10 meters
R7SS-P29	R9LL-P33	5 meters
R7SS-P30	R13LL-P33	10 meters
R7SS-P31	R10LL-P33	10 meters
R7SS-P32	R12LL-P33	10 meters
R8SS-P5	R1UL-P7	10 meters
R8SS-P6	R2UL-P30	10 meters
R8SS-P7	R5UL-P31	10 meters
R8SS-P8	R6LL-P7	5 meters
R8SS-P9	R7UL-P32	5 meters
R8SS-P10	R4UL-P4	10 meters
R8SS-P11	R3UL-P33	10 meters
R8SS-P12	R7LL-P32	5 meters
R8SS-P13	R1LL-P33	10 meters
R8SS-P14	R2LL-P30	10 meters
R8SS-P15	R5LL-P32	10 meters
R8SS-P16	R6UL-P32	5 meters
R8SS-P17	R4LL-P33	10 meters
R8SS-P18	R8LL-P32	3 meters
R8SS-P19	R8UL-P32	3 meters
R8SS-P20	R3LL-P6	10 meters
R8SS-P21	R10UL-P33	5 meters
R8SS-P22	R12UL-P33	10 meters
R8SS-P23	R4LL-P36	10 meters
R8SS-P24	R9UL-P33	5 meters
R8SS-P25	R11UL-P33	10 meters
R8SS-P26	R11UL-P36	10 meters
R8SS-P27	R13UL-P33	10 meters
R8SS-P28	R11LL-P34	10 meters
R8SS-P29	R9LL-P34	5 meters
R8SS-P30	R13LL-P34	10 meters
R8SS-P31	R10LL-P34	5 meters
R8SS-P32	R12LL-P34	10 meters
R9SS-P5	R1UL-P36	10 meters
R9SS-P6	R2UL-P36	10 meters
R9SS-P7	R5UL-P35	10 meters
R9SS-P8	R6LL-P1	10 meters
R9SS-P9	R7UL-P2	5 meters
R9SS-P10	R4UL-P34	10 meters
R9SS-P11	R3UL-P2	10 meters

Table 22-102 (Cont.) Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R9SS-P12	R7LL-P1	5 meters
R9SS-P13	R1LL-P36	10 meters
R9SS-P14	R2LL-P3	10 meters
R9SS-P15	R5LL-P1	10 meters
R9SS-P16	R6UL-P3	10 meters
R9SS-P17	R4LL-P1	10 meters
R9SS-P18	R8LL-P2	5 meters
R9SS-P19	R8UL-P1	5 meters
R9SS-P20	R3LL-P3	10 meters
R9SS-P21	R10UL-P35	5 meters
R9SS-P22	R12UL-P35	10 meters
R9SS-P23	R5UL-P36	10 meters
R9SS-P24	R9UL-P35	3 meters
R9SS-P25	R11UL-P35	5 meters
R9SS-P26	R11LL-P30	5 meters
R9SS-P27	R13UL-P35	10 meters
R9SS-P28	R11LL-P35	5 meters
R9SS-P29	R9LL-P35	3 meters
R9SS-P30	R13LL-P35	10 meters
R9SS-P31	R10LL-P35	5 meters
R9SS-P32	R12LL-P35	10 meters
R10SS-P5	R1UL-P35	15 meters
R10SS-P6	R2UL-P34	10 meters
R10SS-P7	R5UL-P1	10 meters
R10SS-P8	R6LL-P2	10 meters
R10SS-P9	R7UL-P36	10 meters
R10SS-P10	R4UL-P3	10 meters
R10SS-P11	R3UL-P3	10 meters
R10SS-P12	R7LL-P2	10 meters
R10SS-P13	R1LL-P35	15 meters
R10SS-P14	R2LL-P36	10 meters
R10SS-P15	R5LL-P35	10 meters
R10SS-P16	R6UL-P35	10 meters
R10SS-P17	R4LL-P2	10 meters
R10SS-P18	R8LL-P36	5 meters
R10SS-P19	R8UL-P2	5 meters
R10SS-P20	R3LL-P2	10 meters
R10SS-P21	R10UL-P4	3 meters
R10SS-P22	R12UL-P4	5 meters
R10SS-P23	R5LL-P36	10 meters
R10SS-P24	R9UL-P4	5 meters
R10SS-P25	R11UL-P4	5 meters

Table 22-102 (Cont.) Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R10SS-P26	R12UL-P6	5 meters
R10SS-P27	R13UL-P4	10 meters
R10SS-P28	R11LL-P6	5 meters
R10SS-P29	R9LL-P6	5 meters
R10SS-P30	R13LL-P6	10 meters
R10SS-P31	R10LL-P6	3 meters
R10SS-P32	R12LL-P6	5 meters
R11SS-P5	R1UL-P2	15 meters
R11SS-P6	R2UL-P35	15 meters
R11SS-P7	R5UL-P3	10 meters
R11SS-P8	R6LL-P35	10 meters
R11SS-P9	R7UL-P35	10 meters
R11SS-P10	R4UL-P2	10 meters
R11SS-P11	R3UL-P36	10 meters
R11SS-P12	R7LL-P35	10 meters
R11SS-P13	R1LL-P3	15 meters
R11SS-P14	R2LL-P34	15 meters
R11SS-P15	R5LL-P3	10 meters
R11SS-P16	R6UL-P2	10 meters
R11SS-P17	R4LL-P3	10 meters
R11SS-P18	R8LL-P35	10 meters
R11SS-P19	R8UL-P35	10 meters
R11SS-P20	R3LL-P34	10 meters
R11SS-P21	R10UL-P2	5 meters
R11SS-P22	R12UL-P2	5 meters
R11SS-P23	R6UL-P1	10 meters
R11SS-P24	R9UL-P6	5 meters
R11SS-P25	R11UL-P2	3 meters
R11SS-P26	R12LL-P2	5 meters
R11SS-P27	R13UL-P2	5 meters
R11SS-P28	R11LL-P4	3 meters
R11SS-P29	R9LL-P2	5 meters
R11SS-P30	R13LL-P4	5 meters
R11SS-P31	R10LL-P4	5 meters
R11SS-P32	R12LL-P4	5 meters
R12SS-P5	R1UL-P1	15 meters
R12SS-P6	R2UL-P1	15 meters
R12SS-P7	R5UL-P2	10 meters
R12SS-P8	R6LL-P36	10 meters
R12SS-P9	R7UL-P34	10 meters
R12SS-P10	R4UL-P35	10 meters
R12SS-P11	R3UL-P35	15 meters

Table 22-102 (Cont.) Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R12SS-P12	R7LL-P34	10 meters
R12SS-P13	R1LL-P34	15 meters
R12SS-P14	R2LL-P1	15 meters
R12SS-P15	R5LL-P2	10 meters
R12SS-P16	R6UL-P34	10 meters
R12SS-P17	R4LL-P34	10 meters
R12SS-P18	R8LL-P34	10 meters
R12SS-P19	R8UL-P34	10 meters
R12SS-P20	R3LL-P36	15 meters
R12SS-P21	R10UL-P34	5 meters
R12SS-P22	R12UL-P34	3 meters
R12SS-P23	R6LL-P34	10 meters
R12SS-P24	R9UL-P36	10 meters
R12SS-P25	R11UL-P34	5 meters
R12SS-P26	R13UL-P6	5 meters
R12SS-P27	R13UL-P34	5 meters
R12SS-P28	R11LL-P36	5 meters
R12SS-P29	R9LL-P30	10 meters
R12SS-P30	R13LL-P36	5 meters
R12SS-P31	R10LL-P2	5 meters
R12SS-P32	R12LL-P36	3 meters
R13SS-P5	R1UL-P34	15 meters
R13SS-P6	R2UL-P3	15 meters
R13SS-P7	R5UL-P34	10 meters
R13SS-P8	R6LL-P3	10 meters
R13SS-P9	R7UL-P3	10 meters
R13SS-P10	R4UL-P1	15 meters
R13SS-P11	R3UL-P1	15 meters
R13SS-P12	R7LL-P36	10 meters
R13SS-P13	R1LL-P2	15 meters
R13SS-P14	R2LL-P35	15 meters
R13SS-P15	R5LL-P34	10 meters
R13SS-P16	R6UL-P36	10 meters
R13SS-P17	R4LL-P35	15 meters
R13SS-P18	R8LL-P1	10 meters
R13SS-P19	R8UL-P36	10 meters
R13SS-P20	R3LL-P1	15 meters
R13SS-P21	R10UL-P6	10 meters
R13SS-P22	R12UL-P36	5 meters
R13SS-P23	R7UL-P1	10 meters
R13SS-P24	R9UL-P34	10 meters
R13SS-P25	R11UL-P6	5 meters

Table 22-102 (Cont.) Thirteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R13SS-P26	R13LL-P2	3 meters
R13SS-P27	R13UL-P36	3 meters
R13SS-P28	R11LL-P2	5 meters
R13SS-P29	R9LL-P4	10 meters
R13SS-P30	R13LL-P30	3 meters
R13SS-P31	R10LL-P36	10 meters
R13SS-P32	R12LL-P30	5 meters

Fourteen-Rack Cabling for RA21 and Later Model Racks

This section provides the cabling details for RA21 and later models to connect fourteen racks using RoCE Network Fabric.

Note:

- The following conventions are used in the cabling notation for connecting multiple racks together:
 - The abbreviation for the first rack is R1, the second rack is R2, and so on.
 - LL identifies a lower leaf switch and UL identifies an upper leaf switch.
 - SS identifies the spine switch, which is located in U1 on all racks.
 - A specific switch is identified by combining abbreviations. For example, R1LL identifies the lower leaf switch (LL) on the first rack (R1).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in 2-socket systems (ZDLRA Rack X9M and later models).
 - At rack unit 21 (U21) and rack unit 23 (U23) in 8-socket systems (Recovery Appliance X9M-8).
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X9M-8 systems with three database servers and 11 storage servers only, the database servers and storage server require 23 leaf switch ports, which leaves only 13 inter-switch links on each leaf switch. Consequently, these systems cannot support 14 interconnected racks.

The following tables contain details for all of the RoCE Network Fabric cabling connections in a 14 rack system.

Table 22-103 Leaf Switch Connections for the First Rack in a 14 Rack System

Connection	Cable Length
R1UL-P1 to R12SS-P5	15 meters
R1UL-P2 to R11SS-P5	15 meters
R1UL-P3 to R14SS-P5	15 meters
R1UL-P4 to R4SS-P5	10 meters
R1UL-P5 to R1SS-P5	3 meters
R1UL-P6 to R3SS-P5	5 meters
R1UL-P7 to R8SS-P5	10 meters
R1UL-P30 to R5SS-P5	10 meters
R1UL-P31 to R2SS-P5	5 meters
R1UL-P32 to R6SS-P5	10 meters
R1UL-P33 to R7SS-P5	10 meters
R1UL-P34 to R13SS-P5	15 meters
R1UL-P35 to R10SS-P5	15 meters
R1UL-P36 to R9SS-P5	10 meters
R1LL-P1 to R14SS-P13	15 meters
R1LL-P2 to R13SS-P13	15 meters
R1LL-P3 to R11SS-P13	15 meters
R1LL-P4 to R3SS-P13	5 meters
R1LL-P5 to R1SS-P13	3 meters
R1LL-P6 to R7SS-P13	10 meters
R1LL-P7 to R5SS-P13	10 meters
R1LL-P30 to R4SS-P13	10 meters
R1LL-P31 to R2SS-P13	5 meters
R1LL-P32 to R6SS-P13	10 meters
R1LL-P33 to R8SS-P13	10 meters
R1LL-P34 to R12SS-P13	15 meters
R1LL-P35 to R10SS-P13	15 meters
R1LL-P36 to R9SS-P13	10 meters

Table 22-104 Leaf Switch Connections for the Second Rack in a 14 Rack System

Connection	Cable Length
R2UL-P1 to R12SS-P6	15 meters
R2UL-P2 to R14SS-P6	15 meters
R2UL-P3 to R13SS-P6	15 meters
R2UL-P4 to R4SS-P6	5 meters
R2UL-P5 to R1SS-P6	5 meters
R2UL-P6 to R3SS-P6	5 meters
R2UL-P7 to R5SS-P6	10 meters
R2UL-P30 to R8SS-P6	10 meters
R2UL-P31 to R2SS-P6	3 meters

Table 22-104 (Cont.) Leaf Switch Connections for the Second Rack in a 14 Rack System

Connection	Cable Length
R2UL-P32 to R7SS-P6	10 meters
R2UL-P33 to R6SS-P6	10 meters
R2UL-P34 to R10SS-P6	10 meters
R2UL-P35 to R11SS-P6	15 meters
R2UL-P36 to R9SS-P6	10 meters
R2LL-P1 to R12SS-P14	15 meters
R2LL-P2 to R14SS-P14	15 meters
R2LL-P3 to R9SS-P14	10 meters
R2LL-P4 to R4SS-P14	5 meters
R2LL-P5 to R1SS-P14	5 meters
R2LL-P6 to R3SS-P14	5 meters
R2LL-P7 to R6SS-P14	10 meters
R2LL-P30 to R8SS-P14	10 meters
R2LL-P31 to R2SS-P14	3 meters
R2LL-P32 to R7SS-P14	10 meters
R2LL-P33 to R5SS-P14	10 meters
R2LL-P34 to R11SS-P14	15 meters
R2LL-P35 to R13SS-P14	15 meters
R2LL-P36 to R10SS-P14	10 meters

Table 22-105 Leaf Switch Connections for the Third Rack in a 14 Rack System

Connection	Cable Length
R3UL-P1 to R13SS-P11	15 meters
R3UL-P2 to R9SS-P11	10 meters
R3UL-P3 to R10SS-P11	10 meters
R3UL-P4 to R4SS-P11	5 meters
R3UL-P5 to R1SS-P11	5 meters
R3UL-P6 to R2SS-P11	5 meters
R3UL-P7 to R6SS-P11	10 meters
R3UL-P30 to R3SS-P11	3 meters
R3UL-P31 to R7SS-P11	10 meters
R3UL-P32 to R5SS-P11	5 meters
R3UL-P33 to R8SS-P11	10 meters
R3UL-P34 to R14SS-P11	15 meters
R3UL-P35 to R12SS-P11	15 meters
R3UL-P36 to R11SS-P11	10 meters
R3LL-P1 to R13SS-P20	15 meters
R3LL-P2 to R10SS-P20	10 meters
R3LL-P3 to R9SS-P20	10 meters
R3LL-P4 to R1SS-P20	5 meters
R3LL-P5 to R5SS-P20	5 meters

Table 22-105 (Cont.) Leaf Switch Connections for the Third Rack in a 14 Rack System

Connection	Cable Length
R3LL-P6 to R8SS-P20	10 meters
R3LL-P7 to R7SS-P20	10 meters
R3LL-P30 to R6SS-P20	10 meters
R3LL-P31 to R2SS-P20	5 meters
R3LL-P32 to R3SS-P20	3 meters
R3LL-P33 to R4SS-P20	5 meters
R3LL-P34 to R11SS-P20	10 meters
R3LL-P35 to R14SS-P20	15 meters
R3LL-P36 to R12SS-P20	15 meters

Table 22-106 Leaf Switch Connections for the Fourth Rack in a 14 Rack System

Connection	Cable Length
R4UL-P1 to R13SS-P10	15 meters
R4UL-P2 to R11SS-P10	10 meters
R4UL-P3 to R10SS-P10	10 meters
R4UL-P4 to R8SS-P10	10 meters
R4UL-P5 to R5SS-P10	5 meters
R4UL-P6 to R2SS-P10	5 meters
R4UL-P7 to R1SS-P10	10 meters
R4UL-P30 to R6SS-P10	5 meters
R4UL-P31 to R7SS-P10	10 meters
R4UL-P32 to R4SS-P10	3 meters
R4UL-P33 to R3SS-P10	5 meters
R4UL-P34 to R9SS-P10	10 meters
R4UL-P35 to R12SS-P10	10 meters
R4UL-P36 to R14SS-P10	15 meters
R4LL-P1 to R9SS-P17	10 meters
R4LL-P2 to R10SS-P17	10 meters
R4LL-P3 to R11SS-P17	10 meters
R4LL-P4 to R2SS-P17	5 meters
R4LL-P5 to R1SS-P17	10 meters
R4LL-P6 to R6SS-P17	5 meters
R4LL-P7 to R5SS-P17	5 meters
R4LL-P30 to R4SS-P17	3 meters
R4LL-P31 to R3SS-P17	5 meters
R4LL-P32 to R7SS-P17	10 meters
R4LL-P33 to R8SS-P17	10 meters
R4LL-P34 to R12SS-P17	10 meters
R4LL-P35 to R13SS-P17	15 meters
R4LL-P36 to R14SS-P17	15 meters

Table 22-107 Leaf Switch Connections for the Fifth Rack in a 14 Rack System

Connection	Cable Length
R5UL-P1 to R10SS-P7	10 meters
R5UL-P2 to R12SS-P7	10 meters
R5UL-P3 to R11SS-P7	10 meters
R5UL-P4 to R2SS-P8	10 meters
R5UL-P5 to R1SS-P8	10 meters
R5UL-P6 to R3SS-P8	5 meters
R5UL-P7 to R6SS-P7	5 meters
R5UL-P30 to R5SS-P7	3 meters
R5UL-P31 to R8SS-P7	10 meters
R5UL-P32 to R7SS-P7	5 meters
R5UL-P33 to R4SS-P12	5 meters
R5UL-P34 to R13SS-P7	10 meters
R5UL-P35 to R9SS-P7	10 meters
R5UL-P36 to R14SS-P7	15 meters
R5LL-P1 to R9SS-P15	10 meters
R5LL-P2 to R12SS-P15	10 meters
R5LL-P3 to R11SS-P15	10 meters
R5LL-P4 to R6SS-P15	5 meters
R5LL-P5 to R7SS-P15	5 meters
R5LL-P6 to R2SS-P18	10 meters
R5LL-P7 to R1SS-P18	10 meters
R5LL-P30 to R5SS-P15	3 meters
R5LL-P31 to R3SS-P16	5 meters
R5LL-P32 to R8SS-P15	10 meters
R5LL-P33 to R4SS-P19	5 meters
R5LL-P34 to R13SS-P15	10 meters
R5LL-P35 to R10SS-P15	10 meters
R5LL-P36 to R14SS-P15	15 meters

Table 22-108 Leaf Switch Connections for the Sixth Rack in a 14 Rack System

Connection	Cable Length
R6UL-P1 to R14SS-P16	10 meters
R6UL-P2 to R11SS-P16	10 meters
R6UL-P3 to R9SS-P16	10 meters
R6UL-P4 to R2SS-P16	10 meters
R6UL-P5 to R7SS-P16	5 meters
R6UL-P6 to R3SS-P19	10 meters
R6UL-P7 to R1SS-P16	10 meters
R6UL-P30 to R6SS-P16	3 meters
R6UL-P31 to R4SS-P15	5 meters
R6UL-P32 to R8SS-P16	5 meters
R6UL-P33 to R5SS-P18	5 meters

Table 22-108 (Cont.) Leaf Switch Connections for the Sixth Rack in a 14 Rack System

Connection	Cable Length
R6UL-P34 to R12SS-P16	10 meters
R6UL-P35 to R10SS-P16	10 meters
R6UL-P36 to R13SS-P16	10 meters
R6LL-P1 to R9SS-P8	10 meters
R6LL-P2 to R10SS-P8	10 meters
R6LL-P3 to R13SS-P8	10 meters
R6LL-P4 to R2SS-P12	10 meters
R6LL-P5 to R1SS-P12	10 meters
R6LL-P6 to R7SS-P8	5 meters
R6LL-P7 to R8SS-P8	5 meters
R6LL-P30 to R5SS-P8	5 meters
R6LL-P31 to R3SS-P7	10 meters
R6LL-P32 to R6SS-P8	3 meters
R6LL-P33 to R4SS-P9	5 meters
R6LL-P34 to R14SS-P8	10 meters
R6LL-P35 to R11SS-P8	10 meters
R6LL-P36 to R12SS-P8	10 meters

Table 22-109 Leaf Switch Connections for the Seventh Rack in a 14 Rack System

Connection	Cable Length
R7UL-P1 to R14SS-P9	10 meters
R7UL-P2 to R9SS-P9	5 meters
R7UL-P3 to R13SS-P9	10 meters
R7UL-P4 to R3SS-P9	10 meters
R7UL-P5 to R1SS-P9	10 meters
R7UL-P6 to R4SS-P8	10 meters
R7UL-P7 to R2SS-P7	10 meters
R7UL-P30 to R7SS-P9	3 meters
R7UL-P31 to R5SS-P9	5 meters
R7UL-P32 to R8SS-P9	5 meters
R7UL-P33 to R6SS-P9	5 meters
R7UL-P34 to R12SS-P9	10 meters
R7UL-P35 to R11SS-P9	10 meters
R7UL-P36 to R10SS-P9	10 meters
R7LL-P1 to R9SS-P12	5 meters
R7LL-P2 to R10SS-P12	10 meters
R7LL-P3 to R14SS-P12	10 meters
R7LL-P4 to R3SS-P15	10 meters
R7LL-P5 to R1SS-P19	10 meters
R7LL-P6 to R4SS-P18	10 meters
R7LL-P7 to R2SS-P19	10 meters
R7LL-P30 to R7SS-P18	3 meters

Table 22-109 (Cont.) Leaf Switch Connections for the Seventh Rack in a 14 Rack System

Connection	Cable Length
R7LL-P31 to R5SS-P16	5 meters
R7LL-P32 to R8SS-P12	5 meters
R7LL-P33 to R6SS-P18	5 meters
R7LL-P34 to R12SS-P12	10 meters
R7LL-P35 to R11SS-P12	10 meters
R7LL-P36 to R13SS-P12	10 meters

Table 22-110 Leaf Switch Connections for the Eighth Rack in a 14 Rack System

Connection	Cable Length
R8UL-P1 to R9SS-P19	5 meters
R8UL-P2 to R10SS-P19	5 meters
R8UL-P3 to R14SS-P19	10 meters
R8UL-P4 to R3SS-P18	10 meters
R8UL-P5 to R1SS-P15	10 meters
R8UL-P6 to R4SS-P16	10 meters
R8UL-P7 to R2SS-P15	10 meters
R8UL-P30 to R7SS-P19	5 meters
R8UL-P31 to R5SS-P19	10 meters
R8UL-P32 to R8SS-P19	3 meters
R8UL-P33 to R6SS-P19	5 meters
R8UL-P34 to R12SS-P19	10 meters
R8UL-P35 to R11SS-P19	10 meters
R8UL-P36 to R13SS-P19	10 meters
R8LL-P1 to R13SS-P18	10 meters
R8LL-P2 to R9SS-P18	5 meters
R8LL-P3 to R14SS-P18	10 meters
R8LL-P4 to R3SS-P12	10 meters
R8LL-P5 to R1SS-P7	10 meters
R8LL-P6 to R4SS-P7	10 meters
R8LL-P7 to R2SS-P9	10 meters
R8LL-P30 to R7SS-P12	5 meters
R8LL-P31 to R5SS-P12	10 meters
R8LL-P32 to R8SS-P18	3 meters
R8LL-P33 to R6SS-P12	5 meters
R8LL-P34 to R12SS-P18	10 meters
R8LL-P35 to R11SS-P18	10 meters
R8LL-P36 to R10SS-P18	5 meters

Table 22-111 Leaf Switch Connections for the Ninth Rack in a 14 Rack System

Connection	Cable Length
R9UL-P1 to R1SS-P21	10 meters
R9UL-P2 to R14SS-P21	10 meters
R9UL-P3 to R2SS-P21	10 meters
R9UL-P4 to R10SS-P21	5 meters
R9UL-P5 to R3SS-P21	10 meters
R9UL-P6 to R11SS-P21	5 meters
R9UL-P7 to R4SS-P21	10 meters
R9UL-P30 to R5SS-P21	10 meters
R9UL-P31 to R6SS-P21	10 meters
R9UL-P32 to R7SS-P21	5 meters
R9UL-P33 to R8SS-P21	5 meters
R9UL-P34 to R13SS-P21	10 meters
R9UL-P35 to R9SS-P21	3 meters
R9UL-P36 to R12SS-P21	10 meters
R9LL-P1 to R1SS-P22	10 meters
R9LL-P2 to R11SS-P22	5 meters
R9LL-P3 to R2SS-P22	10 meters
R9LL-P4 to R13SS-P22	10 meters
R9LL-P5 to R3SS-P22	10 meters
R9LL-P6 to R10SS-P22	5 meters
R9LL-P7 to R4SS-P22	10 meters
R9LL-P30 to R12SS-P22	10 meters
R9LL-P31 to R5SS-P22	10 meters
R9LL-P32 to R6SS-P22	10 meters
R9LL-P33 to R7SS-P22	5 meters
R9LL-P34 to R8SS-P22	5 meters
R9LL-P35 to R9SS-P22	3 meters
R9LL-P36 to R14SS-P22	10 meters

Table 22-112 Leaf Switch Connections for the Tenth Rack in a 14 Rack System

Connection	Cable Length
R10UL-P1 to R1SS-P23	15 meters
R10UL-P2 to R11SS-P23	5 meters
R10UL-P3 to R2SS-P23	10 meters
R10UL-P4 to R10SS-P23	3 meters
R10UL-P5 to R3SS-P23	10 meters
R10UL-P6 to R13SS-P23	10 meters
R10UL-P7 to R4SS-P23	10 meters
R10UL-P30 to R5SS-P23	10 meters
R10UL-P31 to R6SS-P23	10 meters
R10UL-P32 to R7SS-P23	10 meters
R10UL-P33 to R8SS-P23	5 meters

Table 22-112 (Cont.) Leaf Switch Connections for the Tenth Rack in a 14 Rack System

Connection	Cable Length
R10UL-P34 to R12SS-P23	5 meters
R10UL-P35 to R9SS-P23	5 meters
R10UL-P36 to R14SS-P23	10 meters
R10LL-P1 to R1SS-P24	15 meters
R10LL-P2 to R12SS-P24	5 meters
R10LL-P3 to R2SS-P24	10 meters
R10LL-P4 to R11SS-P24	5 meters
R10LL-P5 to R3SS-P24	10 meters
R10LL-P6 to R10SS-P24	3 meters
R10LL-P7 to R4SS-P24	10 meters
R10LL-P30 to R14SS-P24	10 meters
R10LL-P31 to R5SS-P24	10 meters
R10LL-P32 to R6SS-P24	10 meters
R10LL-P33 to R7SS-P24	10 meters
R10LL-P34 to R8SS-P24	5 meters
R10LL-P35 to R9SS-P24	5 meters
R10LL-P36 to R13SS-P24	10 meters

Table 22-113 Leaf Switch Connections for the Eleventh Rack in a 14 Rack System

Connection	Cable Length
R11UL-P1 to R1SS-P25	15 meters
R11UL-P2 to R11SS-P25	3 meters
R11UL-P3 to R2SS-P25	15 meters
R11UL-P4 to R10SS-P25	5 meters
R11UL-P5 to R3SS-P25	10 meters
R11UL-P6 to R13SS-P25	5 meters
R11UL-P7 to R4SS-P25	10 meters
R11UL-P30 to R5SS-P25	10 meters
R11UL-P31 to R6SS-P25	10 meters
R11UL-P32 to R7SS-P25	10 meters
R11UL-P33 to R8SS-P25	10 meters
R11UL-P34 to R12SS-P25	5 meters
R11UL-P35 to R9SS-P25	5 meters
R11UL-P36 to R14SS-P25	10 meters
R11LL-P1 to R1SS-P26	15 meters
R11LL-P2 to R13SS-P26	5 meters
R11LL-P3 to R2SS-P26	15 meters
R11LL-P4 to R11SS-P26	3 meters
R11LL-P5 to R3SS-P26	10 meters
R11LL-P6 to R10SS-P26	5 meters
R11LL-P7 to R4SS-P26	10 meters
R11LL-P30 to R14SS-P26	10 meters

Table 22-113 (Cont.) Leaf Switch Connections for the Eleventh Rack in a 14 Rack System

Connection	Cable Length
R11LL-P31 to R5SS-P26	10 meters
R11LL-P32 to R6SS-P26	10 meters
R11LL-P33 to R7SS-P26	10 meters
R11LL-P34 to R8SS-P26	10 meters
R11LL-P35 to R9SS-P26	5 meters
R11LL-P36 to R12SS-P26	5 meters

Table 22-114 Leaf Switch Connections for the Twelfth Rack in a 14 Rack System

Connection	Cable Length
R12UL-P1 to R1SS-P27	15 meters
R12UL-P2 to R11SS-P27	5 meters
R12UL-P3 to R2SS-P27	15 meters
R12UL-P4 to R10SS-P27	5 meters
R12UL-P5 to R3SS-P27	15 meters
R12UL-P6 to R14SS-P27	5 meters
R12UL-P7 to R4SS-P27	10 meters
R12UL-P30 to R5SS-P27	10 meters
R12UL-P31 to R6SS-P27	10 meters
R12UL-P32 to R7SS-P27	10 meters
R12UL-P33 to R8SS-P27	10 meters
R12UL-P34 to R12SS-P27	3 meters
R12UL-P35 to R9SS-P27	10 meters
R12UL-P36 to R13SS-P27	5 meters
R12LL-P1 to R1SS-P28	15 meters
R12LL-P2 to R14SS-P28	5 meters
R12LL-P3 to R2SS-P28	15 meters
R12LL-P4 to R11SS-P28	5 meters
R12LL-P5 to R3SS-P28	15 meters
R12LL-P6 to R10SS-P28	5 meters
R12LL-P7 to R4SS-P28	10 meters
R12LL-P30 to R13SS-P28	5 meters
R12LL-P31 to R5SS-P28	10 meters
R12LL-P32 to R6SS-P28	10 meters
R12LL-P33 to R7SS-P28	10 meters
R12LL-P34 to R8SS-P28	10 meters
R12LL-P35 to R9SS-P28	10 meters
R12LL-P36 to R12SS-P28	3 meters

Table 22-115 Leaf Switch Connections for the Thirteenth Rack in a 14 Rack System

Connection	Cable Length
R13UL-P1 to R1SS-P29	15 meters
R13UL-P2 to R11SS-P29	5 meters
R13UL-P3 to R2SS-P29	15 meters
R13UL-P4 to R10SS-P29	10 meters
R13UL-P5 to R3SS-P29	15 meters
R13UL-P6 to R14SS-P29	5 meters
R13UL-P7 to R4SS-P29	15 meters
R13UL-P30 to R5SS-P29	10 meters
R13UL-P31 to R6SS-P29	10 meters
R13UL-P32 to R7SS-P29	10 meters
R13UL-P33 to R8SS-P29	10 meters
R13UL-P34 to R12SS-P29	5 meters
R13UL-P35 to R9SS-P29	10 meters
R13UL-P36 to R13SS-P29	3 meters
R13LL-P1 to R1SS-P30	15 meters
R13LL-P2 to R14SS-P30	5 meters
R13LL-P3 to R2SS-P30	15 meters
R13LL-P4 to R11SS-P30	5 meters
R13LL-P5 to R3SS-P30	15 meters
R13LL-P6 to R10SS-P30	10 meters
R13LL-P7 to R4SS-P30	15 meters
R13LL-P30 to R13SS-P30	3 meters
R13LL-P31 to R5SS-P30	10 meters
R13LL-P32 to R6SS-P30	10 meters
R13LL-P33 to R7SS-P30	10 meters
R13LL-P34 to R8SS-P30	10 meters
R13LL-P35 to R9SS-P30	10 meters
R13LL-P36 to R12SS-P30	5 meters

Table 22-116 Leaf Switch Connections for the Fourteenth Rack in a 14 Rack System

Connection	Cable Length
R14UL-P1 to R1SS-P31	15 meters
R14UL-P2 to R11SS-P31	10 meters
R14UL-P3 to R2SS-P31	15 meters
R14UL-P4 to R10SS-P31	10 meters
R14UL-P5 to R3SS-P31	15 meters
R14UL-P6 to R14SS-P31	3 meters
R14UL-P7 to R4SS-P31	15 meters
R14UL-P30 to R5SS-P31	15 meters
R14UL-P31 to R6SS-P31	10 meters
R14UL-P32 to R7SS-P31	10 meters
R14UL-P33 to R8SS-P31	10 meters

Table 22-116 (Cont.) Leaf Switch Connections for the Fourteenth Rack in a 14 Rack System

Connection	Cable Length
R14UL-P34 to R12SS-P31	5 meters
R14UL-P35 to R9SS-P31	10 meters
R14UL-P36 to R13SS-P31	5 meters
R14LL-P1 to R1SS-P32	15 meters
R14LL-P2 to R14SS-P32	3 meters
R14LL-P3 to R2SS-P32	15 meters
R14LL-P4 to R11SS-P32	10 meters
R14LL-P5 to R3SS-P32	15 meters
R14LL-P6 to R10SS-P32	10 meters
R14LL-P7 to R4SS-P32	15 meters
R14LL-P30 to R13SS-P32	5 meters
R14LL-P31 to R5SS-P32	15 meters
R14LL-P32 to R6SS-P32	10 meters
R14LL-P33 to R7SS-P32	10 meters
R14LL-P34 to R8SS-P32	10 meters
R14LL-P35 to R9SS-P32	10 meters
R14LL-P36 to R12SS-P32	5 meters

The following table contains all of the RoCE Network Fabric cabling connections from the previous tables. In this table, the connections are sorted by the spine switch port location.

Table 22-117 Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P5	R1UL-P5	3 meters
R1SS-P6	R2UL-P5	5 meters
R1SS-P7	R8LL-P5	10 meters
R1SS-P8	R5UL-P5	10 meters
R1SS-P9	R7UL-P5	10 meters
R1SS-P10	R4UL-P7	10 meters
R1SS-P11	R3UL-P5	5 meters
R1SS-P12	R6LL-P5	10 meters
R1SS-P13	R1LL-P5	3 meters
R1SS-P14	R2LL-P5	5 meters
R1SS-P15	R8UL-P5	10 meters
R1SS-P16	R6UL-P7	10 meters
R1SS-P17	R4LL-P5	10 meters
R1SS-P18	R5LL-P7	10 meters
R1SS-P19	R7LL-P5	10 meters
R1SS-P20	R3LL-P4	5 meters
R1SS-P21	R9UL-P1	10 meters
R1SS-P22	R9LL-P1	10 meters

Table 22-117 (Cont.) Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R1SS-P23	R10UL-P1	15 meters
R1SS-P24	R10LL-P1	15 meters
R1SS-P25	R11UL-P1	15 meters
R1SS-P26	R11LL-P1	15 meters
R1SS-P27	R12UL-P1	15 meters
R1SS-P28	R12LL-P1	15 meters
R1SS-P29	R13UL-P1	15 meters
R1SS-P30	R13LL-P1	15 meters
R1SS-P31	R14UL-P1	15 meters
R1SS-P32	R14LL-P1	15 meters
R2SS-P5	R1UL-P31	5 meters
R2SS-P6	R2UL-P31	3 meters
R2SS-P7	R7UL-P7	10 meters
R2SS-P8	R5UL-P4	10 meters
R2SS-P9	R8LL-P7	10 meters
R2SS-P10	R4UL-P6	5 meters
R2SS-P11	R3UL-P6	5 meters
R2SS-P12	R6LL-P4	10 meters
R2SS-P13	R1LL-P31	5 meters
R2SS-P14	R2LL-P31	3 meters
R2SS-P15	R8UL-P7	10 meters
R2SS-P16	R6UL-P4	10 meters
R2SS-P17	R4LL-P4	5 meters
R2SS-P18	R5LL-P6	10 meters
R2SS-P19	R7LL-P7	10 meters
R2SS-P20	R3LL-P31	5 meters
R2SS-P21	R9UL-P3	10 meters
R2SS-P22	R9LL-P3	10 meters
R2SS-P23	R10UL-P3	10 meters
R2SS-P24	R10LL-P3	10 meters
R2SS-P25	R11UL-P3	15 meters
R2SS-P26	R11LL-P3	15 meters
R2SS-P27	R12UL-P3	15 meters
R2SS-P28	R12LL-P3	15 meters
R2SS-P29	R13UL-P3	15 meters
R2SS-P30	R13LL-P3	15 meters
R2SS-P31	R14UL-P3	15 meters
R2SS-P32	R14LL-P3	15 meters
R3SS-P5	R1UL-P6	5 meters
R3SS-P6	R2UL-P6	5 meters
R3SS-P7	R6LL-P31	10 meters
R3SS-P8	R5UL-P6	5 meters

Table 22-117 (Cont.) Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R3SS-P9	R7UL-P4	10 meters
R3SS-P10	R4UL-P33	5 meters
R3SS-P11	R3UL-P30	3 meters
R3SS-P12	R8LL-P4	10 meters
R3SS-P13	R1LL-P4	5 meters
R3SS-P14	R2LL-P6	5 meters
R3SS-P15	R7LL-P4	10 meters
R3SS-P16	R5LL-P31	5 meters
R3SS-P17	R4LL-P31	5 meters
R3SS-P18	R8UL-P4	10 meters
R3SS-P19	R6UL-P6	10 meters
R3SS-P20	R3LL-P32	3 meters
R3SS-P21	R9UL-P5	10 meters
R3SS-P22	R9LL-P5	10 meters
R3SS-P23	R10UL-P5	10 meters
R3SS-P24	R10LL-P5	10 meters
R3SS-P25	R11UL-P5	10 meters
R3SS-P26	R11LL-P5	10 meters
R3SS-P27	R12UL-P5	15 meters
R3SS-P28	R12LL-P5	15 meters
R3SS-P29	R13UL-P5	15 meters
R3SS-P30	R13LL-P5	15 meters
R3SS-P31	R14UL-P5	15 meters
R3SS-P32	R14LL-P5	15 meters
R4SS-P5	R1UL-P4	10 meters
R4SS-P6	R2UL-P4	5 meters
R4SS-P7	R8LL-P6	10 meters
R4SS-P8	R7UL-P6	10 meters
R4SS-P9	R6LL-P33	5 meters
R4SS-P10	R4UL-P32	3 meters
R4SS-P11	R3UL-P4	5 meters
R4SS-P12	R5UL-P33	5 meters
R4SS-P13	R1LL-P30	10 meters
R4SS-P14	R2LL-P4	5 meters
R4SS-P15	R6UL-P31	5 meters
R4SS-P16	R8UL-P6	10 meters
R4SS-P17	R4LL-P30	3 meters
R4SS-P18	R7LL-P6	10 meters
R4SS-P19	R5LL-P33	5 meters
R4SS-P20	R3LL-P33	5 meters
R4SS-P21	R9UL-P7	10 meters
R4SS-P22	R9LL-P7	10 meters

Table 22-117 (Cont.) Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R4SS-P23	R10UL-P7	10 meters
R4SS-P24	R10LL-P7	10 meters
R4SS-P25	R11UL-P7	10 meters
R4SS-P26	R11LL-P7	10 meters
R4SS-P27	R12UL-P7	10 meters
R4SS-P28	R12LL-P7	10 meters
R4SS-P29	R13UL-P7	15 meters
R4SS-P30	R13LL-P7	15 meters
R4SS-P31	R14UL-P7	15 meters
R4SS-P32	R14LL-P7	15 meters
R5SS-P5	R1UL-P30	10 meters
R5SS-P6	R2UL-P7	10 meters
R5SS-P7	R5UL-P30	3 meters
R5SS-P8	R6LL-P30	5 meters
R5SS-P9	R7UL-P31	5 meters
R5SS-P10	R4UL-P5	5 meters
R5SS-P11	R3UL-P32	5 meters
R5SS-P12	R8LL-P31	10 meters
R5SS-P13	R1LL-P7	10 meters
R5SS-P14	R2LL-P33	10 meters
R5SS-P15	R5LL-P30	3 meters
R5SS-P16	R7LL-P31	5 meters
R5SS-P17	R4LL-P7	5 meters
R5SS-P18	R6UL-P33	5 meters
R5SS-P19	R8UL-P31	10 meters
R5SS-P20	R3LL-P5	5 meters
R5SS-P21	R9UL-P30	10 meters
R5SS-P22	R9LL-P31	10 meters
R5SS-P23	R10UL-P30	10 meters
R5SS-P24	R10LL-P31	10 meters
R5SS-P25	R11UL-P30	10 meters
R5SS-P26	R11LL-P31	10 meters
R5SS-P27	R12UL-P30	10 meters
R5SS-P28	R12LL-P31	10 meters
R5SS-P29	R13UL-P30	10 meters
R5SS-P30	R13LL-P31	10 meters
R5SS-P31	R14UL-P30	15 meters
R5SS-P32	R14LL-P31	15 meters
R6SS-P5	R1UL-P32	10 meters
R6SS-P6	R2UL-P33	10 meters
R6SS-P7	R5UL-P7	5 meters
R6SS-P8	R6LL-P32	3 meters

Table 22-117 (Cont.) Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R6SS-P9	R7UL-P33	5 meters
R6SS-P10	R4UL-P30	5 meters
R6SS-P11	R3UL-P7	10 meters
R6SS-P12	R8LL-P33	5 meters
R6SS-P13	R1LL-P32	10 meters
R6SS-P14	R2LL-P7	10 meters
R6SS-P15	R5LL-P4	5 meters
R6SS-P16	R6UL-P30	3 meters
R6SS-P17	R4LL-P6	5 meters
R6SS-P18	R7LL-P33	5 meters
R6SS-P19	R8UL-P33	5 meters
R6SS-P20	R3LL-P30	10 meters
R6SS-P21	R9UL-P31	10 meters
R6SS-P22	R9LL-P32	10 meters
R6SS-P23	R10UL-P31	10 meters
R6SS-P24	R10LL-P32	10 meters
R6SS-P25	R11UL-P31	10 meters
R6SS-P26	R11LL-P32	10 meters
R6SS-P27	R12UL-P31	10 meters
R6SS-P28	R12LL-P32	10 meters
R6SS-P29	R13UL-P31	10 meters
R6SS-P30	R13LL-P32	10 meters
R6SS-P31	R14UL-P31	10 meters
R6SS-P32	R14LL-P32	10 meters
R7SS-P5	R1UL-P33	10 meters
R7SS-P6	R2UL-P32	10 meters
R7SS-P7	R5UL-P32	5 meters
R7SS-P8	R6LL-P6	5 meters
R7SS-P9	R7UL-P30	3 meters
R7SS-P10	R4UL-P31	10 meters
R7SS-P11	R3UL-P31	10 meters
R7SS-P12	R8LL-P30	5 meters
R7SS-P13	R1LL-P6	10 meters
R7SS-P14	R2LL-P32	10 meters
R7SS-P15	R5LL-P5	5 meters
R7SS-P16	R6UL-P5	5 meters
R7SS-P17	R4LL-P32	10 meters
R7SS-P18	R7LL-P30	3 meters
R7SS-P19	R8UL-P30	5 meters
R7SS-P20	R3LL-P7	10 meters
R7SS-P21	R9UL-P32	5 meters
R7SS-P22	R9LL-P33	5 meters

Table 22-117 (Cont.) Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R7SS-P23	R10UL-P32	10 meters
R7SS-P24	R10LL-P33	10 meters
R7SS-P25	R11UL-P32	10 meters
R7SS-P26	R11LL-P33	10 meters
R7SS-P27	R12UL-P32	10 meters
R7SS-P28	R12LL-P33	10 meters
R7SS-P29	R13UL-P32	10 meters
R7SS-P30	R13LL-P33	10 meters
R7SS-P31	R14UL-P32	10 meters
R7SS-P32	R14LL-P33	10 meters
R8SS-P5	R1UL-P7	10 meters
R8SS-P6	R2UL-P30	10 meters
R8SS-P7	R5UL-P31	10 meters
R8SS-P8	R6LL-P7	5 meters
R8SS-P9	R7UL-P32	5 meters
R8SS-P10	R4UL-P4	10 meters
R8SS-P11	R3UL-P33	10 meters
R8SS-P12	R7LL-P32	5 meters
R8SS-P13	R1LL-P33	10 meters
R8SS-P14	R2LL-P30	10 meters
R8SS-P15	R5LL-P32	10 meters
R8SS-P16	R6UL-P32	5 meters
R8SS-P17	R4LL-P33	10 meters
R8SS-P18	R8LL-P32	3 meters
R8SS-P19	R8UL-P32	3 meters
R8SS-P20	R3LL-P6	10 meters
R8SS-P21	R9UL-P33	5 meters
R8SS-P22	R9LL-P34	5 meters
R8SS-P23	R10UL-P33	5 meters
R8SS-P24	R10LL-P34	5 meters
R8SS-P25	R11UL-P33	10 meters
R8SS-P26	R11LL-P34	10 meters
R8SS-P27	R12UL-P33	10 meters
R8SS-P28	R12LL-P34	10 meters
R8SS-P29	R13UL-P33	10 meters
R8SS-P30	R13LL-P34	10 meters
R8SS-P31	R14UL-P33	10 meters
R8SS-P32	R14LL-P34	10 meters
R9SS-P5	R1UL-P36	10 meters
R9SS-P6	R2UL-P36	10 meters
R9SS-P7	R5UL-P35	10 meters
R9SS-P8	R6LL-P1	10 meters

Table 22-117 (Cont.) Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R9SS-P9	R7UL-P2	5 meters
R9SS-P10	R4UL-P34	10 meters
R9SS-P11	R3UL-P2	10 meters
R9SS-P12	R7LL-P1	5 meters
R9SS-P13	R1LL-P36	10 meters
R9SS-P14	R2LL-P3	10 meters
R9SS-P15	R5LL-P1	10 meters
R9SS-P16	R6UL-P3	10 meters
R9SS-P17	R4LL-P1	10 meters
R9SS-P18	R8LL-P2	5 meters
R9SS-P19	R8UL-P1	5 meters
R9SS-P20	R3LL-P3	10 meters
R9SS-P21	R9UL-P35	3 meters
R9SS-P22	R9LL-P35	3 meters
R9SS-P23	R10UL-P35	5 meters
R9SS-P24	R10LL-P35	5 meters
R9SS-P25	R11UL-P35	5 meters
R9SS-P26	R11LL-P35	5 meters
R9SS-P27	R12UL-P35	10 meters
R9SS-P28	R12LL-P35	10 meters
R9SS-P29	R13UL-P35	10 meters
R9SS-P30	R13LL-P35	10 meters
R9SS-P31	R14UL-P35	10 meters
R9SS-P32	R14LL-P35	10 meters
R10SS-P5	R1UL-P35	15 meters
R10SS-P6	R2UL-P34	10 meters
R10SS-P7	R5UL-P1	10 meters
R10SS-P8	R6LL-P2	10 meters
R10SS-P9	R7UL-P36	10 meters
R10SS-P10	R4UL-P3	10 meters
R10SS-P11	R3UL-P3	10 meters
R10SS-P12	R7LL-P2	10 meters
R10SS-P13	R1LL-P35	15 meters
R10SS-P14	R2LL-P36	10 meters
R10SS-P15	R5LL-P35	10 meters
R10SS-P16	R6UL-P35	10 meters
R10SS-P17	R4LL-P2	10 meters
R10SS-P18	R8LL-P36	5 meters
R10SS-P19	R8UL-P2	5 meters
R10SS-P20	R3LL-P2	10 meters
R10SS-P21	R9UL-P4	5 meters
R10SS-P22	R9LL-P6	5 meters

Table 22-117 (Cont.) Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R10SS-P23	R10UL-P4	3 meters
R10SS-P24	R10LL-P6	3 meters
R10SS-P25	R11UL-P4	5 meters
R10SS-P26	R11LL-P6	5 meters
R10SS-P27	R12UL-P4	5 meters
R10SS-P28	R12LL-P6	5 meters
R10SS-P29	R13UL-P4	10 meters
R10SS-P30	R13LL-P6	10 meters
R10SS-P31	R14UL-P4	10 meters
R10SS-P32	R14LL-P6	10 meters
R11SS-P5	R1UL-P2	15 meters
R11SS-P6	R2UL-P35	15 meters
R11SS-P7	R5UL-P3	10 meters
R11SS-P8	R6LL-P35	10 meters
R11SS-P9	R7UL-P35	10 meters
R11SS-P10	R4UL-P2	10 meters
R11SS-P11	R3UL-P36	10 meters
R11SS-P12	R7LL-P35	10 meters
R11SS-P13	R1LL-P3	15 meters
R11SS-P14	R2LL-P34	15 meters
R11SS-P15	R5LL-P3	10 meters
R11SS-P16	R6UL-P2	10 meters
R11SS-P17	R4LL-P3	10 meters
R11SS-P18	R8LL-P35	10 meters
R11SS-P19	R8UL-P35	10 meters
R11SS-P20	R3LL-P34	10 meters
R11SS-P21	R9UL-P6	5 meters
R11SS-P22	R9LL-P2	5 meters
R11SS-P23	R10UL-P2	5 meters
R11SS-P24	R10LL-P4	5 meters
R11SS-P25	R11UL-P2	3 meters
R11SS-P26	R11LL-P4	3 meters
R11SS-P27	R12UL-P2	5 meters
R11SS-P28	R12LL-P4	5 meters
R11SS-P29	R13UL-P2	5 meters
R11SS-P30	R13LL-P4	5 meters
R11SS-P31	R14UL-P2	10 meters
R11SS-P32	R14LL-P4	10 meters
R12SS-P5	R1UL-P1	15 meters
R12SS-P6	R2UL-P1	15 meters
R12SS-P7	R5UL-P2	10 meters
R12SS-P8	R6LL-P36	10 meters

Table 22-117 (Cont.) Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R12SS-P9	R7UL-P34	10 meters
R12SS-P10	R4UL-P35	10 meters
R12SS-P11	R3UL-P35	15 meters
R12SS-P12	R7LL-P34	10 meters
R12SS-P13	R1LL-P34	15 meters
R12SS-P14	R2LL-P1	15 meters
R12SS-P15	R5LL-P2	10 meters
R12SS-P16	R6UL-P34	10 meters
R12SS-P17	R4LL-P34	10 meters
R12SS-P18	R8LL-P34	10 meters
R12SS-P19	R8UL-P34	10 meters
R12SS-P20	R3LL-P36	15 meters
R12SS-P21	R9UL-P36	10 meters
R12SS-P22	R9LL-P30	10 meters
R12SS-P23	R10UL-P34	5 meters
R12SS-P24	R10LL-P2	5 meters
R12SS-P25	R11UL-P34	5 meters
R12SS-P26	R11LL-P36	5 meters
R12SS-P27	R12UL-P34	3 meters
R12SS-P28	R12LL-P36	3 meters
R12SS-P29	R13UL-P34	5 meters
R12SS-P30	R13LL-P36	5 meters
R12SS-P31	R14UL-P34	5 meters
R12SS-P32	R14LL-P36	5 meters
R13SS-P5	R1UL-P34	15 meters
R13SS-P6	R2UL-P3	15 meters
R13SS-P7	R5UL-P34	10 meters
R13SS-P8	R6LL-P3	10 meters
R13SS-P9	R7UL-P3	10 meters
R13SS-P10	R4UL-P1	15 meters
R13SS-P11	R3UL-P1	15 meters
R13SS-P12	R7LL-P36	10 meters
R13SS-P13	R1LL-P2	15 meters
R13SS-P14	R2LL-P35	15 meters
R13SS-P15	R5LL-P34	10 meters
R13SS-P16	R6UL-P36	10 meters
R13SS-P17	R4LL-P35	15 meters
R13SS-P18	R8LL-P1	10 meters
R13SS-P19	R8UL-P36	10 meters
R13SS-P20	R3LL-P1	15 meters
R13SS-P21	R9UL-P34	10 meters
R13SS-P22	R9LL-P4	10 meters

Table 22-117 (Cont.) Fourteen-Rack System Connections Sorted By The Spine Switch Port Location

Spine Switch Port	Leaf Switch Port	Cable Length
R13SS-P23	R10UL-P6	10 meters
R13SS-P24	R10LL-P36	10 meters
R13SS-P25	R11UL-P6	5 meters
R13SS-P26	R11LL-P2	5 meters
R13SS-P27	R12UL-P36	5 meters
R13SS-P28	R12LL-P30	5 meters
R13SS-P29	R13UL-P36	3 meters
R13SS-P30	R13LL-P30	3 meters
R13SS-P31	R14UL-P36	5 meters
R13SS-P32	R14LL-P30	5 meters
R14SS-P5	R1UL-P3	15 meters
R14SS-P6	R2UL-P2	15 meters
R14SS-P7	R5UL-P36	15 meters
R14SS-P8	R6LL-P34	10 meters
R14SS-P9	R7UL-P1	10 meters
R14SS-P10	R4UL-P36	15 meters
R14SS-P11	R3UL-P34	15 meters
R14SS-P12	R7LL-P3	10 meters
R14SS-P13	R1LL-P1	15 meters
R14SS-P14	R2LL-P2	15 meters
R14SS-P15	R5LL-P36	15 meters
R14SS-P16	R6UL-P1	10 meters
R14SS-P17	R4LL-P36	15 meters
R14SS-P18	R8LL-P3	10 meters
R14SS-P19	R8UL-P3	10 meters
R14SS-P20	R3LL-P35	15 meters
R14SS-P21	R9UL-P2	10 meters
R14SS-P22	R9LL-P36	10 meters
R14SS-P23	R10UL-P36	10 meters
R14SS-P24	R10LL-P30	10 meters
R14SS-P25	R11UL-P36	10 meters
R14SS-P26	R11LL-P30	10 meters
R14SS-P27	R12UL-P6	5 meters
R14SS-P28	R12LL-P2	5 meters
R14SS-P29	R13UL-P6	5 meters
R14SS-P30	R13LL-P2	5 meters
R14SS-P31	R14UL-P6	3 meters
R14SS-P32	R14LL-P2	3 meters

Multi-Rack Cabling Tables for Recovery Appliance X8M Models

This section contains multi-rack cabling tables for Recovery Appliance X8M models, which use RoCE Network Fabric.

Understanding Multi-Rack Cabling for X8M Racks

Up to eight racks can be cabled together without external RDMA Network Fabric switches.

When cabling together multiple racks, each rack must contain three RDMA Network Fabric switches.

Two of the switches, known as leaf switches, connect to all of the database servers and storage servers in the rack. The leaf switches are located as follows:

- At rack unit 20 (U20) and 22 (U22) in Recovery Appliance X8M.

The third switch, known as the spine switch, interconnects the leaf switches in a multi-rack configuration. The spine switch is located at rack unit 1 (U1). If any rack does not include a spine switch, then order the expansion switch kit.


The procedures in this section assume the racks are adjacent to each other, standard routing in raised floor is used, and there are no obstacles in the raised floor. If these assumptions are not correct for your environment, then longer cables may be required for the connections.




Note:

By default, Oracle Exadata Database Machine X8M racks do not include spare cables or a third RoCE Network Fabric switch. To extend these racks, you must order the required cables and RoCE Network Fabric switch.

The following diagram shows the default RDMA Network Fabric architecture for a single-rack system. Each rack has two leaf switches, with eight connections between the leaf switches. The database servers and storage servers are each connected to both leaf switches. Each server contains a dual-port RDMA Network Fabric card, with port 1 connected to the lower leaf switch and port 2 connected to the upper leaf switch.



When connecting up to eight racks together, remove the eight existing inter-switch connections between each leaf switch on each rack. From each leaf switch, distribute eight connections over the spine switches in all racks. In multi-rack environments, the leaf switches inside a rack are no longer directly interconnected, as shown in the following graphic:



As shown in the preceding graphic, each leaf switch in rack 1 has the following connections:


- Four connections to its internal spine switch
- Four connections to the spine switch in rack 2

The spine switch in rack 1 has the following connections:

- Four connections to each leaf switch in rack 1
- Four connections to each leaf switch in rack 2

As the number of racks increases from two to eight, the pattern continues as shown in the following graphic:

Figure 23-1 Connections Between Spine Switches and Leaf Switches for up to 8 Racks



As shown in the preceding graphic, each leaf switch has eight inter-switch connections distributed over all spine switches. Each spine switch has 16 inter-switch connections distributed over all leaf switches. The leaf switches are not directly interconnected with other leaf switches, and the spine switches are not directly interconnected with the other spine switches.

Preparing for Multi-Rack Cabling with X8M Racks

Racks can be added together to increase system capacity and performance. When cabling racks together, note the following:

- The cable lengths shown in this document assume the racks are adjacent to each other. If the racks are not adjacent, or there are obstacles in the raised floor, or if you use overhead cabling, then longer cables may be required. For optical cables, the maximum supported cable length is 100 meters. For copper cables, the maximum supported cable length is 5 meters.
- Oracle recommends that the names for the servers include the rack unit number. This helps identify the server during diagnostics.
- When completing Oracle Exadata Deployment Assistant (OEDA) for the additional rack, you are prompted for SCAN addresses. However, these SCAN addresses are not used because the SCAN address from the original rack are used. Manually remove the new SCAN addresses from the generated installation files.
- The software owner account names and group names, as well as their identifiers, must match the names and identifiers of the original rack.
- If the additional grid disks are used with existing disk groups, then ensure the grid disk sizes for the new rack are the same as the original rack.

Note:

Recovery Appliance Eighth Racks, Quarter Racks, and Elastic Configurations are connected to other racks in the same fashion as larger racks are connected to each other. In other words, **a spine switch must exist in each rack in order to interconnect with other racks.**

Perform the following tasks before cabling racks together:

1. Determine the number of racks that will be cabled together.
2. Order the parts needed to connect the racks.

To extend Recovery Appliance racks with RoCE Network Fabric, for each rack being added you must order extra cables, transceivers for longer cables, and a RoCE Network Fabric spine switch, if one is not already available.

When connecting four or more racks, or if you need longer cables for your environment, you must purchase additional 10 meter fiber cables with two QSFP28 SR transceivers to connect each end. The QSFP28 SR transceivers are needed for fiber cables over 5 meters in length.

The following table lists the number of racks, and cables needed to interconnect up to eight racks:

Number of Racks	3 Meter Cables Needed	5 Meter Cables Needed	10 Meter Cables Needed
2	16	16	0
3	16	32	0
4	16	40	8
5	16	46	18
6	17	49	30
7	19	51	42
8	16	52	60

3. Gather the following tools:
 - Pen
 - Note pad
 - Flashlight
 - Labels for cables and rack
 - Cable ties (Oracle recommends hook and loop straps)
4. Determine a naming method for the rack prefixes. For example, if the original rack has the prefix `dbm01`, then use the prefix `dbm02` for the second rack, the prefix `dbm03` for the third rack, and so on.
5. Verify the racks have unique host names and IP addresses. All servers interconnected in the racks must have unique names and IP addresses.

Server names and IP addresses conventions may differ in the following cases:

- **Initial installation of all Recovery Appliance Racks:** System address assignments and host names should be complete.
- **New Recovery Appliance Racks are added to an existing cluster:** The new rack configuration should require unique host names and IP addresses for the new Recovery Appliance. The IP addresses on the same subnet cannot conflict with the existing systems.
- **Two existing Recovery Appliance Racks are clustered together:** You can assign host names and IP addresses only if Recovery Appliance racks are already assigned unique host names and IP addresses, or the entire cluster must be reconfigured. The machines must be on the same subnet and not have conflicting IP addresses.

6. Ensure the IP addresses for the new servers are in the same subnet, and do not overlap with the currently-installed servers.
7. Check that the firmware on the original switches is at the same level as the new switches by using the `sh ver` command.

It is highly recommended, though not mandatory, to use the same firmware version on all of the switches. If the firmware is not at the same level, you can apply a firmware patch to bring the switches up to the same firmware level.

Cabling Multiple Recovery Appliance X8M Racks

A spine switch must exist in the Recovery Appliance rack, regardless of size, to multi-rack with other racks.



Note:

To extend Recovery Appliance racks with RoCE Network Fabric, you must order cables, transceivers if needed, and a RoCE Network Fabric switch, if they are not already available.

In the following steps, the number in parentheses also indicates the number of cables required.

1. Split each leaf switch uplink bundle by the number of spine switches (or racks) in such a way that the total count adds up to 8. This split is represented in parentheses for each example.

Example 1: For two racks, take 8 uplinks from each leaf switch and split evenly by 2. Four uplinks from each leaf switch go to `rack1-spine` switch, and four uplinks from each leaf switch go to `rack2-spine` switch (4 + 4 for each leaf switch).

Example 2: For three racks, take 8 uplinks from each leaf switch and split evenly by 3. Three uplinks go to `rack1-spine` switch, three uplinks go to `rack2-spine` switch, two uplinks go to `rack3-spine` switch (3 + 3 + 2 for each leaf switch).
2. Starting from the first available port on a different spine switch, round-robin the above split scheme for each leaf switch and spine switch.

For example, for three racks:

- `rack1-leaf1` switch starts from `rack1-spine` switch for first split, `rack2-spine` switch for second split, `rack3-spine` switch for third split
 - `rack2-leaf1` switch starts from `rack2-spine` switch for first split, `rack3-spine` switch for second split, `rack1-spine` for third split
 - `rack3-leaf1` switch starts from `rack3-spine` switch for first split, `rack1-spine` switch for second split, `rack2-spine` switch for third split
 - and so on...
3. After walking through all leaf switch uplinks in each case, you will have used all spine switch ports between port 5 and 20 inclusive.

The remaining topics in this section provide detailed cabling information for cabling up to 8 racks together.

Two-Rack Cabling for X8M Racks

This section provides the cabling tables to connect two X8M racks together, both of which use RoCE Network Fabric.




Note:

The following conventions were used in the cabling notation for connecting multiple racks together.

- The spine switch (also referred to as SS) is in U1 for all racks.
- The leaf switches are referred to as Lower Leaf (LL) and Upper Leaf (UL).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in Recovery Appliance X8M or Storage Expansion Rack X8M
 - At rack unit 21 (U21) and rack unit 23 (U23) in Recovery Appliance X8M-8
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X8M-8 systems with three database servers and 11 storage servers only, you must adjust the following multi-rack cabling information. On such systems only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, in the following tables, any connection to port number 30 on any leaf switch must instead connect to port number 34 on the same leaf switch. For example, R1-UL-P30 must change to R1-UL-P34, R1-LL-P30 must change to R1-LL-P34, and so on. These changes only apply to X8M-8 systems with three database servers and 11 storage servers.

The following illustration shows the cable connections for the two spine switches (R1 SS and R2 SS) when cabling two racks together:



The following table describes the cable connections for the first spine switch (R1 SS) when cabling two racks together:

Table 23-1 Leaf Switch Connections for the First Rack in a Two-Rack System

Leaf Switch	Connection	Cable Length
R1 UL within Rack 1	R1-UL-P5 to R1-SS-P5 R1-UL-P7 to R1-SS-P7 R1-UL-P4 to R1-SS-P9 R1-UL-P6 to R1-SS-P11	3 meters
R1 UL to Rack 2	R1-UL-P31 to R2-SS-P5 R1-UL-P33 to R2-SS-P7 R1-UL-P30 to R2-SS-P9 R1-UL-P32 to R2-SS-P11	5 meters
R1 LL within Rack 1	R1-LL-P5 to R1-SS-P13 R1-LL-P7 to R1-SS-P15 R1-LL-P4 to R1-SS-P17 R1-LL-P6 to R1-SS-P19	3 meters
R1 LL to Rack 2	R1-LL-P31 to R2-SS-P13 R1-LL-P33 to R2-SS-P15 R1-LL-P30 to R2-SS-P17 R1-LL-P32 to R2-SS-P19	5 meters

The following table describes the cable connections for the second spine switch (R2 SS) when cabling two racks together:

Table 23-2 Leaf Switch Connections for the Second Rack in a Two-Rack System

Leaf Switch	Connection	Cable Length
R2 UL to Rack 1	R2-UL-P5 to R1-SS-P6 R2-UL-P7 to R1-SS-P8 R2-UL-P4 to R1-SS-P10 R2-UL-P6 to R1-SS-P12	5 meters
R2 UL within Rack 2	R2-UL-P31 to R2-SS-P6 R2-UL-P33 to R2-SS-P8 R2-UL-P30 to R2-SS-P10 R2-UL-P32 to R2-SS-P12	3 meters
R2 LL to Rack 1	R2-LL-P5 to R1-SS-P14 R2-LL-P7 to R1-SS-P16 R2-LL-P4 to R1-SS-P18 R2-LL-P6 to R1-SS-P20	5 meters
R2 LL within Rack 2	R2-LL-P31 to R2-SS-P14 R2-LL-P33 to R2-SS-P16 R2-LL-P30 to R2-SS-P18 R2-LL-P32 to R2-SS-P20	3 meters

Three-Rack Cabling for X8M Racks


This section provides the cabling tables to connect three X8M racks together using RoCE Network Fabric.

**Note:**

The following conventions were used in the cabling notation for connecting multiple racks together.

- The spine switch (also referred to as SS) is in U1 for all racks.
- The leaf switches are referred to as Lower Leaf (LL) and Upper Leaf (UL).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in Recovery Appliance X8M or Storage Expansion Rack X8M
 - At rack unit 21 (U21) and rack unit 23 (U23) in Recovery Appliance X8M-8
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X8M-8 systems with three database servers and 11 storage servers only, you must adjust the following multi-rack cabling information. On such systems only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, in the following tables, any connection to port number 30 on any leaf switch must instead connect to port number 34 on the same leaf switch. For example, R1-UL-P30 must change to R1-UL-P34, R1-LL-P30 must change to R1-LL-P34, and so on. These changes only apply to X8M-8 systems with three database servers and 11 storage servers.

The following illustration shows the cable connections for the three spine switches (Rack1-spine, Rack2-spine, and Rack3-spine) when cabling three racks together:



The following table describes the cable connections for the first spine switch (R1-SS) when cabling three racks together:

Table 23-3 Leaf Switch Connections for the First Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R1 UL within Rack 1	R1-UL-P5 to R1-SS-P5 R1-UL-P7 to R1-SS-P7 R1-UL-P4 to R1-SS-P9	3 meters
R1 UL to Rack 2	R1-UL-P31 to R2-SS-P5 R1-UL-P33 to R2-SS-P7 R1-UL-P30 to R2-SS-P9	5 meters
R1 UL to Rack 3	R1-UL-P6 to R3-SS-P5 R1-UL-P32 to R3-SS-P7	5 meters
R1 LL within Rack 1	R1-LL-P5 to R1-SS-P13 R1-LL-P7 to R1-SS-P15	3 meters
R1 LL to Rack 2	R1-LL-P31 to R2-SS-P13 R1-LL-P33 to R2-SS-P15 R1-LL-P30 to R2-SS-P17	5 meters

Table 23-3 (Cont.) Leaf Switch Connections for the First Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R1 LL to Rack 3	R1-LL-P4 to R3-SS-P13 R1-LL-P6 to R3-SS-P15 R1-LL-P32 to R3-SS-P17	5 meters

The following table describes the cable connections for the second spine switch (R2-SS) when cabling three racks together:

Table 23-4 Leaf Switch Connections for the Second Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R2 UL to Rack 1	R2-UL-P5 to R1-SS-P6 R2-UL-P7 to R1-SS-P8 R2-UL-P4 to R1-SS-P10	5 meters
R2 UL within Rack 2	R2-UL-P31 to R2-SS-P6 R2-UL-P33 to R2-SS-P8	3 meters
R2 UL to Rack 3	R2-UL-P6 to R3-SS-P6 R2-UL-P30 to R3-SS-P8 R2-UL-P32 to R3-SS-P10	5 meters
R2 LL to Rack 1	R2-LL-P5 to R1-SS-P14 R2-LL-P7 to R1-SS-P16 R2-LL-P4 to R1-SS-P18	5 meters
R2 LL within Rack 2	R2-LL-P31 to R2-SS-P14 R2-LL-P33 to R2-SS-P16 R2-LL-P30 to R2-SS-P18	3 meters
R2 LL to Rack 3	R2-LL-P6 to R3-SS-P14 R2-LL-P32 to R3-SS-P16	5 meters

The following table describes the cable connections for the third spine switch (R3-SS) when cabling three racks together:

Table 23-5 Leaf Switch Connections for the Third Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R3 UL to Rack 1	R3-UL-P5 to R1-SS-P11 R3-UL-P7 to R1-SS-P12	5 meters
R3 UL to Rack 2	R3-UL-P4 to R2-SS-P10 R3-UL-P6 to R2-SS-P11 R3-UL-P31 to R2-SS-P12	5 meters
R3 UL within Rack 3	R3-UL-P33 to R3-SS-P9 R3-UL-P30 to R3-SS-P11 R3-UL-P32 to R3-SS-P12	3 meters

Table 23-5 (Cont.) Leaf Switch Connections for the Third Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R3 LL to Rack 1	R3-LL-P5 to R1-SS-P17 R3-LL-P7 to R1-SS-P19 R3-LL-P4 to R1-SS-P20	5 meters
R3 LL to Rack 2	R3-LL-P6 to R2-SS-P19 R3-LL-P31 to R2-SS-P20	5 meters
R3 LL within Rack 3	R3-LL-P33 to R3-SS-P18 R3-LL-P30 to R3-SS-P19 R3-LL-P32 to R3-SS-P20	3 meters

Four-Rack Cabling for X8M Racks


This section provides the cabling tables to connect four X8M racks together, all of which use RoCE Network Fabric.

Note:

The following conventions were used in the cabling notation for connecting multiple racks together.

- The spine switch (also referred to as SS) is in U1 for all racks.
- The leaf switches are referred to as Lower Leaf (LL) and Upper Leaf (UL).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in Recovery Appliance X8M or Storage Expansion Rack X8M
 - At rack unit 21 (U21) and rack unit 23 (U23) in Recovery Appliance X8M-8
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X8M-8 systems with three database servers and 11 storage servers only, you must adjust the following multi-rack cabling information. On such systems only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, in the following tables, any connection to port number 30 on any leaf switch must instead connect to port number 34 on the same leaf switch. For example, R1-UL-P30 must change to R1-UL-P34, R1-LL-P30 must change to R1-LL-P34, and so on. These changes only apply to X8M-8 systems with three database servers and 11 storage servers.

The following illustration shows the cable connections for the four spine switches (Rack1-spine, Rack2-spine, Rack3-spine, and Rack4-spine) when cabling two racks together:



The following table describes the cable connections for the first spine switch (R1-SS) when cabling four racks together:

Table 23-6 Leaf Switch Connections for the First Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R1 UL within Rack 1	R1-UL-P5 to R1-SS-P5 R1-UL-P7 to R1-SS-P7	3 meters
R1 UL to Rack 2	R1-UL-P31 to R2-SS-P5 R1-UL-P33 to R2-SS-P7	5 meters
R1 UL to Rack 3	R1-UL-P6 to R3-SS-P5 R1-UL-P32 to R3-SS-P7	5 meters

Table 23-6 (Cont.) Leaf Switch Connections for the First Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R1 UL to Rack 4	R1-UL-P4 to R4-SS-P5 R1-UL-P30 to R4-SS-P7	10 meters
R1 LL within Rack 1	R1-LL-P5 to R1-SS-P13 R1-LL-P7 to R1-SS-P15	3 meters
R1 LL to Rack 2	R1-LL-P31 to R2-SS-P13 R1-LL-P33 to R2-SS-P15	5 meters
R1 LL to Rack 3	R1-LL-P4 to R3-SS-P13 R1-LL-P6 to R3-SS-P15	5 meters
R1 LL to Rack 4	R1-LL-P30 to R4-SS-P13 R1-LL-P32 to R4-SS-P15	10 meters

The following table describes the cable connections for the second spine switch (R2-SS) when cabling four full racks together:

Table 23-7 Leaf Switch Connections for the Second Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R2 UL to Rack 1	R2-UL-P5 to R1-SS-P6 R2-UL-P7 to R1-SS-P8	5 meters
R2 UL within Rack 2	R2-UL-P31 to R2-SS-P6 R2-UL-P33 to R2-SS-P8	3 meters
R2 UL to Rack 3	R2-UL-P6 to R3-SS-P6 R2-UL-P30 to R3-SS-P8	5 meters
R2 UL to Rack 4	R2-UL-P4 to R4-SS-P6 R2-UL-P32 to R4-SS-P8	5 meters
R2 LL to Rack 1	R2-LL-P5 to R1-SS-P14 R2-LL-P7 to R1-SS-P16	5 meters
R2 LL within Rack 2	R2-LL-P31 to R2-SS-P14 R2-LL-P33 to R2-SS-P16	3 meters
R2 LL to Rack 3	R2-LL-P6 to R3-SS-P14 R2-LL-P32 to R3-SS-P16	5 meters
R2 LL to Rack 4	R2-LL-P4 to R4-SS-P14 R2-LL-P30 to R4-SS-P16	5 meters

The following table describes the cable connections for the third spine switch (R3-SS) when cabling four full racks together:

Table 23-8 Leaf Switch Connections for the Third Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R3 UL to Rack 1	R3-UL-P5 to R1-SS-P11 R3-UL-P7 to R1-SS-P12	5 meters

Table 23-8 (Cont.) Leaf Switch Connections for the Third Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R3 UL to Rack 2	R3-UL-P6 to R2-SS-P11 R3-UL-P31 to R2-SS-P12	5 meters
R3 UL within Rack 3	R3-UL-P30 to R3-SS-P11 R3-UL-P32 to R3-SS-P12	3 meters
R3 UL to Rack 4	R3-UL-P4 to R4-SS-P11 R3-UL-P33 to R4-SS-P12	5 meters
R3 LL to Rack 1	R3-LL-P7 to R1-SS-P19 R3-LL-P4 to R1-SS-P20	5 meters
R3 LL to Rack 2	R3-LL-P6 to R2-SS-P19 R3-LL-P31 to R2-SS-P20	5 meters
R3 LL within Rack 3	R3-LL-P30 to R3-SS-P19 R3-LL-P32 to R3-SS-P20	3 meters
R3 LL to Rack 4	R3-LL-P5 to R4-SS-P19 R3-LL-P33 to R4-SS-P20	5 meters

The following table describes the cable connections for the fourth spine switch (R4-SS) when cabling four full racks together:

Table 23-9 Leaf Switch Connections for the Fourth Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R4 UL to Rack 1	R4-UL-P5 to R1-SS-P9 R4-UL-P7 to R1-SS-P10	10 meters
R4 UL to Rack 2	R4-UL-P4 to R2-SS-P9 R4-UL-P6 to R2-SS-P10	5 meters
R4 UL to Rack 3	R4-UL-P31 to R3-SS-P9 R4-UL-P33 to R3-SS-P10	5 meters
R4 UL within Rack 4	R4-UL-P30 to R4-SS-P9 R4-UL-P32 to R4-SS-P10	3 meters
R4 LL to Rack 1	R4-LL-P5 to R1-SS-P17 R4-LL-P7 to R1-SS-P18	10 meters
R4 LL to Rack 2	R4-LL-P4 to R2-SS-P17 R4-LL-P6 to R2-SS-P18	5 meters
R4 LL to Rack 3	R4-LL-P31 to R3-SS-P17 R4-LL-P33 to R3-SS-P18	5 meters
R4 LL within Rack 4	R4-LL-P30 to R4-SS-P17 R4-LL-P32 to R4-SS-P18	3 meters

Five-Rack Cabling for X8M Racks

This section provides the cabling tables to connect five (5) X8M racks together, all of which use RoCE Network Fabric.

 Note:

The following conventions were used in the cabling notation for connecting multiple racks together.

- The spine switch (also referred to as SS) is in U1 for all racks.
- The leaf switches are referred to as Lower Leaf (LL) and Upper Leaf (UL).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in Recovery Appliance X8M or Storage Expansion Rack X8M
 - At rack unit 21 (U21) and rack unit 23 (U23) in Recovery Appliance X8M-8
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X8M-8 systems with three database servers and 11 storage servers only, you must adjust the following multi-rack cabling information. On such systems only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, in the following tables, any connection to port number 30 on any leaf switch must instead connect to port number 34 on the same leaf switch. For example, R1-UL-P30 must change to R1-UL-P34, R1-LL-P30 must change to R1-LL-P34, and so on. These changes only apply to X8M-8 systems with three database servers and 11 storage servers.

The following illustration shows the cable connections for the five spine switches when cabling five racks together:

Rack 1 Spine (R1SS)																	
1	3	5 R1 UL P5	7 R1 UL P7	9 R5 UL P7	11 R3 UL P5	13 R1 LL P5	15 R5 LL P5	17 R4 LL P5	19 R3 LL P7	21	23	25	27	29	31	33	35
2	4	6 R2 UL P5	8 R5 UL P5	10 R4 UL P7	12 R3 UL P7	14 R2 LL P5	16 R2 LL P7	18 R5 LL P7	20 R3 LL P4	22	24	26	28	30	32	34	36

Rack 2 Spine (R2SS)																	
1	3	5 R1 UL P31	7 R1 UL P33	9 R4 UL P4	11 R3 UL P6	13 R1 LL P31	15 R1 LL P33	17 R4 LL P4	19 R3 LL P6	21	23	25	27	29	31	33	35
2	4	6 R2 UL P31	8 R5 UL P4	10 R4 UL P6	12 R3 UL P31	14 R2 LL P31	16 R5 LL P4	18 R5 LL P6	20 R3 LL P31	22	24	26	28	30	32	34	36

Rack 3 Spine (R3SS)																	
1	3	5 R1 UL P6	7 R1 UL P32	9 R4 UL P31	11 R3 UL P30	13 R1 LL P4	15 R1 LL P6	17 R4 LL P31	19 R3 LL P30	21	23	25	27	29	31	33	35
2	4	6 R2 UL P6	8 R2 UL P30	10 R4 UL P33	12 R5 UL P6	14 R2 LL P6	16 R5 LL P31	18 R4 LL P33	20 R3 LL P32	22	24	26	28	30	32	34	36

Rack 4 Spine (R4SS)																	
1	3	5 R1 UL P4	7 R5 UL P31	9 R4 UL P30	11 R3 UL P4	13 R1 LL P30	15 R1 LL P32	17 R4 LL P30	19 R5 LL P33	21	23	25	27	29	31	33	35
2	4	6 R2 UL P4	8 R2 UL P32	10 R4 UL P32	12 R5 UL P33	14 R2 LL P4	16 R2 LL P30	18 R4 LL P32	20 R3 LL P33	22	24	26	28	30	32	34	36

Rack 5 Spine (R5SS)																	
1	3	5 R1 UL P30	7 R5 UL P30	9 R5 UL P32	11 R3 UL P32	13 R1 LL P7	15 R5 LL P30	17 R4 LL P7	19 R5 LL P32	21	23	25	27	29	31	33	35
2	4	6 R2 UL P7	8 R2 UL P33	10 R4 UL P5	12 R3 UL P33	14 R2 LL P33	16 R2 LL P32	18 R4 LL P6	20 R3 LL P5	22	24	26	28	30	32	34	36

The following table describes the cable connections for the first spine switch (R1-SS) when cabling five racks together:

Table 23-10 Leaf Switch Connections for the First Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R1 UL within Rack 1	R1-UL-P5 to R1-SS-P5 R1-UL-P7 to R1-SS-P7	3 meters
R1 UL to Rack 2	R1-UL-P31 to R2-SS-P5 R1-UL-P33 to R2-SS-P7	5 meters
R1 UL to Rack 3	R1-UL-P6 to R3-SS-P5 R1-UL-P32 to R3-SS-P7	5 meters
R1 UL to Rack 4	R1-UL-P4 to R4-SS-P5	10 meters
R1 UL to Rack 5	R1-UL-P30 to R5-SS-P5	10 meters
R1 LL within Rack 1	R1-LL-P5 to R1-SS-P13	3 meters
R1 LL to Rack 2	R1-LL-P31 to R2-SS-P13 R1-LL-P33 to R2-SS-P15	5 meters
R1 LL to Rack 3	R1-LL-P4 to R3-SS-P13 R1-LL-P6 to R3-SS-P15	5 meters
R1 LL to Rack 4	R1-LL-P30 to R4-SS-P13 R1-LL-P32 to R4-SS-P15	10 meters
R1 LL to Rack 5	R1-LL-P7 to R5-SS-P13	10 meters

The following table describes the cable connections for the second spine switch (R2-SS) when cabling five full racks together:

Table 23-11 Leaf Switch Connections for the Second Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R2 UL to Rack 1	R2-UL-P5 to R1-SS-P6	5 meters
R2 UL within Rack 2	R2-UL-P31 to R2-SS-P6	3 meters
R2 UL to Rack 3	R2-UL-P6 to R3-SS-P6 R2-UL-P30 to R3-SS-P8	5 meters
R2 UL to Rack 4	R2-UL-P4 to R4-SS-P6 R2-UL-P32 to R4-SS-P8	5 meters
R2 UL to Rack 5	R2-UL-P7 to R5-SS-P6 R2-UL-P33 to R5-SS-P8	10 meters
R2 LL to Rack 1	R2-LL-P5 to R1-SS-P14 R2-LL-P7 to R1-SS-P16	5 meters
R2 LL within Rack 2	R2-LL-P31 to R2-SS-P14	3 meters
R2 LL to Rack 3	R2-LL-P6 to R3-SS-P14	5 meters
R2 LL to Rack 4	R2-LL-P4 to R4-SS-P14 R2-LL-P30 to R4-SS-P16	5 meters
R2 LL to Rack 5	R2-LL-P33 to R5-SS-P14 R2-LL-P32 to R5-SS-P16	10 meters

The following table describes the cable connections for the third spine switch (R3-SS) when cabling five full racks together:

Table 23-12 Leaf Switch Connections for the Third Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R3 UL to Rack 1	R3-UL-P5 to R1-SS-P11 R3-UL-P7 to R1-SS-P12	5 meters
R3 UL to Rack 2	R3-UL-P6 to R2-SS-P11 R3-UL-P31 to R2-SS-P12	5 meters
R3 UL within Rack 3	R3-UL-P30 to R3-SS-P11	3 meters
R3 UL to Rack 4	R3-UL-P4 to R4-SS-P11	5 meters
R3 UL to Rack 5	R3-UL-P32 to R5-SS-P11 R3-UL-P33 to R5-SS-P12	5 meters
R3 LL to Rack 1	R3-LL-P7 to R1-SS-P19 R3-LL-P4 to R1-SS-P20	5 meters
R3 LL to Rack 2	R3-LL-P6 to R2-SS-P19 R3-LL-P31 to R2-SS-P20	5 meters
R3 LL within Rack 3	R3-LL-P30 to R3-SS-P19 R3-LL-P32 to R3-SS-P20	3 meters
R3 LL to Rack 4	R3-LL-P33 to R4-SS-P20	5 meters
R3 LL to Rack 5	R3-LL-P5 to R5-SS-P20	5 meters

The following table describes the cable connections for the fourth spine switch (R4-SS) when cabling five full racks together:

Table 23-13 Leaf Switch Connections for the Fourth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R4 UL to Rack 1	R4-UL-P7 to R1-SS-P10	10 meters
R4 UL to Rack 2	R4-UL-P4 to R2-SS-P9 R4-UL-P6 to R2-SS-P10	5 meters
R4 UL to Rack 3	R4-UL-P31 to R3-SS-P9 R4-UL-P33 to R3-SS-P10	5 meters
R4 UL within Rack 4	R4-UL-P30 to R4-SS-P9 R4-UL-P32 to R4-SS-P10	3 meters
R4 UL to Rack 5	R4-UL-P5 to R5-SS-P10	5 meters
R4 LL to Rack 1	R4-LL-P5 to R1-SS-P17	10 meters
R4 LL to Rack 2	R4-LL-P4 to R2-SS-P17	5 meters
R4 LL to Rack 3	R4-LL-P31 to R3-SS-P17 R4-LL-P33 to R3-SS-P18	5 meters
R4 LL within Rack 4	R4-LL-P30 to R4-SS-P17 R4-LL-P32 to R4-SS-P18	3 meters
R4 LL to Rack 5	R4-LL-P7 to R5-SS-P17 R4-LL-P6 to R5-SS-P18	5 meters

The following table describes the cable connections for the fifth spine switch (R5-SS) when cabling five full racks together:

Table 23-14 Leaf Switch Connections for the Fifth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R5 UL to Rack 1	R5-UL-P5 to R1-SS-P8 R5-UL-P7 to R1-SS-P9	10 meters
R5 UL to Rack 2	R5-UL-P4 to R2-SS-P8	10 meters
R5 UL to Rack 3	R5-UL-P6 to R3-SS-P12	5 meters
R5 UL to Rack 4	R5-UL-P31 to R4-SS-P7 R5-UL-P33 to R4-SS-P12	5 meters
R5 UL within Rack 5	R5-UL-P30 to R5-SS-P7 R5-UL-P32 to R5-SS-P9	3 meters
R5 LL to Rack 1	R5-LL-P5 to R1-SS-P15 R5-LL-P7 to R1-SS-P18	10 meters
R5 LL to Rack 2	R5-LL-P4 to R2-SS-P16 R5-LL-P6 to R2-SS-P18	10 meters
R5 LL to Rack 3	R5-LL-P31 to R3-SS-P16	5 meters
R5 LL to Rack 4	R5-LL-P33 to R4-SS-P19	5 meters
R5 LL within Rack 5	R5-LL-P30 to R5-SS-P15 R5-LL-P32 to R5-SS-P19	3 meters

Six-Rack Cabling for X8M Racks

This section provides the cabling tables to connect six (6) X8M racks together, all of which use RoCE Network Fabric.



Note:

The following conventions were used in the cabling notation for connecting multiple racks together.

- The spine switch (also referred to as SS) is in U1 for all racks.
- The leaf switches are referred to as Lower Leaf (LL) and Upper Leaf (UL).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in Recovery Appliance X8M or Storage Expansion Rack X8M
 - At rack unit 21 (U21) and rack unit 23 (U23) in Recovery Appliance X8M-8
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X8M-8 systems with three database servers and 11 storage servers only, you must adjust the following multi-rack cabling information. On such systems only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, in the following tables, any connection to port number 30 on any leaf switch must instead connect to port number 34 on the same leaf switch. For example, R1-UL-P30 must change to R1-UL-P34, R1-LL-P30 must change to R1-LL-P34, and so on. These changes only apply to X8M-8 systems with three database servers and 11 storage servers.

The following illustration shows the cable connections for the 6 spine switches when cabling six racks together:

Rack 1 Spine (R1SS)

1	3	5 R1 UL P5	7 R1 UL P7	9 R6 UL P5	11 R3 UL P5	13 R1 LL P5	15 R6 LL P7	17 R4 LL P5	19 R3 LL P7	21	23	25	27	29	31	33	35
2	4	6 R2 UL P5	8 R5 UL P5	10 R4 UL P7	12 R6 LL P5	14 R2 LL P5	16 R6 UL P7	18 R5 LL P7	20 R3 LL P4	22	24	26	28	30	32	34	36

Rack 2 Spine (R2SS)

1	3	5 R1 UL P31	7 R1 UL P33	9 R4 UL P4	11 R3 UL P6	13 R1 LL P31	15 R1 LL P33	17 R4 LL P4	19 R6 LL P6	21	23	25	27	29	31	33	35
2	4	6 R2 UL P31	8 R5 UL P4	10 R4 UL P6	12 R6 LL P4	14 R2 LL P31	16 R6 UL P4	18 R5 LL P6	20 R3 LL P31	22	24	26	28	30	32	34	36

Rack 3 Spine (R3SS)

1	3	5 R1 UL P6	7 R6 LL P31	9 R4 UL P31	11 R3 UL P30	13 R1 LL P4	15 R1 LL P6	17 R4 LL P31	19 R6 UL P6	21	23	25	27	29	31	33	35
2	4	6 R2 UL P6	8 R2 UL P30	10 R4 UL P33	12 R5 UL P6	14 R2 LL P6	16 R5 LL P31	18 R4 LL P33	20 R3 LL P32	22	24	26	28	30	32	34	36

Rack 4 Spine (R4SS)

1	3	5 R1 UL P4	7 R5 UL P31	9 R6 LL P33	11 R3 UL P4	13 R1 LL P30	15 R6 UL P31	17 R4 LL P30	19 R5 LL P33	21	23	25	27	29	31	33	35
2	4	6 R2 UL P4	8 R2 UL P32	10 R4 UL P32	12 R5 UL P33	14 R2 LL P4	16 R2 LL P30	18 R4 LL P32	20 R3 LL P33	22	24	26	28	30	32	34	36

Rack 5 Spine (R5SS)

1	3	5 R1 UL P30	7 R5 UL P30	9 R5 UL P32	11 R3 UL P32	13 R1 LL P7	15 R5 LL P30	17 R4 LL P7	19 R5 LL P32	21	23	25	27	29	31	33	35
2	4	6 R2 UL P7	8 R6 LL P30	10 R4 UL P5	12 R3 UL P33	14 R2 LL P33	16 R2 LL P32	18 R6 UL P33	20 R3 LL P5	22	24	26	28	30	32	34	36

Rack 6 Spine (R6SS)

1	3	5 R1 UL P32	7 R5 UL P7	9 R5 LL P5	11 R3 UL P7	13 R1 LL P32	15 R5 LL P4	17 R4 LL P6	19 R3 LL P6	21	23	25	27	29	31	33	35
2	4	6 R2 UL P33	8 R6 LL P32	10 R4 UL P30	12 R3 UL P31	14 R2 LL P7	16 R6 UL P30	18 R6 UL P32	20 R3 LL P30	22	24	26	28	30	32	34	36

The following table describes the cable connections for the first spine switch (R1-SS) when cabling six racks together:

Table 23-15 Leaf Switch Connections for the First Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R1 UL within Rack 1	R1-UL-P5 to R1-SS-P5 R1-UL-P7 to R1-SS-P7	3 meters
R1 UL to Rack 2	R1-UL-P31 to R2-SS-P5 R1-UL-P33 to R2-SS-P7	5 meters
R1 UL to Rack 3	R1-UL-P6 to R3-SS-P5	5 meters
R1 UL to Rack 4	R1-UL-P4 to R4-SS-P5	10 meters
R1 UL to Rack 5	R1-UL-P30 to R5-SS-P5	10 meters
R1 UL to Rack 6	R1-UL-P32 to R6-SS-P5	10 meters
R1 LL within Rack 1	R1-LL-P5 to R1-SS-P13	3 meters
R1 LL to Rack 2	R1-LL-P31 to R2-SS-P13 R1-LL-P33 to R2-SS-P15	5 meters
R1 LL to Rack 3	R1-LL-P4 to R3-SS-P13 R1-LL-P6 to R3-SS-P15	5 meters
R1 LL to Rack 4	R1-LL-P30 to R4-SS-P13	10 meters
R1 LL to Rack 5	R1-LL-P7 to R5-SS-P13	10 meters
R1 LL to Rack 6	R1-LL-P32 to R6-SS-P13	10 meters

The following table describes the cable connections for the second spine switch (R2-SS) when cabling six full racks together:

Table 23-16 Leaf Switch Connections for the Second Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R2 UL to Rack 1	R2-UL-P5 to R1-SS-P6	5 meters
R2 UL within Rack 2	R2-UL-P31 to R2-SS-P6	3 meters
R2 UL to Rack 3	R2-UL-P6 to R3-SS-P6 R2-UL-P30 to R3-SS-P8	5 meters
R2 UL to Rack 4	R2-UL-P4 to R4-SS-P6 R2-UL-P32 to R4-SS-P8	5 meters
R2 UL to Rack 5	R2-UL-P7 to R5-SS-P6	10 meters
R2 UL to Rack 6	R2-UL-P33 to R6-SS-P6	10 meters
R2 LL to Rack 1	R2-LL-P5 to R1-SS-P14	5 meters
R2 LL within Rack 2	R2-LL-P31 to R2-SS-P14	3 meters
R2 LL to Rack 3	R2-LL-P6 to R3-SS-P14	5 meters
R2 LL to Rack 4	R2-LL-P4 to R4-SS-P14 R2-LL-P30 to R4-SS-P16	5 meters
R2 LL to Rack 5	R2-LL-P33 to R5-SS-P14 R2-LL-P32 to R5-SS-P16	10 meters

Table 23-16 (Cont.) Leaf Switch Connections for the Second Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R2 LL to Rack 6	R2-LL-P7 to R6-SS-P14	10 meters

The following table describes the cable connections for the third spine switch (R3-SS) when cabling six full racks together:

Table 23-17 Leaf Switch Connections for the Third Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R3 UL to Rack 1	R3-UL-P5 to R1-SS-P11	5 meters
R3 UL to Rack 2	R3-UL-P6 to R2-SS-P11	5 meters
R3 UL within Rack 3	R3-UL-P30 to R3-SS-P11	3 meters
R3 UL to Rack 4	R3-UL-P4 to R4-SS-P11	5 meters
R3 UL to Rack 5	R3-UL-P32 to R5-SS-P11 R3-UL-P33 to R5-SS-P12	5 meters
R3 UL to Rack 6	R3-UL-P7 to R6-SS-P11 R3-UL-P31 to R6-SS-P12	10 meters
R3 LL to Rack 1	R3-LL-P7 to R1-SS-P19 R3-LL-P4 to R1-SS-P20	5 meters
R3 LL to Rack 2	R3-LL-P31 to R2-SS-P20	5 meters
R3 LL within Rack 3	R3-LL-P32 to R3-SS-P20	3 meters
R3 LL to Rack 4	R3-LL-P33 to R4-SS-P20	5 meters
R3 LL to Rack 5	R3-LL-P5 to R5-SS-P20	5 meters
R3 LL to Rack 6	R3-LL-P6 to R6-SS-P19 R3-LL-P30 to R6-SS-P20	10 meters

The following table describes the cable connections for the fourth spine switch (R4-SS) when cabling six full racks together:

Table 23-18 Leaf Switch Connections for the Fourth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R4 UL to Rack 1	R4-UL-P7 to R1-SS-P10	10 meters
R4 UL to Rack 2	R4-UL-P4 to R2-SS-P9 R4-UL-P6 to R2-SS-P10	5 meters
R4 UL to Rack 3	R4-UL-P31 to R3-SS-P9 R4-UL-P33 to R3-SS-P10	5 meters
R4 UL within Rack 4	R4-UL-P32 to R4-SS-P10	3 meters
R4 UL to Rack 5	R4-UL-P5 to R5-SS-P10	5 meters
R4 UL to Rack 6	R4-UL-P30 to R6-SS-P10	5 meters
R4 LL to Rack 1	R4-LL-P5 to R1-SS-P17	10 meters

Table 23-18 (Cont.) Leaf Switch Connections for the Fourth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R4 LL to Rack 2	R4-LL-P4 to R2-SS-P17	5 meters
R4 LL to Rack 3	R4-LL-P31 to R3-SS-P17 R4-LL-P33 to R3-SS-P18	5 meters
R4 LL within Rack 4	R4-LL-P30 to R4-SS-P17 R4-LL-P32 to R4-SS-P18	3 meters
R4 LL to Rack 5	R4-LL-P7 to R5-SS-P17	5 meters
R4 LL to Rack 6	R4-LL-P6 to R6-SS-P17	5 meters

The following table describes the cable connections for the fifth spine switch (R5-SS) when cabling six full racks together:

Table 23-19 Leaf Switch Connections for the Fifth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R5 UL to Rack 1	R5-UL-P5 to R1-SS-P8	10 meters
R5 UL to Rack 2	R5-UL-P4 to R2-SS-P8	10 meters
R5 UL to Rack 3	R5-UL-P6 to R3-SS-P12	5 meters
R5 UL to Rack 4	R5-UL-P31 to R4-SS-P7 R5-UL-P33 to R4-SS-P12	5 meters
R5 UL within Rack 5	R5-UL-P30 to R5-SS-P7 R5-UL-P32 to R5-SS-P9	3 meters
R5 UL to Rack 6	R5-UL-P7 to R6-SS-P7	5 meters
R5 LL to Rack 1	R5-LL-P7 to R1-SS-P18	10 meters
R5 LL to Rack 2	R5-LL-P6 to R2-SS-P18	10 meters
R5 LL to Rack 3	R5-LL-P31 to R3-SS-P16	5 meters
R5 LL to Rack 4	R5-LL-P33 to R4-SS-P19	5 meters
R5 LL within Rack 5	R5-LL-P30 to R5-SS-P15 R5-LL-P32 to R5-SS-P19	3 meters
R5 LL to Rack 6	R5-LL-P5 to R6-SS-P9 R5-LL-P4 to R6-SS-P15	5 meters

The following table describes the cable connections for the sixth spine switch (R6-SS) when cabling six full racks together:

Table 23-20 Leaf Switch Connections for the Sixth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R6 UL to Rack 1	R6-UL-P5 to R1-SS-P9 R6-UL-P7 to R1-SS-P16	10 meters
R6 UL to Rack 2	R6-UL-P4 to R2-SS-P16	10 meters
R6 UL to Rack 3	R6-UL-P6 to R3-SS-P19	10 meters

Table 23-20 (Cont.) Leaf Switch Connections for the Sixth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R6 UL to Rack 4	R6-UL-P31 to R4-SS-P15	5 meters
R6 UL to Rack 5	R6-UL-P33 to R5-SS-P18	5 meters
R6 UL within Rack 6	R6-UL-P30 to R6-SS-P16 R6-UL-P32 to R6-SS-P18	3 meters
R6 LL to Rack 1	R6-LL-P5 to R1-SS-P12 R6-LL-P7 to R1-SS-P15	10 meters
R6 LL to Rack 2	R6-LL-P4 to R2-SS-P12 R6-LL-P6 to R2-SS-P19	10 meters
R6 LL to Rack 3	R6-LL-P31 to R3-SS-P7	10 meters
R6 LL to Rack 4	R6-LL-P33 to R4-SS-P9	5 meters
R6 LL to Rack 5	R6-LL-P30 to R5-SS-P8	5 meters
R6 LL within Rack 6	R6-LL-P32 to R6-SS-P8	3 meters

Seven-Rack Cabling for X8M Racks

This section provides the cabling tables to connect seven (7) X8M racks together, all of which use RoCE Network Fabric.

**Note:**

The following conventions were used in the cabling notation for connecting multiple racks together.

- The spine switch (also referred to as SS) is in U1 for all racks.
- The leaf switches are referred to as Lower Leaf (LL) and Upper Leaf (UL).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in Recovery Appliance X8M or Storage Expansion Rack X8M
 - At rack unit 21 (U21) and rack unit 23 (U23) in Recovery Appliance X8M-8
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X8M-8 systems with three database servers and 11 storage servers only, you must adjust the following multi-rack cabling information. On such systems only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, in the following tables, any connection to port number 30 on any leaf switch must instead connect to port number 34 on the same leaf switch. For example, R1-UL-P30 must change to R1-UL-P34, R1-LL-P30 must change to R1-LL-P34, and so on. These changes only apply to X8M-8 systems with three database servers and 11 storage servers.


The following diagrams show the cable connections for the 7 spine switches when cabling seven racks together:

Rack 1 Spine (R1SS)																	
1	3	5 R1 UL P5	7 R1 UL P7	9 R7 UL P5	11 R3 UL P5	13 R1 LL P5	15 R6 LL P7	17 R4 LL P5	19 R7 LL P5	21	23	25	27	29	31	33	35
2	4	6 R2 UL P5	8 R5 UL P5	10 R4 UL P7	12 R6 LL P5	14 R2 LL P5	16 R6 UL P7	18 R5 LL P7	20 R3 LL P4	22	24	26	28	30	32	34	36

Rack 2 Spine (R2SS)																	
1	3	5 R1 UL P31	7 R7 UL P7	9 R4 UL P4	11 R3 UL P6	13 R1 LL P31	15 R1 LL P33	17 R4 LL P4	19 R7 LL P7	21	23	25	27	29	31	33	35
2	4	6 R2 UL P31	8 R5 UL P4	10 R4 UL P6	12 R6 LL P4	14 R2 LL P31	16 R6 UL P4	18 R5 LL P6	20 R3 LL P31	22	24	26	28	30	32	34	36

Rack 3 Spine (R3SS)																	
1	3	5 R1 UL P6	7 R6 LL P31	9 R7 UL P4	11 R3 UL P30	13 R1 LL P4	15 R7 LL P4	17 R4 LL P31	19 R6 UL P6	21	23	25	27	29	31	33	35
2	4	6 R2 UL P6	8 R2 UL P30	10 R4 UL P33	12 R5 UL P6	14 R2 LL P6	16 R5 LL P31	18 R4 LL P33	20 R3 LL P32	22	24	26	28	30	32	34	36

Rack 4 Spine (R4SS)																	
1	3	5 R1 UL P4	7 R5 UL P31	9 R6 LL P33	11 R3 UL P4	13 R1 LL P30	15 R6 UL P31	17 R4 LL P30	19 R5 LL P33	21	23	25	27	29	31	33	35
2	4	6 R2 UL P4	8 R7 UL P6	10 R4 UL P32	12 R5 UL P33	14 R2 LL P4	16 R2 LL P30	18 R7 LL P6	20 R3 LL P33	22	24	26	28	30	32	34	36



The following table describes the cable connections for the first spine switch (R1-SS) when cabling seven racks together:

Table 23-21 Leaf Switch Connections for the First Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R1 UL within Rack 1	R1-UL-P5 to R1-SS-P5 R1-UL-P7 to R1-SS-P7	3 meters
R1 UL to Rack 2	R1-UL-P31 to R2-SS-P5	5 meters
R1 UL to Rack 3	R1-UL-P6 to R3-SS-P5	5 meters
R1 UL to Rack 4	R1-UL-P4 to R4-SS-P5	10 meters
R1 UL to Rack 5	R1-UL-P30 to R5-SS-P5	10 meters
R1 UL to Rack 6	R1-UL-P32 to R6-SS-P5	10 meters
R1 UL to Rack 7	R1-UL-P33 to R7-SS-P5	10 meters
R1 LL within Rack 1	R1-LL-P5 to R1-SS-P13	3 meters
R1 LL to Rack 2	R1-LL-P31 to R2-SS-P13 R1-LL-P33 to R2-SS-P15	5 meters
R1 LL to Rack 3	R1-LL-P4 to R3-SS-P13	5 meters
R1 LL to Rack 4	R1-LL-P30 to R4-SS-P13	10 meters
R1 LL to Rack 5	R1-LL-P7 to R5-SS-P13	10 meters

Table 23-21 (Cont.) Leaf Switch Connections for the First Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R1 LL to Rack 6	R1-LL-P32 to R6-SS-P13	10 meters
R1 LL to Rack 7	R1-LL-P6 to R7-SS-P13	10 meters

The following table describes the cable connections for the second spine switch (R2-SS) when cabling seven full racks together:

Table 23-22 Leaf Switch Connections for the Second Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R2 UL to Rack 1	R2-UL-P5 to R1-SS-P6	5 meters
R2 UL within Rack 2	R2-UL-P31 to R2-SS-P6	3 meters
R2 UL to Rack 3	R2-UL-P6 to R3-SS-P6 R2-UL-P30 to R3-SS-P8	5 meters
R2 UL to Rack 4	R2-UL-P4 to R4-SS-P6	5 meters
R2 UL to Rack 5	R2-UL-P7 to R5-SS-P6	10 meters
R2 UL to Rack 6	R2-UL-P33 to R6-SS-P6	10 meters
R2 UL to Rack 7	R2-UL-P32 to R7-SS-P6	10 meters
R2 LL to Rack 1	R2-LL-P5 to R1-SS-P14	5 meters
R2 LL within Rack 2	R2-LL-P31 to R2-SS-P14	3 meters
R2 LL to Rack 3	R2-LL-P6 to R3-SS-P14	5 meters
R2 LL to Rack 4	R2-LL-P4 to R4-SS-P14 R2-LL-P30 to R4-SS-P16	5 meters
R2 LL to Rack 5	R2-LL-P33 to R5-SS-P14	10 meters
R2 LL to Rack 6	R2-LL-P7 to R6-SS-P14	10 meters
R2 LL to Rack 7	R2-LL-P32 to R7-SS-P14	10 meters

The following table describes the cable connections for the third spine switch (R3-SS) when cabling seven full racks together:

Table 23-23 Leaf Switch Connections for the Third Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R3 UL to Rack 1	R3-UL-P5 to R1-SS-P11	5 meters
R3 UL to Rack 2	R3-UL-P6 to R2-SS-P11	5 meters
R3 UL within Rack 3	R3-UL-P30 to R3-SS-P11	3 meters
R3 UL to Rack 4	R3-UL-P4 to R4-SS-P11	5 meters
R3 UL to Rack 5	R3-UL-P32 to R5-SS-P11 R3-UL-P33 to R5-SS-P12	5 meters
R3 UL to Rack 6	R3-UL-P7 to R6-SS-P11	10 meters
R3 UL to Rack 7	R3-UL-P31 to R7-SS-P11	10 meters

Table 23-23 (Cont.) Leaf Switch Connections for the Third Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R3 LL to Rack 1	R3-LL-P4 to R1-SS-P20	5 meters
R3 LL to Rack 2	R3-LL-P31 to R2-SS-P20	5 meters
R3 LL within Rack 3	R3-LL-P32 to R3-SS-P20	3 meters
R3 LL to Rack 4	R3-LL-P33 to R4-SS-P20	5 meters
R3 LL to Rack 5	R3-LL-P5 to R5-SS-P20	5 meters
R3 LL to Rack 6	R3-LL-P6 to R6-SS-P19 R3-LL-P30 to R6-SS-P20	10 meters
R3 LL to Rack 7	R3-LL-P7 to R7-SS-P20	10 meters

The following table describes the cable connections for the fourth spine switch (R4-SS) when cabling seven full racks together:

Table 23-24 Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 UL to Rack 1	R4-UL-P7 to R1-SS-P10	10 meters
R4 UL to Rack 2	R4-UL-P4 to R2-SS-P9 R4-UL-P6 to R2-SS-P10	5 meters
R4 UL to Rack 3	R4-UL-P33 to R3-SS-P10	5 meters
R4 UL within Rack 4	R4-UL-P32 to R4-SS-P10	3 meters
R4 UL to Rack 5	R4-UL-P5 to R5-SS-P10	5 meters
R4 UL to Rack 6	R4-UL-P30 to R6-SS-P10	5 meters
R4 UL to Rack 7	R4-UL-P31 to R7-SS-P10	10 meters
R4 LL to Rack 1	R4-LL-P5 to R1-SS-P17	10 meters
R4 LL to Rack 2	R4-LL-P4 to R2-SS-P17	5 meters
R4 LL to Rack 3	R4-LL-P31 to R3-SS-P17 R4-LL-P33 to R3-SS-P18	5 meters
R4 LL within Rack 4	R4-LL-P30 to R4-SS-P17	3 meters
R4 LL to Rack 5	R4-LL-P7 to R5-SS-P17	5 meters
R4 LL to Rack 6	R4-LL-P6 to R6-SS-P17	5 meters
R4 LL to Rack 7	R4-LL-P32 to R7-SS-P17	10 meters

The following table describes the cable connections for the fifth spine switch (R5-SS) when cabling seven full racks together:

Table 23-25 Leaf Switch Connections for the Fifth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R5 UL to Rack 1	R5-UL-P5 to R1-SS-P8	10 meters
R5 UL to Rack 2	R5-UL-P4 to R2-SS-P8	10 meters

Table 23-25 (Cont.) Leaf Switch Connections for the Fifth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R5 UL to Rack 3	R5-UL-P6 to R3-SS-P12	5 meters
R5 UL to Rack 4	R5-UL-P31 to R4-SS-P7 R5-UL-P33 to R4-SS-P12	5 meters
R5 UL within Rack 5	R5-UL-P30 to R5-SS-P7	3 meters
R5 UL to Rack 6	R5-UL-P7 to R6-SS-P7	5 meters
R5 UL to Rack 7	R5-UL-P32 to R7-SS-P7	5 meters
R5 LL to Rack 1	R5-LL-P7 to R1-SS-P18	10 meters
R5 LL to Rack 2	R5-LL-P6 to R2-SS-P18	10 meters
R5 LL to Rack 3	R5-LL-P31 to R3-SS-P16	5 meters
R5 LL to Rack 4	R5-LL-P33 to R4-SS-P19	5 meters
R5 LL within Rack 5	R5-LL-P30 to R5-SS-P15 R5-LL-P32 to R5-SS-P19	3 meters
R5 LL to Rack 6	R5-LL-P4 to R6-SS-P15	5 meters
R5 LL to Rack 7	R5-LL-P5 to R7-SS-P15	5 meters

The following table describes the cable connections for the sixth spine switch (R6-SS) when cabling seven full racks together:

Table 23-26 Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R6 UL to Rack 1	R6-UL-P7 to R1-SS-P16	10 meters
R6 UL to Rack 2	R6-UL-P4 to R2-SS-P16	10 meters
R6 UL to Rack 3	R6-UL-P6 to R3-SS-P19	10 meters
R6 UL to Rack 4	R6-UL-P31 to R4-SS-P15	5 meters
R6 UL to Rack 5	R6-UL-P33 to R5-SS-P18	5 meters
R6 UL within Rack 6	R6-UL-P30 to R6-SS-P16 R6-UL-P32 to R6-SS-P18	3 meters
R6 UL to Rack 7	R6-UL-P5 to R7-SS-P16	5 meters
R6 LL to Rack 1	R6-LL-P5 to R1-SS-P12 R6-LL-P7 to R1-SS-P15	10 meters
R6 LL to Rack 2	R6-LL-P4 to R2-SS-P12	10 meters
R6 LL to Rack 3	R6-LL-P31 to R3-SS-P7	10 meters
R6 LL to Rack 4	R6-LL-P33 to R4-SS-P9	5 meters
R6 LL to Rack 5	R6-LL-P30 to R5-SS-P8	5 meters
R6 LL within Rack 6	R6-LL-P32 to R6-SS-P8	3 meters
R6 LL to Rack 7	R6-LL-P6 to R7-SS-P8	5 meters

The following table describes the cable connections for the seventh spine switch (R7-SS) when cabling seven full racks together:

Table 23-27 Leaf Switch Connections for the Seventh Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R7 UL to Rack 1	R7-UL-P5 to R1-SS-P9	10 meters
R7 UL to Rack 2	R7-UL-P7 to R2-SS-P7	10 meters
R7 UL to Rack 3	R7-UL-P4 to R3-SS-P9	10 meters
R7 UL to Rack 4	R7-UL-P6 to R4-SS-P8	10 meters
R7 UL to Rack 5	R7-UL-P31 to R5-SS-P9	5 meters
R7 UL to Rack 6	R7-UL-P33 to R6-SS-P9	5 meters
R7 UL within Rack 7	R7-UL-P30 to R7-SS-P9 R7-UL-P32 to R7-SS-P12	3 meters
R7 LL to Rack 1	R7-LL-P5 to R1-SS-P19	10 meters
R7 LL to Rack 2	R7-LL-P7 to R2-SS-P19	10 meters
R7 LL to Rack 3	R7-LL-P4 to R3-SS-P15	10 meters
R7 LL to Rack 4	R7-LL-P6 to R4-SS-P18	10 meters
R7 LL to Rack 5	R7-LL-P31 to R5-SS-P16	5 meters
R7 LL to Rack 6	R7-LL-P33 to R6-SS-P12	5 meters
R7 LL within Rack 7	R7-LL-P30 to R7-SS-P18 R7-LL-P32 to R7-SS-P19	3 meters

Eight-Rack Cabling for X8M Racks

This section provides the cabling tables to connect eight (8) X8M racks together, all of which use RoCE Network Fabric.

**Note:**

The following conventions were used in the cabling notation for connecting multiple racks together.

- The spine switch (also referred to as SS) is in U1 for all racks.
- The leaf switches are referred to as Lower Leaf (LL) and Upper Leaf (UL).
- The leaf switches are located as follows:
 - At rack unit 20 (U20) and 22 (U22) in Recovery Appliance X8M or Storage Expansion Rack X8M
 - At rack unit 21 (U21) and rack unit 23 (U23) in Recovery Appliance X8M-8
- The cable lengths shown in the following lists assume that the racks are adjacent to each other, the cables are routed through a raised floor, and there are no obstacles in the routing between the racks. If the racks are not adjacent, or use overhead cabling trays, then they may require longer cable lengths. Cable lengths up to 100 meters are supported.
- Only optical cables (with additional transceivers) are supported for lengths greater than 5 meters.
- For X8M-8 systems with three database servers and 11 storage servers only, you must adjust the following multi-rack cabling information. On such systems only, port 30 on the leaf switches is connected to a database server and is not used as an inter-switch link. Consequently, in the following tables, any connection to port number 30 on any leaf switch must instead connect to port number 34 on the same leaf switch. For example, R1-UL-P30 must change to R1-UL-P34, R1-LL-P30 must change to R1-LL-P34, and so on. These changes only apply to X8M-8 systems with three database servers and 11 storage servers.


The following diagrams show the cable connections for the 8 spine switches when cabling eight racks together:

Rack 1 Spine (R1SS)																	
1	3	5 R1 UL P5	7 R8 LL P5	9 R7 UL P5	11 R3 UL P5	13 R1 LL P5	15 R8 UL P5	17 R4 LL P5	19 R7 LL P5	21	23	25	27	29	31	33	35
2	4	6 R2 UL P5	8 R5 UL P5	10 R4 UL P7	12 R6 LL P5	14 R2 LL P5	16 R6 UL P7	18 R5 LL P7	20 R3 LL P4	22	24	26	28	30	32	34	36

Rack 2 Spine (R2SS)																	
1	3	5 R1 UL P31	7 R7 UL P7	9 R8 LL P7	11 R3 UL P6	13 R1 LL P31	15 R8 UL P7	17 R4 LL P4	19 R7 LL P7	21	23	25	27	29	31	33	35
2	4	6 R2 UL P31	8 R5 UL P4	10 R4 UL P6	12 R6 LL P4	14 R2 LL P31	16 R6 UL P4	18 R5 LL P6	20 R3 LL P31	22	24	26	28	30	32	34	36

Rack 3 Spine (R3SS)																	
1	3	5 R1 UL P6	7 R6 LL P31	9 R7 UL P4	11 R3 UL P30	13 R1 LL P4	15 R7 LL P4	17 R4 LL P31	19 R6 UL P6	21	23	25	27	29	31	33	35
2	4	6 R2 UL P6	8 R5 UL P6	10 R4 UL P33	12 R8 LL P4	14 R2 LL P6	16 R5 LL P31	18 R8 UL P4	20 R3 LL P32	22	24	26	28	30	32	34	36

Rack 4 Spine (R4SS)																	
1	3	5 R1 UL P4	7 R8 LL P6	9 R6 LL P33	11 R3 UL P4	13 R1 LL P30	15 R6 UL P31	17 R4 LL P30	19 R5 LL P33	21	23	25	27	29	31	33	35
2	4	6 R2 UL P4	8 R7 UL P6	10 R4 UL P32	12 R5 UL P33	14 R2 LL P4	16 R8 UL P6	18 R7 LL P6	20 R3 LL P33	22	24	26	28	30	32	34	36



The following table describes the cable connections for the first spine switch (R1-SS) when cabling eight racks together:

Table 23-28 Leaf Switch Connections for the First Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R1 UL within Rack 1	R1-UL-P5 to R1-SS-P5	3 meters
R1 UL to Rack 2	R1-UL-P31 to R2-SS-P5	5 meters
R1 UL to Rack 3	R1-UL-P6 to R3-SS-P5	5 meters
R1 UL to Rack 4	R1-UL-P4 to R4-SS-P5	10 meters
R1 UL to Rack 5	R1-UL-P30 to R5-SS-P5	10 meters
R1 UL to Rack 6	R1-UL-P32 to R6-SS-P5	10 meters

Table 23-28 (Cont.) Leaf Switch Connections for the First Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R1 UL to Rack 7	R1-UL-P33 to R7-SS-P5	10 meters
R1 UL to Rack 8	R1-UL-P7 to R8-SS-P5	10 meters
R1 LL within Rack 1	R1-LL-P5 to R1-SS-P13	3 meters
R1 LL to Rack 2	R1-LL-P31 to R2-SS-P13	5 meters
R1 LL to Rack 3	R1-LL-P4 to R3-SS-P13	5 meters
R1 LL to Rack 4	R1-LL-P30 to R4-SS-P13	10 meters
R1 LL to Rack 5	R1-LL-P7 to R5-SS-P13	10 meters
R1 LL to Rack 6	R1-LL-P32 to R6-SS-P13	10 meters
R1 LL to Rack 7	R1-LL-P6 to R7-SS-P13	10 meters
R1 LL to Rack 8	R1-LL-P33 to R8-SS-P13	10 meters

The following table describes the cable connections for the second spine switch (R2-SS) when cabling eight full racks together:

Table 23-29 Leaf Switch Connections for the Second Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R2 UL to Rack 1	R2-UL-P5 to R1-SS-P6	5 meters
R2 UL within Rack 2	R2-UL-P31 to R2-SS-P6	3 meters
R2 UL to Rack 3	R2-UL-P6 to R3-SS-P6	5 meters
R2 UL to Rack 4	R2-UL-P4 to R4-SS-P6	5 meters
R2 UL to Rack 5	R2-UL-P7 to R5-SS-P6	10 meters
R2 UL to Rack 6	R2-UL-P33 to R6-SS-P6	10 meters
R2 UL to Rack 7	R2-UL-P32 to R7-SS-P6	10 meters
R2 UL to Rack 8	R2-UL-P30 to R8-SS-P6	10 meters
R2 LL to Rack 1	R2-LL-P5 to R1-SS-P14	5 meters
R2 LL within Rack 2	R2-LL-P31 to R2-SS-P14	3 meters
R2 LL to Rack 3	R2-LL-P6 to R3-SS-P14	5 meters
R2 LL to Rack 4	R2-LL-P4 to R4-SS-P14	5 meters
R2 LL to Rack 5	R2-LL-P33 to R5-SS-P14	10 meters
R2 LL to Rack 6	R2-LL-P7 to R6-SS-P14	10 meters
R2 LL to Rack 7	R2-LL-P32 to R7-SS-P14	10 meters
R2 LL to Rack 8	R2-LL-P30 to R8-SS-P14	10 meters

The following table describes the cable connections for the third spine switch (R3-SS) when cabling eight full racks together:

Table 23-30 Leaf Switch Connections for the Third Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R3 UL to Rack 1	R3-UL-P5 to R1-SS-P11	5 meters
R3 UL to Rack 2	R3-UL-P6 to R2-SS-P11	5 meters
R3 UL within Rack 3	R3-UL-P30 to R3-SS-P11	3 meters
R3 UL to Rack 4	R3-UL-P4 to R4-SS-P11	5 meters
R3 UL to Rack 5	R3-UL-P32 to R5-SS-P11	5 meters
R3 UL to Rack 6	R3-UL-P7 to R6-SS-P11	10 meters
R3 UL to Rack 7	R3-UL-P31 to R7-SS-P11	10 meters
R3 UL to Rack 8	R3-UL-P33 to R8-SS-P11	10 meters
R3 LL to Rack 1	R3-LL-P4 to R1-SS-P20	5 meters
R3 LL to Rack 2	R3-LL-P31 to R2-SS-P20	5 meters
R3 LL within Rack 3	R3-LL-P32 to R3-SS-P20	3 meters
R3 LL to Rack 4	R3-LL-P33 to R4-SS-P20	5 meters
R3 LL to Rack 5	R3-LL-P5 to R5-SS-P20	5 meters
R3 LL to Rack 6	R3-LL-P30 to R6-SS-P20	10 meters
R3 LL to Rack 7	R3-LL-P7 to R7-SS-P20	10 meters
R3 LL to Rack 8	R3-LL-P6 to R8-SS-P20	10 meters

The following table describes the cable connections for the fourth spine switch (R4-SS) when cabling eight full racks together:

Table 23-31 Leaf Switch Connections for the Fourth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R4 UL to Rack 1	R4-UL-P7 to R1-SS-P10	10 meters
R4 UL to Rack 2	R4-UL-P6 to R2-SS-P10	5 meters
R4 UL to Rack 3	R4-UL-P33 to R3-SS-P10	5 meters
R4 UL within Rack 4	R4-UL-P32 to R4-SS-P10	3 meters
R4 UL to Rack 5	R4-UL-P5 to R5-SS-P10	5 meters
R4 UL to Rack 6	R4-UL-P30 to R6-SS-P10	5 meters
R4 UL to Rack 7	R4-UL-P31 to R7-SS-P10	10 meters
R4 UL to Rack 8	R4-UL-P4 to R8-SS-P10	10 meters
R4 LL to Rack 1	R4-LL-P5 to R1-SS-P17	10 meters
R4 LL to Rack 2	R4-LL-P4 to R2-SS-P17	5 meters
R4 LL to Rack 3	R4-LL-P31 to R3-SS-P17	5 meters
R4 LL within Rack 4	R4-LL-P30 to R4-SS-P17	3 meters
R4 LL to Rack 5	R4-LL-P7 to R5-SS-P17	5 meters
R4 LL to Rack 6	R4-LL-P6 to R6-SS-P17	5 meters
R4 LL to Rack 7	R4-LL-P32 to R7-SS-P17	10 meters
R4 LL to Rack 8	R4-LL-P33 to R8-SS-P17	10 meters

The following table describes the cable connections for the fifth spine switch (R5-SS) when cabling eight full racks together:

Table 23-32 Leaf Switch Connections for the Fifth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R5 UL to Rack 1	R5-UL-P5 to R1-SS-P8	10 meters
R5 UL to Rack 2	R5-UL-P4 to R2-SS-P8	10 meters
R5 UL to Rack 3	R5-UL-P6 to R3-SS-P8	5 meters
R5 UL to Rack 4	R5-UL-P33 to R4-SS-P12	5 meters
R5 UL within Rack 5	R5-UL-P30 to R5-SS-P7	3 meters
R5 UL to Rack 6	R5-UL-P7 to R6-SS-P7	5 meters
R5 UL to Rack 7	R5-UL-P32 to R7-SS-P7	5 meters
R5 UL to Rack 8	R5-UL-P31 to R8-SS-P7	10 meters
R5 LL to Rack 1	R5-LL-P7 to R1-SS-P18	10 meters
R5 LL to Rack 2	R5-LL-P6 to R2-SS-P18	10 meters
R5 LL to Rack 3	R5-LL-P31 to R3-SS-P16	5 meters
R5 LL to Rack 4	R5-LL-P33 to R4-SS-P19	5 meters
R5 LL within Rack 5	R5-LL-P30 to R5-SS-P15	3 meters
R5 LL to Rack 6	R5-LL-P4 to R6-SS-P15	5 meters
R5 LL to Rack 7	R5-LL-P5 to R7-SS-P15	5 meters
R5 LL to Rack 8	R5-LL-P32 to R8-SS-P15	10 meters

The following table describes the cable connections for the sixth spine switch (R6-SS) when cabling eight full racks together:

Table 23-33 Leaf Switch Connections for the Sixth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R6 UL to Rack 1	R6-UL-P7 to R1-SS-P16	10 meters
R6 UL to Rack 2	R6-UL-P4 to R2-SS-P16	10 meters
R6 UL to Rack 3	R6-UL-P6 to R3-SS-P19	10 meters
R6 UL to Rack 4	R6-UL-P31 to R4-SS-P15	5 meters
R6 UL to Rack 5	R6-UL-P33 to R5-SS-P18	5 meters
R6 UL within Rack 6	R6-UL-P30 to R6-SS-P16	3 meters
R6 UL to Rack 7	R6-UL-P5 to R7-SS-P16	5 meters
R6 UL to Rack 8	R6-UL-P32 to R8-SS-P16	5 meters
R6 LL to Rack 1	R6-LL-P5 to R1-SS-P12	10 meters
R6 LL to Rack 2	R6-UL-P4 to R2-SS-P12	10 meters
R6 LL to Rack 3	R6-LL-P31 to R3-SS-P7	10 meters
R6 LL to Rack 4	R6-LL-P33 to R4-SS-P9	5 meters
R6 LL to Rack 5	R6-LL-P30 to R5-SS-P8	5 meters
R6 LL within Rack 6	R6-LL-P32 to R6-SS-P8	3 meters

Table 23-33 (Cont.) Leaf Switch Connections for the Sixth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R6 LL to Rack 7	R6-LL-P6 to R7-SS-P8	5 meters
R6 LL to Rack 8	R6-LL-P7 to R8-SS-P8	5 meters

The following table describes the cable connections for the seventh spine switch (R7-SS) when cabling eight full racks together:

Table 23-34 Leaf Switch Connections for the Seventh Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R7 UL to Rack 1	R7-UL-P5 to R1-SS-P9	10 meters
R7 UL to Rack 2	R7-UL-P7 to R2-SS-P7	10 meters
R7 UL to Rack 3	R7-UL-P4 to R3-SS-P9	10 meters
R7 UL to Rack 4	R7-UL-P6 to R4-SS-P8	10 meters
R7 UL to Rack 5	R7-UL-P31 to R5-SS-P9	5 meters
R7 UL to Rack 6	R7-UL-P33 to R6-SS-P9	5 meters
R7 UL within Rack 7	R7-UL-P30 to R7-SS-P9	3 meters
R7 UL to Rack 8	R7-UL-P32 to R8-SS-P9	5 meters
R7 LL to Rack 1	R7-LL-P5 to R1-SS-P19	10 meters
R7 LL to Rack 2	R7-LL-P7 to R2-SS-P19	10 meters
R7 LL to Rack 3	R7-LL-P4 to R3-SS-P15	10 meters
R7 LL to Rack 4	R7-LL-P6 to R4-SS-P18	10 meters
R7 LL to Rack 5	R7-LL-P31 to R5-SS-P16	5 meters
R7 LL to Rack 6	R7-LL-P33 to R6-SS-P18	5 meters
R7 LL within Rack 7	R7-LL-P30 to R7-SS-P18	3 meters
R7 LL to Rack 8	R7-LL-P32 to R8-SS-P12	5 meters

The following table describes the cable connections for the eighth spine switch (R8-SS) when cabling eight full racks together:

Table 23-35 Leaf Switch Connections for the Eighth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R8 UL to Rack 1	R8-UL-P5 to R1-SS-P15	10 meters
R8 UL to Rack 2	R8-UL-P7 to R2-SS-P15	10 meters
R8 UL to Rack 3	R8-UL-P4 to R3-SS-P18	10 meters
R8 UL to Rack 4	R8-UL-P6 to R4-SS-P16	10 meters
R8 UL to Rack 5	R8-UL-P31 to R5-SS-P19	10 meters
R8 UL to Rack 6	R8-UL-P33 to R6-SS-P19	5 meters
R8 UL to Rack 7	R8-UL-P30 to R7-SS-P19	5 meters
R8 UL within Rack 8	R8-UL-P32 to R8-SS-P19	3 meters

Table 23-35 (Cont.) Leaf Switch Connections for the Eighth Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R8 LL to Rack 1	R8-LL-P5 to R1-SS-P7	10 meters
R8 LL to Rack 2	R8-LL-P7 to R2-SS-P9	10 meters
R8 LL to Rack 3	R8-LL-P4 to R3-SS-P12	10 meters
R8 LL to Rack 4	R8-LL-P6 to R4-SS-P7	10 meters
R8 LL to Rack 5	R8-LL-P31 to R5-SS-P12	10 meters
R8 LL to Rack 6	R8-LL-P33 to R6-SS-P12	5 meters
R8 LL to Rack 7	R8-LL-P30 to R7-SS-P12	5 meters
R8 LL within Rack 8	R8-LL-P32 to R8-SS-P18	3 meters

Multi-Rack Cabling Tables for ZDLRA Rack Models with InfiniBand Network Fabric (X2 to X8)

This section contains multi-rack cabling tables for ZDLRA Rack models that use InfiniBand Network Fabric. This includes ZDLRA Rack models from X2 to X8.

Understanding Multi-Rack Cabling for Racks with InfiniBand Network Fabric

Up to eight racks can be cabled together without external RDMA Network Fabric switches.

Three RDMA Network Fabric switches are needed when cabling racks together. If any rack does not include a third switch, then order the expansion switch kit. These switches attach to standard RDMA Network Fabric connectors at the end of the cables. The procedures in this section assume the racks are adjacent to each other. If they are not, then longer cables may be required for the connections.



Note:

Only for InfiniBand Network Fabric (X8 and earlier).

- For Eighth or Quarter Racks, which are the smallest Elastic Configurations, follow the instructions in "[Cabling Oracle Exadata Quarter Racks and Oracle Exadata Eighth Racks with InfiniBand Network Fabric](#)" for direct connection without spine switch.
- For other racks (Half Rack, Full Rack, Elastic Configurations larger than Eighth or Quarter Rack) install a spine switch and follow the standard *N*-rack cabling for two or more interconnected racks, for example "[Two-Rack Cabling with InfiniBand Network Fabric](#)."

The switch at rack unit 1 (U1) is referred to as the spine switch.

The following switches are referred to as leaf switches:


- The switches at rack unit 20 (U20) and rack unit 24 (U24) in Recovery Appliance X2-2 racks and Oracle Exadata Storage Expansion Racks
- The switches at rack unit 21 (U21) and rack unit 23 (U23) in Recovery Appliance X2-8 and later.
- The switches at rack unit 20 (U20) and 22 (U22) in Recovery Appliance X3-2 and later racks or Storage Expansion Rack X3-2 and later.

Note:

Oracle Exadata Database Machine X4-2 and later racks or Oracle Exadata Database Machine X3-8 Full Racks with Exadata Storage Server X4-2L Servers do not include spare cables or a third Sun Datacenter InfiniBand Switch 36 switch. To extend Oracle Exadata Database Machine X4-2 and later racks or Oracle Exadata Database Machine X3-8 Full Racks with Exadata Storage Server X4-2L Servers, you must order cables and a Sun Datacenter InfiniBand Switch 36 switch.


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 the following graphic:

Figure 24-1 Connections Between Spine Switch and Leaf Switches in a Single Rack




The Oracle Database servers and Exadata Storage Servers connect to the leaf switches as shown in the following graphic:

Figure 24-2 Connections Between Database Servers and Storage Servers and Leaf Switches



When connecting up to eight racks together, remove the seven existing inter-switch connections between each leaf switch, as well as 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 multi-rack environments, the leaf switches inside a rack are no longer directly interconnected, as shown in the following graphic:

Figure 24-3 Connections Between Spine Switches and Leaf Switches Across Two Racks



As shown in the preceding graphic, each leaf switch in rack 1 connects to the following switches:


- Four connections to its internal spine switch
- Four connections to the spine switch in rack 2

The spine switch in rack 1 connects to the following switches:

- Eight connections to both internal leaf switches
- Eight connections to both leaf switches in rack 2

As the number of racks increases from two to eight, the pattern continues as shown in the following graphic:

Figure 24-4 Connections Between Spine Switches and Leaf Switches for up to 8 Racks



As shown in the preceding graphic, each leaf switch has eight inter-switch connections distributed over all spine switches. Each spine switch has 16 inter-switch connections distributed over all leaf switches. The leaf switches are not directly interconnected with other leaf switches, and the spine switches are not directly interconnected with the other spine switches.

**Note:**

Cable lengths are specified for racks 1 through 8.

Preparing for Multi-Rack Cabling with InfiniBand Network Fabric

Racks can be added together to increase system capacity and performance. When cabling racks together, note the following:

- The cable lengths shown in this document assume the racks are adjacent to each other. If the racks are not adjacent, or there are obstacles in the raised floor, or if you use overhead cabling, then longer cables may be required. For optical cables, the maximum supported cable length is 100 meters. For copper cables, the maximum supported cable length is 5 meters.
- Oracle recommends that the names for the servers include the rack unit number. This helps identify the server during diagnostics.
- When completing Oracle Exadata Deployment Assistant for the additional rack, you are prompted for SCAN addresses. However, these SCAN addresses are not used because the SCAN address from the original rack are used. Manually remove the new SCAN addresses from the generated installation files.
- The software owner account names and group names, as well as their identifiers, must match the names and identifiers of the original rack.
- If the additional grid disks are used with existing disk groups, then ensure the grid disk sizes for the new rack are the same as the original rack.
- If the InfiniBand network consists of four or more racks cabled together, then disable the Subnet Manager on the leaf switches.
- Verify the Master Subnet Manager is located on the spine switch.
- Oracle Exadata Database Machine Quarter Racks can be extended as follows:
 - Connect two Oracle Exadata Database Machine Quarter Racks together. At least four of the six ports reserved for external connectivity are open on each leaf switch. The six ports are 5B, 6A, 6B, 7A, 7B, and 12A in each leaf switch. Maintain the existing seven inter-switch links between the leaf switches within each rack. Connect the leaf switches between the racks with two links each, using the ports reserved for external connectivity.
 - Connect one Oracle Exadata Database Machine Quarter Rack with one Oracle Exadata Database Machine Half Rack or one Oracle Exadata Database Machine Full Rack. At least four ports reserved for external connectivity are open on each leaf switch. The spine switch in the Oracle Exadata Database Machine Half Rack or Oracle Exadata Database Machine Full Rack remains as the spine switch. Maintain the existing seven inter-switch links between the leaf switches within each rack. Connect the leaf switches between the racks with two links each, using the ports reserved for external connectivity.

- Connect one Oracle Exadata Database Machine Quarter Rack with two or more Oracle Exadata Database Machine Half Racks or Oracle Exadata Database Machine Full Racks. The racks are interconnected using a fat-tree topology. Connect each leaf switch in the quarter rack to the spine switch of each half rack or full rack using two links each. If there are more than four racks, then use one link instead of two. The seven inter-switch links between the leaf switches in the quarter rack are removed.

 Note:

To connect more than one quarter rack to additional racks, it is necessary to purchase Sun Datacenter InfiniBand Switch 36 switches for the quarter racks.

- If you are extending Oracle Exadata Database Machine X4-2 or later, or Oracle Exadata Database Machine X3-8 Full Rack, or Oracle Exadata Database Machine X2-2 (with X4170 and X4275 servers) half rack, then order the expansion kit that includes a Sun Datacenter InfiniBand Switch 36 switch.

Perform the following tasks before cabling racks together:

1. Determine the number of racks that will be cabled together.
2. Count the spare cables from the kit, and existing inter-switch cables.

 Note:

Oracle Exadata Database Machine X4-2 and later racks or Oracle Exadata Database Machine X3-8 Full Racks with Exadata Storage Server X4-2L Servers do not include spare cables or a third Sun Datacenter InfiniBand Switch 36 switch. To extend Oracle Exadata Database Machine X4-2 and later racks or Oracle Exadata Database Machine X3-8 Full Racks with Exadata Storage Server X4-2L Servers, you must order cables and a Sun Datacenter InfiniBand Switch 36 switch.

For Oracle Exadata Racks earlier than Oracle Exadata Database Machine X4-2, no additional InfiniBand cables need to be purchased when connecting up to three Oracle Exadata Database Machine Full Racks. The following table lists the spare cables for the switch:

Cable	Spare	Inter-switch	Total
InfiniBand 2 meter cable	0	7	7
InfiniBand 3 meter cable	6	2	8
InfiniBand 5 meter cable	10	0	10

When connecting four or more racks, it is necessary to purchase additional 10 meter optical QSFP-to-QSFP QDR InfiniBand cables. Two QSFP SR transceivers are needed for each optical cable. The following table lists the number of racks, and cables needed to interconnect the racks.

Number of Racks	3 Meter Cables Needed	5 Meter Cables Needed	10 Meter Cables Needed
2	16	16	0
3	18	30	0
4	16	40	8
5	20	42	18
6	24	46	26
7	28	44	40
8	16	52	60

3. Gather the following tools:

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

4. Determine a naming method for the rack prefixes. For example, if the original rack has the prefix `dbm01`, then use the prefix `dbm02` for the second rack, the prefix `dbm03` for the third rack, and so on.

5. Verify the racks have unique host names and IP addresses. All servers interconnected in the racks must have unique names and IP addresses.

Server names and IP addresses conventions may differ in the following cases:

- Initial installation of all Oracle Exadata Database Machine Full Racks: System address assignments and host names should be complete.
 - New Oracle Exadata Database Machine Full Racks are added to an existing cluster: The new rack configuration should require unique host names and IP addresses for the new Oracle Exadata Database Machines. The IP addresses on the same subnet cannot conflict with the existing systems.
 - Two existing Oracle Exadata Database Machine Full Racks are clustered together: You can assign host names and IP addresses only if Oracle Exadata Database Machines are already assigned unique host names and IP addresses, or the entire cluster must be reconfigured. The machines must be on the same subnet and not have conflicting IP addresses.
6. Ensure the IP addresses for the new servers are in the same subnet, and do not overlap with the currently-installed servers.
7. Ensure the firmware on the original switches are at the same level as the new switches using the `nm2version` command. If the firmware is not at the same level, then apply a firmware patch.

Cabling Oracle Exadata Quarter Racks and Oracle Exadata Eighth Racks with InfiniBand Network Fabric

Oracle Exadata Quarter Racks and Oracle Exadata Eighth Racks with InfiniBand Network Fabric can be cabled as follows:

- Oracle Exadata Quarter Rack to Oracle Exadata Quarter Rack


- Oracle Exadata Quarter Rack to Oracle Exadata Half Rack, or multiple Oracle Exadata Half Racks or Oracle Exadata Full Racks
- Oracle Exadata Quarter Rack to Oracle Exadata Full Rack, or multiple Oracle Exadata Full Racks or Oracle Exadata Half Racks
- Oracle Exadata Eighth Rack to Oracle Exadata Eighth Rack
- Oracle Exadata Eighth Rack to Oracle Exadata Half Rack, or multiple Oracle Exadata Half Racks or Oracle Exadata Full Racks
- Oracle Exadata Eighth Rack to Oracle Exadata Full Rack, or multiple Oracle Exadata Full Racks or Oracle Exadata Half Racks

 Note:

- To cable additional Oracle Exadata Quarter Racks or Oracle Exadata Eighth Racks, it is necessary to purchase additional Sun Datacenter InfiniBand Switch 36 switches.
- It is assumed that four of the six ports dedicated for external connectivity in Oracle Exadata Half Rack or Oracle Exadata Full Rack are unused and available for switch interconnections with Oracle Exadata Quarter Rack or Oracle Exadata Eighth Rack. The six ports are 5B, 6A, 6B, 7A, 7B, and 12A in each leaf switch.


The following graphic shows the cable connections for two Oracle Exadata Quarter Racks. The leaf switches within each rack maintain their existing seven connections. The leaf switches interconnect between the racks with two links each using the ports reserved for external connectivity.

Figure 24-5 Leaf Switch Connections Between 2 Quarter Racks




The following graphic shows the cable connections from Oracle Exadata Quarter Rack to Oracle Exadata Half Rack or Oracle Exadata Full Rack. The leaf switches within each rack maintain their existing seven connections. The leaf switches interconnect between the racks with two links each using the ports reserved for external connectivity.

Figure 24-6 Leaf and Spine Switch Connections Between a Quarter Rack and a Half or Full Rack



The following graphic shows the cable connections from Oracle Exadata Quarter Rack to two or more racks. The racks that connect to Oracle Exadata Quarter Rack must be all Oracle Exadata Half Racks or Oracle Exadata Full Racks, interconnected using a fat-tree topology. Each leaf switch in Oracle Exadata Quarter Rack connects to the spine switches in the other half racks or full racks with two links each. If there are more than four racks, then use one link instead of two. The seven inter-switch links between the leaf switches in the quarter rack are removed.

Figure 24-7 Leaf and Spine Switch Connections for a Quarter Rack Connected to One or More Half or Full Racks



Two-Rack Cabling with InfiniBand Network Fabric

Review this information before cabling two racks together with InfiniBand Network Fabric.

- The spine switch is in U1 (also referred to as IB1) for all racks.

- In the following lists, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine 8-socket (X8-8, X7-8, X6-8, X5-8, X4-8, X3-8 and X2-8) Full Rack systems, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, Oracle Exadata Database Machine X3-2 Full Racks and Oracle Exadata Storage Expansion Rack X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3, respectively.
- The cable lengths shown in the following lists assume that the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Cables lengths up to 100 meters are supported.

The following table shows the cable connections for the first spine switch (R1-U1) when cabling two full racks together.

Table 24-1 Leaf Switch Connections for the First Rack 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

The following table shows the cable connections for the second spine switch (R2-U1) when cabling two full racks together.

Table 24-2 Leaf Switch Connections for the Second Rack 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

Related Topics

- [Cabling Two Racks Together](#)
Choose from the available methods based on your system specifications and operational requirements.

Three-Rack Cabling with InfiniBand Network Fabric

Note the following when cabling three racks together with InfiniBand Network Fabric:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata Database Machine X3-2 Full Racks or Oracle Exadata Storage Expansion Rack X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3, respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 24-3 shows the cable connections for the first spine switch (R1-U1) when cabling three racks together:

Table 24-3 Leaf Switch Connections for the First Rack 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 24-4 shows the cable connections for the second spine switch (R2-U1) when cabling three racks together:

Table 24-4 Leaf Switch Connections for the Second Rack 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 24-5 shows the cable connections for the third spine switch (R3-U1) when cabling three full racks together:

Table 24-5 Leaf Switch Connections for the Third Rack 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

Four-Rack Cabling with InfiniBand Network Fabric

Note the following when cabling four racks together with InfiniBand Network Fabric:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata Database Machine X3-2 Full Racks or Oracle Exadata Storage Expansion Rack X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3, respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 24-6 shows the cable connections for the first spine switch (R1-U1) when cabling four racks together:

Table 24-6 Leaf Switch Connections for the First Rack 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 24-7 shows the cable connections for the second spine switch (R2-U1) when cabling four full racks together:

Table 24-7 Leaf Switch Connections for the Second Rack 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
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 24-8 shows the cable connections for the third spine switch (R3-U1) when cabling four full racks together:

Table 24-8 Leaf Switch Connections for the Third Rack 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 24-9 shows the cable connections for the fourth spine switch (R4-U1) when cabling four full racks together:

Table 24-9 Leaf Switch Connections for the Fourth Rack 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
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

Five-Rack Cabling with InfiniBand Network Fabric

Note the following when cabling five racks together with InfiniBand Network Fabric:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later racks, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata Database Machine X3-2 Full Racks and Oracle Exadata Storage Expansion Rack X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3, respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 24-10 shows the cable connections for the first spine switch (R1-U1) when cabling five racks together:

Table 24-10 Leaf Switch Connections for the First Rack 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 24-11 shows the cable connections for the second spine switch (R2-U1) when cabling five full racks together:

Table 24-11 Leaf Switch Connections for the Second Rack 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 24-12 shows the cable connections for the third spine switch (R3-U1) when cabling five full racks together:

Table 24-12 Leaf Switch Connections for the Third Rack 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 24-13 shows the cable connections for the fourth spine switch (R4-U1) when cabling five full racks together:

Table 24-13 Leaf Switch Connections for the Fourth Rack 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 24-14 shows the cable connections for the fifth spine switch (R5-U1) when cabling five full racks together:

Table 24-14 Leaf Switch Connections for the Fifth Rack 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-IB2-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

Six-Rack Cabling with InfiniBand Network Fabric

Note the following when cabling six racks together with InfiniBand Network Fabric:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later 8-socket racks, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later two-socket racks, or Oracle Exadata Database Machine X3-2 Full Racks and Oracle Exadata Storage Expansion Rack X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3, respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 24-15 shows the cable connections for the first spine switch (R1-U1) when cabling six racks together:

Table 24-15 Leaf Switch Connections for the First Rack 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 24-16 shows the cable connections for the second spine switch (R2-U1) when cabling six full racks together:

Table 24-16 Leaf Switch Connections for the Second Rack 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
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 24-17 shows the cable connections for the third spine switch (R3-U1) when cabling six full racks together:

Table 24-17 Leaf Switch Connections for the Third Rack 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 24-18 shows the cable connections for the fourth spine switch (R4-U1) when cabling six full racks together:

Table 24-18 Leaf Switch Connections for the Fourth Rack 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 24-19 shows the cable connections for the fifth spine switch (R5-U1) when cabling six full racks together:

Table 24-19 Leaf Switch Connections for the Fifth Rack 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-IB2-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 24-20 shows the cable connections for the sixth spine switch (R6-U1) when cabling six full racks together:

Table 24-20 Leaf Switch Connections for the Sixth Rack 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

Seven-Rack Cabling with InfiniBand Network Fabric

Note the following when cabling seven racks together with InfiniBand Network Fabric:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.
- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later racks, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata Database Machine X3-2 Full Racks and Oracle Exadata Storage Expansion Rack X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3, respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 24-21 shows the cable connections for the first spine switch (R1-U1) when cabling seven racks together:

Table 24-21 Leaf Switch Connections for the First Rack 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
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 24-22 shows the cable connections for the second spine switch (R2-U1) when cabling seven full racks together:

Table 24-22 Leaf Switch Connections for the Second Rack 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 24-23 shows the cable connections for the third spine switch (R3-U1) when cabling seven full racks together:

Table 24-23 Leaf Switch Connections for the Third Rack 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
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 24-24 shows the cable connections for the fourth spine switch (R4-U1) when cabling seven full racks together:

Table 24-24 Leaf Switch Connections for the Fourth Rack 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

Table 24-24 (Cont.) Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 IB2 to Rack 7	R4-IB2-P10A to R7-U1-P7B	10 meters

Table 24-25 shows the cable connections for the fifth spine switch (R5-U1) when cabling seven full racks together:

Table 24-25 Leaf Switch Connections for the Fifth Rack 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-IB2-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 24-26 shows the cable connections for the sixth spine switch (R6-U1) when cabling seven full racks together:

Table 24-26 Leaf Switch Connections for the Sixth Rack 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

Table 24-26 (Cont.) Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
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-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
R6 IB2 to Rack 7	R6-IB2-P9A to R7-U1-P5B	5 meters

Table 24-27 shows the cable connections for the seventh spine switch (R7-U1) when cabling seven full racks together:

Table 24-27 Leaf Switch Connections for the Seventh Rack 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

Eight-Rack Cabling with InfiniBand Network Fabric

Note the following when cabling eight racks together with InfiniBand Network Fabric:

- The spine switch is in U1 (also referred to as IB1) for all racks.
- In the following tables, the leaf switches are referred to as IB2 and IB3. Use the rack unit that is appropriate for your system.

- In Oracle Exadata Database Machine X2-2 racks and Oracle Exadata Storage Expansion Racks with Exadata Storage Server with Sun Fire X4270 M2 Servers, the leaf switches are in U20 and U24, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X2-8 and later racks, the leaf switches are in U21 and U23, also referred to as IB2 and IB3, respectively.
- In Oracle Exadata Database Machine X4-2 and later racks, or Oracle Exadata Database Machine X3-2 Full Racks and Oracle Exadata Storage Expansion Rack X3-2 racks, the leaf switches are in U20 and U22, also referred to as IB2 and IB3, respectively.
- The cable lengths shown in the tables assume the racks are adjacent to each other. If the racks are not adjacent or use overhead cabling trays, then they may require longer cables lengths. Up to 100 meters is supported.
- Only optical cables are supported for lengths greater than 5 meters.

Table 24-28 shows the cable connections for the first spine switch (R1-U1) when cabling eight racks together:

Table 24-28 Leaf Switch Connections for the First Rack in a 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
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 24-29 shows the cable connections for the second spine switch (R2-U1) when cabling eight full racks together:

Table 24-29 Leaf Switch Connections for the Second Rack in a 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

Table 24-29 (Cont.) Leaf Switch Connections for the Second Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
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 24-30 shows the cable connections for the third spine switch (R3-U1) when cabling eight full racks together:

Table 24-30 Leaf Switch Connections for the Third Rack in a 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
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 24-31 shows the cable connections for the fourth spine switch (R4-U1) when cabling eight full racks together:

Table 24-31 Leaf Switch Connections for the Fourth Rack in a 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 24-32 shows the cable connections for the fifth spine switch (R5-U1) when cabling eight full racks together:

Table 24-32 Leaf Switch Connections for the Fifth Rack 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-IB2-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 24-33 shows the cable connections for the sixth spine switch (R6-U1) when cabling eight full racks together:

Table 24-33 Leaf Switch Connections for the Sixth Rack 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-IB2-P8B to R7-U1-P4B	5 meters
R6 IB2 to Rack 8	R6-IB2-P9A to R8-U1-P5B	5 meters

Table 24-34 shows the cable connections for the seventh spine switch (R7-U1) when cabling eight full racks together:

Table 24-34 Leaf Switch Connections for the Seventh Rack 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-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

Table 24-34 (Cont.) Leaf Switch Connections for the Seventh Rack in a Eight-Rack System

Leaf Switch	Connection	Cable Length
R7 IB2 to Rack 6	R7-IB2-P11B to R6-U1-P10B	5 meters
R7 IB2 to Rack 8	R7-IB2-P8B to R8-U1-P4B	5 meters

Table 24-35 shows the cable connections for the eighth spine switch (R8-U1) when cabling eight full racks together:

Table 24-35 Leaf Switch Connections for the Eighth Rack 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
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-IB2-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-IB2-P11A to R6-U1-P9B	5 meters
R8 IB2 to Rack 7	R8-IB2-P11B to R7-U1-P10B	5 meters

Part VII

Appendixes

Zero Data Loss Recovery Appliance Owner's Guide appendixes are:

- [Error Messages](#)
- [Replacement Units](#)

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](#)
- [Generic Alert Messages: HALRT-00001 to HALRT-00500](#)
- [Temperature Alert Messages: HALRT-00513 to HALRT-00600](#)
- [Fan Alert Messages: HALRT-00616 to HALRT-00700](#)
- [Power Supply/Converters Alert Messages: HALRT-00717 to HALRT-00800](#)
- [Network Alert Messages: HALRT-00912 to HALRT-01000](#)
- [Hard Disk and Flash Disk Alert Messages: HALRT-01015 to HALRT-01200](#)
- [Accelerator/Controller Alert Messages: HALRT-01209 to HALRT-01300](#)
- [Voltage Alert Messages: HALRT-01301 to HALRT-01400](#)
- [Physical Disk Alert Messages: HALRT-01401 to HALRT-01500](#)
- [Auto Service Request Alert Messages for Storage Servers: HALRT-02001 to HALRT-02013](#)
- [HDD Controller Battery Alert Messages: HALRT-02101 to HALRT-02200](#)
- [Generic Actions for Alert Messages: HALRT-10001 to HALRT-10500](#)
- [Generic Actions for Alert Messages: HALRT-10503 to HALRT-11000](#)
- [Generic Actions for Alert Messages: HALRT-11001 to HALRT-11100](#)



See Also:

Oracle Database Error Messages for information about Oracle Database SNMP error messages

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

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](#)".

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](#)".

HALRT-00004, CPU Internal Error

Cause: Detected an Internal CPU error.

Action: See "[HALRT-10025](#)".

HALRT-00005, CPU Internal Error Cleared

Cause: Detected a clear event for CPU internal error.

Action: See "[HALRT-10001](#)".

HALRT-00006, CPU Configuration Error

Cause: Detected a CPU configuration error.

Action: See "[HALRT-10025](#)".

HALRT-00007, CPU Configuration Error Cleared

Cause: Detected a clear event for a CPU configuration error.

Action: See "[HALRT-10001](#)".

HALRT-00008, CPU Presence (Processor Presence detected)

Cause: Detected presence event for a new CPU.

Action: See "[HALRT-10025](#)".

HALRT-00009, CPU Not Present (Processor Not Present)

Cause: Detected a CPU is not present.

Action: See "[HALRT-10001](#)".

HALRT-00010, System Event Log (SEL) Full (Logging Disabled)

Cause: System event log detected to be full.

Action: See "[HALRT-10001](#)".

HALRT-00011, ASR (Automatic System Recovery) Timer Expired

Cause: Automatic System Recovery timer expired.

Action: See "[HALRT-10001](#)".

HALRT-00012, ASR (Automatic System Recovery) Reset Occurred

Cause: Automatic System Recovery reset occurred.

Action: See "[HALRT-10001](#)".

HALRT-00013, ASR (Automatic System Recovery) Power Down Occurred

Cause: Automatic System Recovery was powered down.

Action: See "[HALRT-10001](#)".

HALRT-00014, ASR (Automatic System Recovery) Power Cycle Occurred

Cause: Automatic System Recovery power was cycled.

Action: See "[HALRT-10001](#)".

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.](#)".

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](#)".

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](#)".

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.](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

HALRT-00028, This is a test trap.

Cause: User generated a test trap, and an alert was created.

Action: See "[HALRT-10001](#)".

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.

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](#)"

HALRT-00514, Under-Temperature Warning Cleared.

Cause: The low temperature warning was cleared.

Action: See "[HALRT-10001](#)".

HALRT-00515, Under-Temperature Problem (Lower Critical - going low)

Cause: A low temperature problem was detected.

Action: See "[HALRT-10004](#)".

HALRT-00516, Under-Temperature Problem Cleared

Cause: The low temperature problem was cleared.

Action: See "[HALRT-10001](#)".

HALRT-00517, Over-Temperature warning (Upper non-critical, going high)

Cause: The temperature was above the specified range.

Action: "[HALRT-10015](#)".

HALRT-00518, Over-Temperature warning Cleared

Cause: The high temperature warning was cleared.

Action: See "[HALRT-10001](#)".

HALRT-00519, Over-Temperature Problem (Upper Critical - going high)

Cause: A high temperature was detected.

Action: "[HALRT-10004](#)".

HALRT-00520, Over-Temperature Problem Cleared

Cause: The high temperature problem was cleared.

Action: See "[HALRT-10001](#)".

HALRT-00521, CPU Thermal Trip (Over Temperature Shutdown)

Cause: The CPU high temperature shut down was tripped.

Action: See "[HALRT-10004](#)".

HALRT-00522, CPU Thermal Trip (Over Temperature Shutdown) Cleared

Cause: The CPU high temperature shut down was cleared.

Action: See "[HALRT-10001](#)".

Fan Alert Messages: HALRT-00616 to HALRT-00700

HALRT-00616, Generic Critical Fan Failure

Cause: A fan has failed.

Action: See "[HALRT-10005](#)".

HALRT-00617, Generic Critical Fan Failure Cleared

Cause: A fan failure was cleared.

Action: See "[HALRT-10001](#)".

HALRT-00618, Fan Redundancy Lost

Cause: A redundant fan has stopped functioning.

Action: See "[HALRT-10001](#)".

HALRT-00619, Fan redundancy Has Returned to Normal

Cause: A redundant fan has started to function normally.

Action: See "[HALRT-10001](#)".

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](#)".

HALRT-00718, Power Supply Redundancy Has Returned to Normal

Cause: A redundant power supply has started to function normally.

Action: See "[HALRT-10001](#)".

HALRT-00719, Power Supply Inserted

Cause: A power supply was inserted into the system.

Action: See "[HALRT-10001](#)".

HALRT-00720, Power Supply Removed

Cause: A power supply was removed from the system.

Action: "[HALRT-10007](#)".

HALRT-00721, Power Supply Failure

Cause: A power supply has failed.

Action: "[HALRT-10007](#)".

HALRT-00722, Power Supply Failure Cleared

Cause: A power supply failure was cleared.

Action: See "[HALRT-10001](#)".

HALRT-00723, Power Supply Warning

Cause: A power supply problem has set off a warning.

Action: "[HALRT-10007](#)".

HALRT-00724, Power Supply Warning Cleared

Cause: A power supply warning was cleared.

Action: See "[HALRT-10001](#)".

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](#)".

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](#)".

HALRT-00913, InfiniBand port {0} indicates invalid rate: {1} Gps

Cause: InfiniBand device indicated invalid rate.

Action: See "[HALRT-11003](#)".

HALRT-00914, InfiniBand port status is OK

Cause: InfiniBand device status was all right.

Action: See "[HALRT-10001](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

HALRT-01038, {0} output has duplicate serial number: {1}

Cause: The output had a duplicate serial number.

Action: See "[HALRT-10073](#)".

HALRT-01039, {0} output has duplicate OS device name: {1}

Cause: The output had a duplicate operating system device name.

Action: "HALRT-10073".

HALRT-01040, Cannot find device name for the disk with target ID: {0}

Cause: The device was not found.

Action: "HALRT-10073".

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 Disk: {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](#)".

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](#)".

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](#)".

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](#)".

HALRT-01115, BGI completed with uncorrectable errors on Adapter {0} VD Target {1}

Cause: BGI completed with errors.

Action: See "[HALRT-10001](#)".

HALRT-01116, BGI failed on Adapter {0} VD Target {1}

Cause: BGI failed on adapter.

Action: See "[HALRT-10039](#)".

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](#)".

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](#)".

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](#)".

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](#)"[HALRT-10001](#)".

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](#)".

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](#)".

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](#)".

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](#)".

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](#)".

HALRT-01153, Disk drive {0} is back in WriteBack caching mode.

Cause: The disk drive entered write back caching mode.

Action: See "[HALRT-10506](#)".

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"](#).

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](#)".

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](#)".

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](#)".

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](#)".

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.

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](#)".

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](#)".

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](#)".

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](#)".

HALRT-01213, Cache Data Recovered for Adapter {0}.

Cause: Cache data was recovered for the adapter.

Action: See "[HALRT-10001](#)".

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](#)".

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](#)".

HALRT-01216, Flash downloaded image corrupted for Adapter {0}.

Cause: The flash memory card image was corrupted.

Action: See "[HALRT-10034](#)".

HALRT-01217, Flash erase error for Adapter {0}.

Cause: There was a flash erase error for the adapter.

Action: See "[HALRT-10034](#)".

HALRT-01218, Flash timeout during erase for Adapter {0}.

Cause: The flash timed out during the erase procedure.

Action: See "[HALRT-10034](#)".

HALRT-01219, Flash error for Adapter {0}.

Cause: There was a flash error for the adapter.

Action: See "[HALRT-10034](#)".

HALRT-01220, Flashing image: {0} for Adapter {1}.

Cause: The controller flash was upgraded with a new image.

Action: See "[HALRT-10001](#)".

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](#)".

HALRT-01222, Flash programming error for Adapter {0}.

Cause: There was a flash programming error for the adapter.

Action: See "[HALRT-10034](#)".

HALRT-01223, Flash timeout during programming for Adapter {0}.

Cause: The flash timed out while programming of the adapter.

Action: See "[HALRT-10034](#)".

HALRT-01224, Fatal error received for Adapter {0}.

Cause: A fatal error occurred.

Action: See "[HALRT-10035](#)".

HALRT-01225, Rebuild rate changed to {0} for Adapter {1}.

Cause: The rebuild rate changed for the adapter.

Action: See "[HALRT-10001](#)".

HALRT-01226, Alarm has been disabled by user for Adapter {0}.

Cause: The alarm for the adapter was disabled.

Action: See "[HALRT-10001](#)".

HALRT-01227, Configuration cleared for Adapter {0}.

Cause: The configuration for the adapter was cleared.

Action: See "[HALRT-10036](#)".

HALRT-01228, Factory defaults restored for Adapter {0}.

Cause: The factory default settings were restored for the adapter.

Action: See "[HALRT-10036](#)".

HALRT-01229, HDD disk controller battery not Present for Adapter {0}.

Cause: The battery was not present for the adapter.

Action: See "[HALRT-10036](#)".

HALRT-01230, New HDD disk controller battery Detected for Adapter {0}.

Cause: A new battery was detected for the adapter.

Action: See "[HALRT-10001](#)".

HALRT-01231, HDD disk controller battery has been replaced for Adapter {0}.

Cause: The battery was replaced for the adapter.

Action: See "[HALRT-10001](#)".

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](#)".

HALRT-01233, HDD disk controller battery voltage low for Adapter {0}.

Cause: The battery voltage was low for the adapter.

Action: See "[HALRT-10038](#)".

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](#)".

HALRT-01235, HDD disk controller battery needs replacement, SOH Bad for Adapter {0}.

Cause: The battery needed replacement.

Action: See "[HALRT-10038](#)".

HALRT-01236, HDD disk controller battery removed for Adapter {0}.

Cause: The battery was removed for the adapter.

Action: See "[HALRT-10036](#)".

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](#)".

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](#)".

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](#)".

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"

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](#)".

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.

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](#)".

HALRT-01302, Under-Voltage Problem Cleared

Cause: The under-voltage problem was cleared.

Action: See "[HALRT-10001](#)".

HALRT-01303, Generic Critical Voltage Problem

Cause: There was a generic voltage problem.

Action: See "[HALRT-10001](#)".

HALRT-01304, Generic Critical Voltage Problem Cleared

Cause: The generic voltage problem was cleared.

Action: See "[HALRT-10001](#)".

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.

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.

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.

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.

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.

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.

Replacement Units

This appendix lists the replacement units for Recovery Appliance. It contains the following topics:

- [About Warranty Coverage](#)
- [Understanding Repair Categories](#)
- [LED Status Descriptions](#)
- [Parts for Compute Servers](#)
- [Parts for Storage Servers](#)
- [Parts for InfiniBand Switches](#)
- [Parts for Cisco Switch](#)
- [Parts for Power Distribution Units](#)
- [Cables and Cable Accessories for Recovery Appliance](#)

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.

Understanding Repair Categories

Repair categories describe the potential impact of a repair action. [Table 26-1](#) describes the categories and actions.

Table 26-1 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.	• 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.	• Ethernet switch • External cables • InfiniBand switch

Table 26-1 (Cont.) Recovery Appliance Repair Categories

Category	Description	Examples
Compute Server Offline (DBO)	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. If the server is currently running, then you must shut it down. See "Powering Down the Servers". 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.	• Memory • PCIe cards • PDUs • Processors • System boards
Storage server Offline (EBO)	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. If the server is currently running, then you must shut it down. See "Powering Down the Servers". 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.	• Memory • PCIe cards • Processors • System boards

LED Status Descriptions

This section describes the LEDs.

Oracle Server LEDs

[Table 26-2](#) identifies the LED status indicators that can indicate faults on the compute servers and the storage servers.

Table 26-2 Oracle Server LED Status Descriptions

Component	LED Status
General Status	The following LEDs are located on the front and back panels of each server: Service Required (amber) Off indicates normal operation. Steady on indicates a fault in the server. Power/OK (green) Steady on indicates that the server is powered on. A slow blink indicates that the server is starting up. A steady blink indicates that standby power is on, but chassis power is off. Off indicates no AC power or an incomplete startup of Oracle ILOM.
Fans	Each fan module has one bicolored LED, which you can see when the top cover fan door is open. 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. Green indicates that the system is powered on, and the fan module is functioning correctly.
Drives	The following LEDs are located on the storage and boot drives: Service Required (amber) Off indicates normal operation. Steady on indicates a fault in the drive. OK/Activity (green) Steady on indicates normal inactivity. A steady blink indicates normal activity. Off indicates that the power is off, or the system does not recognize the drive.

Table 26-2 (Cont.) Oracle Server LED Status Descriptions

Component	LED Status
Power supply	The following LEDs are on each power supply. You can see them from the back of the server. Service Required (amber) Off indicates normal operation. Steady on indicates that the system has detected a problem in the power supply regarding the voltage, current, temperature, or fans. AC OK/DC OK (green) Steady on indicates normal operation for both AC and DC power. A slow blink indicates that input AC power is normal and output DC power is disabled. Off indicates no input AC power.

**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

Sun Datacenter InfiniBand Switch 36 LEDs

Table 26-3 identifies the LED status indicators that can indicate faults on Sun Datacenter InfiniBand Switch 36.

Table 26-3 Sun InfiniBand Switch LED Status Descriptions

Component	LED Status
Chassis	The following LEDs are located on the left side of the rear panel. Attention (amber) Steady on indicates a fault. Off indicates normal operation. Flashing indicates no function. OK (green) Steady on indicates normal operation. Off indicates that the switch is off or initializing. Flashing indicates no function.
Link status	The link status LEDs are located by the data cable connectors on the rear panel. Link (green) Steady on indicates that a link is established. Off indicates that the link is down. Flashing indicates that there are symbol errors.
Network management ports	The following LEDs are located on the network management connector on the rear panel. Link Speed (amber/green): Green steady on indicates a 1000BASE-T link. Amber steady on indicates a 100BASE-T link. Off indicates no link. Flashing indicates no function. Activity (green): Flashing indicates packet activity. On indicates no function. Off indicates no activity.

Table 26-3 (Cont.) Sun InfiniBand Switch LED Status Descriptions

Component	LED Status
Power supply	The following LEDs are located on the power supply at the front of the chassis. Attention (amber) Steady on indicates a fault, and 12 VDC is shut down. Off indicates normal operation. Flashing indicates no function. OK (green) Steady on indicates 12 VDC. Off indicates no DC voltage. Flashing indicates no function. AC LED is green Steady on indicates AC power is supplied. Off indicates that AC power is not present. Flashing indicates no function.
Fan	The following LED is located on the air inlet on the front of the chassis. FT On indicates a problem with the fan. Off indicates normal operation

**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

Parts for Compute Servers

The following topics identify the replacement parts for the Recovery Appliance compute servers:

- [Database Server Components of Recovery Appliance X9M](#)
- [Database Server Components of Recovery Appliance X8M](#)
- [Oracle Server X8 and X8M Database Server Parts](#)

Database Server Components of Recovery Appliance X9M

Each Recovery Appliance X9M database server includes the following components:

- CPU - 32-core Intel Xeon 8358 processors (2.6 GHz):
 - Flexible configuration: 2 CPUs
 - Eighth Rack: 1 CPU
- RAM:
 - Flexible configuration: 512 GB RAM (16 x 32 GB DIMMs), expandable to 1 TB (16 x 64 GB DIMMs), 1.5 TB (16 x 32 GB and 16 x 64 GB DIMMs), or 2 TB (32 x 64 GB DIMMs) with memory expansion kit
 - Eighth Rack: 384 GB RAM (12 x 32 GB DIMMs), expandable to 1 TB (16 x 64 GB DIMMs) with memory expansion kit
- System storage: 2 x 3.84 TB PCIe 4.0 NVMe storage drives, expandable to 4 drives
- RDMA Network Fabric: 1 x dual-port CX5 RDMA Network Fabric card; PCIe 4.0, 2 x 100 Gb/s ports active-active (200 Gb/s combined throughput)
- Administration network: 1 x 1 GbE Base-T Ethernet port
- Integrated Lights Out Manager (ILOM) port for remote management: 1 x 1 GbE Base-T Ethernet port
- Client and additional network connectivity:
 - Flexible configuration options:
 - * Two factory-installed network cards. Each card contains 2 x SFP+/SFP28 ports (10/25 GbE) or 4 x 10GBASE-T ports (10 GbE)
 - * One optional field-installed network card containing 2 x SFP+/SFP28 ports (10/25 GbE) or 4 x 10GBASE-T ports (10 GbE) or 2 x QSFP28 ports (100GbE). Each QSFP28 port can support one 10 GbE or 25 GbE link with appropriate breakout cables.
 - Eighth Rack: Two factory-installed network cards only. Each card contains 2 x SFP+/SFP28 ports (10/25 GbE) or 4 x 10GBASE-T ports (10 GbE)

Database Server Components of Recovery Appliance X8M

Recovery Appliance X8M database servers include the following components:

- 2 x 24-core Intel Xeon Platinum 8260 Processors (2.4 GHz)
- 384 GB RAM (12 x 32 GB DIMMs), expandable to 1.5 TB (24 x 64 GB DIMMs) with memory expansion kit
- 4 x 1.2 TB boot drive/hard disks, hot swappable, expandable to 8 drives
- Disk controller HBA with 2 GB cache (no batteries)
- RDMA Network Fabric: 1 x dual-port QSFP28 100 Gb/s PCIe 3.0 RDMA Network Fabric card - both ports active
- Client network: 2 x SFP+/SFP28 ports (10/25 GbE) or 2 x 10GBASE-T ports (10 GbE)
- Backup/optional networks: 2 x SFP+/SFP28 ports (10/25 GbE) or 4 x 10GBASE-T ports (10 GbE)
- Administration network: 1 x 1 GbE Base-T Ethernet port
- Integrated Lights Out Manager (ILOM) port for remote management: 1 x 1 GbE Base-T Ethernet port

Oracle Server X8 and X8M Database Server Parts


The following image shows the front view of Oracle Server X8 and X8M Database Servers.

Figure 26-1 Front View of Oracle Server X8 and X8M Database Servers



The following image shows the rear view of the Oracle Server.

Figure 26-2 Rear View of Oracle Server X8 and X8M Database Servers



The following table lists the replaceable parts for Oracle Server X8 and X8M database servers. 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 26-4 Replacement Parts for Oracle Server X8 and X8M Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
370-5621-01	3 volt lithium battery grover	DBO	FRU	See "Servicing the Battery (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqviw.html#scrolltoc
7363537	ASSY DRV 1.2TB 10KRPM SAS3 512N SFF MARLIN	DBO	FRU	See "Servicing Storage Drives (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gquak.html#scrolltoc
7322171	ASM LFIM	DBO	FRU	See "Servicing Other Cables (FRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvev.html#scrolltoc
7339763	BD LP DUAL SFP28 ADAPTER WHITNEY+ PCIe3.0x8	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvjk.html#scrolltoc

Table 26-4 (Cont.) Replacement Parts for Oracle Server X8 and X8M Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
8200663	HEATSINK CPU HEIGHT:23.5 MM	DBO	FRU	See "Servicing Processors (FRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqtmk.html#scrolltoc
7343599	Assembly, riser, 1-slot, 1U	DBO	FRU	See "Servicing PCIe Risers (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvft.html#scrolltoc
7330698	DIMM, 32 GB, RDIMM, DDR4, 2666 2Rx4 1.2V	DBO	CRU	See "Servicing the DIMMs (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at https://docs.oracle.com/en/servers/x86/x8-2/exadata-db-service-manual/gqvkr.html
8200116	Power supply, AC, A266, F, 12V, 1200W	HS	CRU	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 Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/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 Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/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 Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvmq.html#scrolltoc
7340897	Cable, SAS3 to HBA	DBO	CRU	See "Servicing the Internal HBA Cables" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvav.html#scrolltoc
7331123	Assembly, fan module	HS	CRU	See "Servicing Fan Modules (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gquhg.html#scrolltoc
8200428	CPU INTEL CLX P8260 (B1) 2.4 GHz 165W 24c 8S 1TB (S-RF9H)	DBO	FRU	See "Servicing Processors (FRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqtmk.html#scrolltoc

Table 26-4 (Cont.) Replacement Parts for Oracle Server X8 and X8M Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7104073	Assembly, Dual 40Gb/Sec (4x) QDR RDMA Fabric Host Channel Adapter Module M3	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvjk.html#scrolltoc 3. See "Configuring the InfiniBand Switches"
7046442 (X8)	Assembly, Dual port, QDR CX3, HCA	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvjk.html#scrolltoc 3. See "Configuring the InfiniBand Switches"
7119792 (X8M)	CX5 100GbE Adapter Dual Port QSFP28	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvjk.html#scrolltoc 3. See Verifying the RoCE Network Fabric Configuration.
7360829	Assembly, motherboard, 1U	DBO	FRU	1. See "Servicing the Motherboard Assembly (FRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvfa.html#scrolltoc 2. Press F2 after restarting the server. 3. Set the boot order.
7331100	Assembly, riser, 2-slot, 1U	DBO	FRU	See "Servicing PCIe Risers (CRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvft.html#scrolltoc

Table 26-4 (Cont.) Replacement Parts for Oracle Server X8 and X8M Database Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7331142	Assembly, Front LED indicator module	DBO	FRU	See "Servicing the Front LED Indicator Module (FRU)" in <i>Oracle Exadata X8-2 Database Server Service Manual</i> at http://docs.oracle.com/cd/E93359_01/html/E93386/gqvok.html#scrolltoc

Parts for Storage Servers

The following topics identify the replacement parts for the Recovery Appliance storage servers:

- [Oracle Exadata Storage Server X9M, X8M, and X8 High Capacity and Extended \(XT\) Server Parts](#)

Oracle Exadata Storage Server X9M, X8M, and X8 High Capacity and Extended (XT) Server Parts

The following figure shows the front view of the Oracle Exadata Storage Server X9M, X8M and X8 High Capacity and XT servers.



Note:

The Recovery Appliance supports Oracle Exadata Storage Server high capacity servers and does not use Oracle Exadata Storage Server XT servers. Both are discussed for completeness and compatibility with Oracle Exadata documentation.

Figure 26-3 Front View of Oracle Exadata Storage Server X9M, X8M and X8 High Capacity and XT servers




The following figure shows the rear view of the Oracle Exadata Storage Server X8M and X8 High Capacity and XT servers.



Note:

The Recovery Appliance does not support XT servers.

Figure 26-4 Rear View of Oracle Exadata Storage Server X9M, X8M and X8 High Capacity and XT servers

The following table lists the replaceable parts for Oracle Exadata Storage Server X9M, X8M and X8 High Capacity and XT servers. 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.

**Note:**

The Recovery Appliance supports Oracle Exadata Storage Server high capacity servers and does not use Oracle Exadata Storage Server XT servers. Both are discussed for completeness and compatibility with Oracle Exadata.

Table 26-5 Replacement Parts for Oracle Exadata Storage Server X8M and X8 High Capacity and XT Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
370-5621-01	3 volt lithium battery grover	EBO	FRU	See "Servicing the Battery" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtff.html#scrolltoc
7322171	Assembly, Front LED Indicator Module (LFIM)	EBO	FRU	See "Servicing the Front LED Indicator Module" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtpg.html#scrolltoc
7341141	Assembly, 12-slot disk backplane	EBO	FRU	See "Servicing the Disk Backplane" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtsb.html#scrolltoc
8200116	Power Supply, A266 F, 1200W, AC Input	HS	FRU	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 X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtag.html#scrolltoc 3. Replace the CMA.

Table 26-5 (Cont.) Replacement Parts for Oracle Exadata Storage Server X8M and X8 High Capacity and XT Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7322166	Fan assembly, 2U	HS	FRU	See "Servicing Fan Modules" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/ggtak.html#scrolltoc
7046442 (X8)	Assembly, dual port, QDR CX3, Host Channel Adapter	EBO	FRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/ggtcq.html#scrolltoc 3. See "Configuring the InfiniBand Switches" 4. Verify the new hardware meets the hardware and firmware specifications using the following command: <code>/opt/oracle.SupportTools/CheckHWnFWprofile</code>
7119792 (X8M-8)	CX5 100GbE Adapter Dual Port QSFP28	DBO	CRU	1. If the server is currently running, then it should be shut down gracefully. See "Stopping Recovery Appliance Services" for additional information. 2. See "Servicing PCIe Cards" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/ggtcq.html#scrolltoc 3. See "Configuring the InfiniBand Switches" 4. Verify the new hardware meets the hardware and firmware specifications using the following command: <code>/opt/oracle.SupportTools/CheckHWnFWprofile</code>
7322160	Heat sink, CPU	EBO	FRU	See "Servicing Processors" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/ggtmk.html#scrolltoc
7357761	Assembly, Drive, 14TB 3.5" SAS3/7200 rpm Coral with encryption	HS	FRU	See "Servicing Storage Drives" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/ggtcs.html#scrolltoc

Table 26-5 (Cont.) Replacement Parts for Oracle Exadata Storage Server X8M and X8 High Capacity and XT Servers

Part Number	Description	Repair Category	Coverage under Basic Warranty	Repair Procedure
7353916	RDIMM 16GB, DDR4 2666, 1Rx4 1.2V	EBO	FRU	See "Servicing the DIMMs" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at https://docs.oracle.com/en/servers/x86/x8-2l/exadata-storage-service-manual/gqtcn.html
7332895	HBA, 12 Gbps SAS PCIe 16-port internal RAID with flash module	EBO	FRU	See "Servicing the Internal HBA Card and HBA Super Capacitor" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtjv.html#scrolltoc
7086345	Assembly, super cap, 13.5V, 6.4F, SAS3 HBA	EBO	FRU	See "Servicing the Internal HBA Card and HBA Super Capacitor" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtjv.html#scrolltoc
7315464	SAS-3 Data Cable Assembly	EBO	FRU	See "Servicing SAS Cables" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqthb.html#scrolltoc
7361253	M.2,240GB,SATA,22x80 mm,SSD,INTEL, DC S4510	EBO	FRU	See "Servicing the Internal M.2 Flash SSDs" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/grlly.html#scrolltoc
7335943	6.4TB FLASH F640 NVMe AIC AURA 7	EBO	FRU	See "Servicing PCIe Cards" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtcq.html#scrolltoc
7361454	6.4TB FLASH F640 V2 NVMe AIC AURA 8	EBO	FRU	See "Servicing PCIe Cards" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtcq.html#scrolltoc
8200426	CPU INTEL CLX G5218 (B1) 2.3 GHz 125W 16c 4S 1TB (S-RF8T)	EBO	FRU	See "Servicing Processors" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtmk.html#scrolltoc
7360978	Assembly, System board	EBO	FRU	See "Servicing the Motherboard Assembly" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/gqtwc.html#scrolltoc
7341466	Assembly, M.2 form factor riser	EBO	FRU	See "Servicing the Internal M.2 Flash SSDs" in <i>Oracle Exadata Storage Server X8-2 EF, HC, and XT Service Manual</i> at http://docs.oracle.com/cd/E93361_01/html/E93395/grlly.html#scrolltoc

Parts for Cisco Switch

Table 26-6 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 26-6 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

Parts for InfiniBand Switches

Table 26-7 lists the replaceable parts for the InfiniBand switches used in Recovery Appliance. See "[Cables and Cable Accessories for Recovery Appliance](#)" 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 26-7 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

Parts for Power Distribution Units

Table 26-8 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 26-8 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" 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" 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" of this guide.
371-5039	Single-phase PDU: 22k VA Hubbell CS8265C PlugNo X-Option	IR	See "Replacing a Power Distribution Unit" 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" 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" of this guide.

Cables and Cable Accessories for Recovery Appliance

[Table 26-9](#) 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 26-9 Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
350-1287	Slim rail, CMA, 1U-2U for a storage server	EBO	1. If the server is currently running, then it should be shut down gracefully. See "Powering On and Off a Recovery Appliance Rack" 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	1. If the server is currently running, then it should be shut down gracefully. See "Powering On and Off a Recovery Appliance Rack" 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: • 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
350-1638	Label assembly EV2, 3-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1639	Label assembly EV2, 4-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1640	Label assembly EV2, 5-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1641	Label assembly EV2, 6-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1642	Label assembly EV2, 7-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
350-1643	Label assembly EV2, 8-way, service upgrade, label	HS	See In-Rack Cabling Tables for Recovery Appliance X4
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.

Table 26-9 (Cont.) Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
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
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

Table 26-9 (Cont.) Replacement Cables for Recovery Appliance

Part Number	Description	Repair Category	Repair
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" of this guide.
7042273	Slim rail, cable management arm, 1U-2U for a storage server	EBO	1. If the server is currently running, then shut it down. See "Shutting Down Recovery Appliance". 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	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 Note: Purchase two QSFP SR Transceivers (Part Number X2124A-N) for each passive fiber cable.
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.

Index

A

access route requirements, [2-3](#)

add

 racli Command, [15-7](#)

add admin_user

 racli Command, [15-7](#)

add certificate

 racli Command, [15-7](#)

add cloud_key

 racli Command, [15-8](#)

add cloud_location

 racli Command, [15-9](#)

add cloud_user

 racli Command, [15-10](#)

add credential

 racli Command, [15-11](#)

add credential_wallet

 racli Command, [15-12](#)

add db_user

 racli Command, [15-12](#)

add host

 racli Command, [15-13](#)

add keystore

 racli Command, [15-13](#)

add nfs

 racli Command, [15-14](#)

add rasys_wallet

 racli Command, [15-14](#)

add remote_syslog

 racli Command, [15-15](#)

add storage_cells

 racli Command, [15-15](#)

add storage_disks

 racli Command, [15-16](#)

add tag

 racli Command, [15-16](#)

add vpc_user

 racli Command, [15-17](#)

airflow, [2-12](#), [2-14](#)

alert messages

 battery learn cycle, [13-33](#)

 HALRT-00002 to HALRT-00500, [25-2](#)

 HALRT-00513 to HALRT-00600, [25-5](#)

 HALRT-00616 to HALRT-00700, [25-6](#)

 HALRT-00717 to HALRT-00800, [25-6](#)

alert messages (*continued*)

 HALRT-00912 to HALRT-01000, [25-7](#)

 HALRT-01015 to HALRT-01100, [25-7](#)

 HALRT-01209 to HALRT-01300, [25-27](#)

 HALRT-01301 to HALRT-01400, [25-30](#)

 HALRT-02001 to HALRT-02008, [25-33](#)

 HALRT-02101 to HALRT-02200, [25-34](#)

 HALRT-10001 to HALRT-10500, [25-37](#)

 HALRT-10503 to HALRT-11000, [25-42](#)

 HALRT-11001 to HALRT-11100, [25-43](#)

 hardware, [25-1](#)

ALLOW_BACKUP_DELETION, [10-5](#)

alter

 racli Command, [15-18](#)

alter admin_user

 racli Command, [15-18](#)

alter adr_parameter

 racli Command, [15-19](#)

alter cloud_key

 racli Command, [15-19](#)

alter cloud_location

 racli Command, [15-20](#)

alter db_user

 racli Command, [15-20](#)

alter keystore

 racli Command, [15-21](#)

alter network

 racli Command, [15-22](#)

alter parameter

 racli Command, [15-23](#)

alter passphrase

 racli Command, [15-23](#)

alter rasys_user

 racli Command, [15-24](#)

alter vpc_user

 racli Command, [15-25](#)

ambient relative humidity range, [2-12](#)

ambient temperature ranges, [2-12](#)

approve

 racli Command, [15-25](#)

approve request

 racli Command, [15-25](#)

ASR,

 about, [4-1](#)

 configuration, [4-2](#)

 troubleshooting, [4-5](#)

ASR Manager
 installation, [4-3](#)
 ASR network connections, [4-2](#)
 ASR trap destinations, [4-4](#)
 Auto Service Request, [4-1](#)
See also ASR

B

batteries
 learn cycle, [13-33](#)
 life expectancy, [13-8](#)
 monitoring, [13-9](#)
 replacing, [13-9](#)
 replacing disk controller BBUs, [13-46](#)
 bend radius requirements, [7-5](#)

C

cable management arm, [7-5](#)
 cable management arms
 replacement part, [26-17](#)
 cables
 color coding, [16-1](#), [17-1](#), [18-1](#), [19-1](#), [20-1](#)
 replacement, [26-17](#)
 cabling
 bend radius, [7-5](#)
 bundling InfiniBand, [7-5](#)
 estimating cable length, [7-2](#)
 InfiniBand precautions, [7-5](#)
 new storage servers, [7-18](#)
 CELLBOOT USB flash drive, [13-48](#)
 channel adapter, [1-5–1-8](#)
 checking
 cell configuration, [7-2](#)
 cell disk configuration, [7-2](#)
 firmware version, [7-3](#)
 flash logs, [7-2](#)
 HugePages memory configuration, [7-3](#)
 image history, [7-2](#)
 InfiniBand switch information, [7-3](#)
 IORM plans, [7-2](#)
 IP addresses, [7-2](#)
 checkip script, [9-2](#)
 chrony, [14-8](#)
 ciphers, [11-1](#)
 Cisco switch
 configuring default gateway, [8-14](#), [8-16](#), [8-21](#)
 configuring host name, [8-14](#), [8-21](#)
 default port speed, [8-14](#), [8-15](#)
 model, [1-5–1-8](#)
 replacement parts, [26-15](#)
 serial number, [11-4](#), [11-5](#)
 setting password, [8-14](#), [8-21](#), [14-4](#)
 Cisco switches
 model, [1-5](#), [1-6](#)

cloning Oracle Database homes
 Linux-based servers, [13-26](#)
 cloning Oracle Grid infrastructure
 Linux-based servers, [13-24](#)
 Cloud Control discovery, [9-5](#)
 cluster verification utility, [13-18](#)
 CMA guidelines, [7-5](#)
 compliance
 DBMS_RA, [10-5](#)
 racli add db_user, [10-5](#)
 Compliance
 Quorum, [10-9](#)
 Compliance Configuration, [10-5](#)
 compute servers
 components, [1-8](#)
 model, [1-5–1-8](#)
 configuration, [5-1](#)
 configure
 racli Command, [15-26](#)
 configure export
 racli Command, [15-26](#)
 configuring
 InfiniBand interface, [8-4](#)
 PDUs, [8-26](#)
 configuring ASR, [4-2](#)
 create
 racli Command, [15-26](#)
 create certificate
 racli Command, [15-26](#)
 CSI, [7-8](#)
 Customer Support Identifier, [7-8](#)

D

default IP addresses, [3-13](#)
 default passwords, [10-1](#)
 default serial port speed
 Cisco switch, [8-14](#), [8-15](#)
 deny
 racli Command, [15-27](#)
 deny request
 racli Command, [15-27](#)
 Deployment Assistant, [5-1](#)
 disable
 racli Command, [15-27](#)
 disable aide
 racli Command, [15-28](#)
 disable archived_log_grouping
 racli Command, [15-28](#)
 disable check
 racli Command, [15-28](#)
 disable cloud_location
 racli Command, [15-29](#)
 disable keystore
 racli Command, [15-29](#)

- disable nfs
 - racli Command, [15-29](#)
- disable rasys_user
 - racli Command, [15-30](#)
- disable root_access
 - racli Command, [15-30](#)
- disable rpc
 - racli Command, [15-30](#)
- disable ssh
 - racli Command, [15-31](#)
- disable sys_remote_access
 - racli Command, [15-31](#)
- disk controller battery BBUs, [13-46](#)
- disk controllers, [13-16](#)
- disk drives
 - server status indicators, [26-3](#)
- disk failures
 - replacing physical disks, [13-34](#)
- DNS, [3-13](#)
- domain name system, [3-13](#)

E

- electrostatic discharge, [2-12](#)
- email accounts, My Oracle Support (ASR), [4-1](#)
- enable
 - racli Command, [15-31](#)
- enable acfs
 - racli Command, [15-32](#)
- enable aide
 - racli Command, [15-32](#)
- enable archived_log_grouping
 - racli Command, [15-32](#)
- enable check
 - racli Command, [15-33](#)
- enable cloud_location
 - racli Command, [15-33](#)
- enable keystore
 - racli Command, [15-33](#)
- enable nfs
 - racli Command, [15-34](#)
- enable rasys_user
 - racli Command, [15-34](#)
- enable rpc
 - racli Command, [15-34](#)
- enable ssh
 - racli Command, [15-35](#)
- enable sys_remote_access
 - racli Command, [15-35](#)
- error messages
 - contacting Oracle Support Services, [25-1](#)
- Ethernet cabling connections, [16-3](#), [17-6](#), [18-4](#), [19-4](#), [20-4](#)
- Ethernet switch
 - model, [1-5–1-8](#)
 - replacement part, [26-15](#)

- exachk utility, [7-2](#)
- Exadata software installation, [9-3](#)
- export
 - racli Command, [15-35](#)
- export appliance
 - racli Command, [15-35](#)

F

- fan modules
 - server status indicators, [26-3](#)
 - switch status indicators, [26-6](#)
- FDOM, [13-40](#)
- fiber channel adapters, [3-10](#)
- FIPS 140-2, [11-1](#)
- firewalls
 - port assignments, [3-14](#)
- firmware levels, [13-35](#), [13-36](#)
- flash disks, [13-40](#)

G

- getmaster command, [8-35](#)

H

- HALRT-00002 to HALRT-00500 alert messages, [25-2](#)
- HALRT-00513 to HALRT-00600 alert messages, [25-5](#)
- HALRT-00616 to HALRT-00700 alert messages, [25-6](#)
- HALRT-00717 to HALRT-00800 alert messages, [25-6](#)
- HALRT-00912 to HALRT-01000 alert messages, [25-7](#)
- HALRT-01015 to HALRT-01100 alert messages, [25-7](#)
- HALRT-01209 to HALRT-01300 alert messages, [25-27](#)
- HALRT-01301 to HALRT-01400 alert messages, [25-30](#)
- HALRT-02001 to HALRT-02008 alert messages, [25-33](#)
- HALRT-02101 to HALRT-02200 alert messages, [25-34](#)
- HALRT-10001 to HALRT-10500 alert messages, [25-37](#)
- HALRT-10503 to HALRT-11000 alert messages, [25-42](#)
- HALRT-11001 to HALRT-11100 alert messages, [25-43](#)
- hardware alert messages, [25-1](#)
- hardware components, [1-4](#)
- hardware failures, automatic reporting, [4-1](#)
- hardware modifications, [1-19](#)

hardware replacement
 X8 High Capacity, [26-11](#)
 X8M High Capacity, [26-11](#)
 humidity range, [2-12](#)

I

identifying
 predictive failure flash disk, [13-42](#)
 ILOM cabling connections, [16-2](#), [17-6](#), [18-3](#), [19-3](#),
 [20-3](#)
 InfiniBand
 switch time zone, [14-9](#)
 InfiniBand cabling, [19-9](#), [20-7](#)
 InfiniBand switch cabling, [20-7](#)
 InfiniBand switches
 LEDs, [26-4](#)
 model, [1-6–1-8](#)
 ingest network, [3-2](#)
 install
 racli Command, [15-36](#)
 install appliance
 racli Command, [15-36](#)
 install okv_endpoint
 racli Command, [15-37](#)
 install utility
 step descriptions, [9-8](#)
 syntax, [9-6](#)
 installation
 safety guidelines, [7-4](#)
 installation directory, [9-2](#)
 installing software, [9-3](#)
 IP addresses
 factory default IP addresses, [3-13](#)
 ipconf utility, [13-49](#)
 IPMItool
 resetting the Service Processor, [13-16](#)

K

KEEP_COMPLIANCE, [10-5](#)
 Kerberos Authentication, [10-6](#)

L

LDAP Authentication, [10-7](#)
 leaf switches, [8-8](#), [8-35](#)
 LED color codes, [26-2](#), [26-4](#)
 LED status, [26-2](#)
 license restrictions, [1-19](#)
 list
 racli Command, [15-38](#)
 list admin_user
 racli Command, [15-38](#)
 list certificate
 racli Command, [15-38](#)

list check
 racli Command, [15-38](#)
 list cloud_key
 racli Command, [15-39](#)
 list cloud_location
 racli Command, [15-39](#)
 list cloud_user
 racli Command, [15-40](#)
 list compute_nodes
 racli Command, [15-40](#)
 list db_user
 racli Command, [15-41](#)
 list env
 racli Command, [15-41](#)
 list export
 racli Command, [15-41](#)
 list grid_home
 racli Command, [15-42](#)
 list host
 racli Command, [15-42](#)
 list ip_addr
 racli Command, [15-42](#)
 list listener
 racli Command, [15-43](#)
 list nfs
 racli Command, [15-43](#)
 list okv_endpoint
 racli Command, [15-43](#)
 list oracle_home
 racli Command, [15-44](#)
 list parameter
 racli Command, [15-44](#)
 list remote_syslog
 racli Command, [15-45](#)
 list san
 racli Command, [15-45](#)
 list scan
 racli Command, [15-45](#)
 list scan_listener
 racli Command, [15-45](#)
 list storage_cells
 racli Command, [15-46](#)
 list tag
 racli Command, [15-46](#)
 listing PCI slot numbers, [13-40](#)

M

media managers
 installation of third-party, [1-19](#)
 model numbers, [1-4](#)
 moving drives, [13-37](#)
 My Oracle Support email account, [4-1](#)

N

net weights
 Oracle Exadata Storage Expansion Rack, [2-3](#)
 Recovery Appliance, [2-3](#)
 network components, [3-4](#), [3-6](#)
 network connections
 ASR, [4-2](#)
 network diagram, [3-8](#)
 network preparation, [3-1](#)
 Network Time Protocol, [14-8](#)
 networks
 ingest, [3-2](#)
 RDMA Network Fabric, [3-2](#)
 NTP servers, [14-8](#)

O

Oracle Database security products, [11-6](#)
 Oracle Enterprise Manager Cloud Control
 discovery, [9-5](#)
 Oracle Enterprise Manager installation, [1-19](#)
 Oracle Exadata Deployment Assistant, [5-1](#)
 Oracle Exadata Deployment Assistant (OEDA),
 [10-10](#)
 Oracle Management Server, [9-2](#)
 Oracle Maximum Availability Architecture, [13-30](#)
 Oracle Secure Backup, [1-19](#)
 Oracle Support Services
 contacting, [25-1](#)

P

password policies, [10-10](#)
 passwords
 changing
 Cisco Ethernet password, [14-4](#)
 database server passwords, [14-1](#)
 ILOM password, [14-3](#)
 InfiniBand switch passwords, [14-3](#)
 PDU passwords, [14-2](#)
 storage server passwords, [14-2](#)
 PCI slot numbers, [13-40](#)
 PCIe cards, [13-40](#)
 PDU cabling, [20-5](#), [20-6](#)
 RA21, [16-5](#), [16-6](#), [17-8](#), [17-9](#)
 single-phase, [16-5](#), [17-8](#), [18-5](#), [19-5](#), [19-6](#)
 three-phase, [16-6](#), [17-9](#), [18-6](#), [19-7](#), [19-8](#)
 X7 and earlier, [19-6](#), [19-8](#)
 X8, [19-5](#), [19-7](#)
 X8M, [18-5](#), [18-6](#)
 PDUs
 configuration, [8-26](#)
 PDU-A location, [13-10](#)
 PDU-B location, [13-10](#)
 replacement parts, [26-16](#)

perforated tiles, [2-14](#)
 physical disks
 removing in error, [13-39](#)
 replacing, [13-34](#)
 status, [13-30](#)
 port assignments, [3-14](#)
 power off
 emergency, [13-2](#)
 switch, [13-2](#)
 power on
 sequence for servers, [13-6](#)
 using ILOM, [13-6](#)
 power supplies
 server status indicators, [26-4](#)
 switch status indicators, [26-6](#)
 problem reporting, [4-1](#)

Q

Quorum
 Compliance, [10-9](#)

R

rack
 serial number, [11-1](#)
 rack components, [1-4](#)
 serial numbers, [11-2](#)
 rack configuration, [5-1](#)
 rack layout
 RA21 and X8M full rack, [18-1](#)
 RA21 and X8M minimum configuration, [18-2](#)
 RA21 full rack, [17-2](#)
 RA21 minimum configuration, [17-4](#)
 X4 full rack, [20-1](#)
 X4 minimum configuration, [20-2](#)
 X7, X6, or X5 full rack, [19-1](#)
 X8, X7, X6, or X5 minimum configuration,
 [19-2](#)
 racli Commands
 add, [15-7](#)
 add admin_user, [15-7](#)
 add certificate, [15-7](#)
 add cloud_key, [15-8](#)
 add cloud_location, [15-9](#)
 add cloud_user, [15-10](#)
 add credential, [15-11](#)
 add credential_wallet, [15-12](#)
 add db_user, [15-12](#)
 add host, [15-13](#)
 add keystore, [15-13](#)
 add nfs, [15-14](#)
 add rasys_wallet, [15-14](#)
 add remote_syslog, [15-15](#)
 add storage_cells, [15-15](#)
 add storage_disks, [15-16](#)

racli Commands (*continued*)

add tag, [15-16](#)
 add vpc_user, [15-17](#)
 alter, [15-18](#)
 alter admin_user, [15-18](#)
 alter adr_parameter, [15-19](#)
 alter cloud_key, [15-19](#)
 alter cloud_location, [15-20](#)
 alter db_user, [15-20](#)
 alter keystore, [15-21](#)
 alter network, [15-22](#)
 alter parameter, [15-23](#)
 alter passphrase, [15-23](#)
 alter rasys_user, [15-24](#)
 alter vpc_user, [15-25](#)
 approve, [15-25](#)
 approve request, [15-25](#)
 configure, [15-26](#)
 configure export, [15-26](#)
 create, [15-26](#)
 create certificate, [15-26](#)
 deny, [15-27](#)
 deny request, [15-27](#)
 disable, [15-27](#)
 disable aide, [15-28](#)
 disable archived_log_grouping, [15-28](#)
 disable check, [15-28](#)
 disable cloud_location, [15-29](#)
 disable keystore, [15-29](#)
 disable nfs, [15-29](#)
 disable rasys_user, [15-30](#)
 disable root_access, [15-30](#)
 disable rpc, [15-30](#)
 disable ssh, [15-31](#)
 disable sys_remote_access, [15-31](#)
 enable, [15-31](#)
 enable acfs, [15-32](#)
 enable aide, [15-32](#)
 enable archived_log_grouping, [15-32](#)
 enable check, [15-33](#)
 enable cloud_location, [15-33](#)
 enable keystore, [15-33](#)
 enable nfs, [15-34](#)
 enable rasys_user, [15-34](#)
 enable rpc, [15-34](#)
 enable ssh, [15-35](#)
 enable sys_remote_access, [15-35](#)
 export, [15-35](#)
 export appliance, [15-35](#)
 install, [15-36](#)
 install appliance, [15-36](#)
 install okv_endpoint, [15-37](#)
 list, [15-38](#)
 list admin_user, [15-38](#)
 list certificate, [15-38](#)
 list check, [15-38](#)

racli Commands (*continued*)

list cloud_key, [15-39](#)
 list cloud_location, [15-39](#)
 list cloud_user, [15-40](#)
 list compute_nodes, [15-40](#)
 list db_user, [15-41](#)
 list env, [15-41](#)
 list export, [15-41](#)
 list grid_home, [15-42](#)
 list host, [15-42](#)
 list ip_addr, [15-42](#)
 list listener, [15-43](#)
 list nfs, [15-43](#)
 list okv_endpoint, [15-43](#)
 list oracle_home, [15-44](#)
 list parameter, [15-44](#)
 list remote_syslog, [15-45](#)
 list san, [15-45](#)
 list scan, [15-45](#)
 list scan_listener, [15-45](#)
 list storage_cells, [15-46](#)
 list tag, [15-46](#)
 remove, [15-46](#)
 remove acfs, [15-47](#)
 remove admin_user, [15-47](#)
 remove certificate, [15-47](#)
 remove cloud_key, [15-48](#)
 remove cloud_location, [15-48](#)
 remove cloud_user, [15-48](#)
 remove db_user, [15-49](#)
 remove grid_home, [15-49](#)
 remove host, [15-49](#)
 remove keystore, [15-50](#)
 remove nfs, [15-50](#)
 remove oracle_home, [15-50](#)
 remove osb_pieces, [15-51](#)
 remove remote_syslog, [15-51](#)
 remove tag, [15-51](#)
 remove vpc_user, [15-51](#)
 run, [15-52](#)
 run check, [15-52](#)
 run diagnostics, [15-53](#)
 run preinstall, [15-54](#)
 start, [15-55](#)
 start appliance, [15-55](#)
 start nfs, [15-55](#)
 status, [15-56](#)
 status aide, [15-56](#)
 status appliance, [15-56](#)
 status archived_log_grouping, [15-56](#)
 status cloud_location, [15-57](#)
 status compliance, [15-57](#)
 status disk_group, [15-57](#)
 status keystore, [15-58](#)
 status listener, [15-58](#)
 status nfs, [15-58](#)

- racli Commands (*continued*)
 - status okv_endpoint, [15-59](#)
 - status request, [15-59](#)
 - status scan, [15-59](#)
 - status scan_listener, [15-59](#)
 - stop, [15-60](#)
 - stop appliance, [15-60](#)
 - stop nfs, [15-60](#)
 - uninstall, [15-60](#)
 - uninstall appliance, [15-61](#)
 - uninstall okv_endpoint, [15-62](#)
 - update, [15-62](#)
 - update aide, [15-63](#)
 - update appliance, [15-63](#)
 - update parameter, [15-64](#)
 - update ssh_keys, [15-64](#)
 - update ssh_known_hosts, [15-64](#)
 - update storage_location, [15-64](#)
 - update support_tools, [15-65](#)
 - version, [15-65](#)
- RDMA Network Fabric
 - private network, [3-2](#)
- RDMA Network Fabric cabling, [16-7](#), [17-10](#), [18-7](#)
- RDMA Network Fabric switch cabling, [16-8](#), [17-11](#), [18-8](#), [19-10](#)
- Recovery Appliance software installation, [9-3](#)
- RECOVERY_WINDOW_COMPLIANCE, [10-5](#)
- recreating
 - CELLBOOT USB flash drive, [13-51](#)
- reenabling a rejected disk, [13-39](#)
- relative humidity, [2-12](#), [2-14](#)
- relocating voting disks
 - Linux-based servers, [13-19](#)
- remove
 - racli Command, [15-46](#)
- remove acfs
 - racli Command, [15-47](#)
- remove admin_user
 - racli Command, [15-47](#)
- remove certificate
 - racli Command, [15-47](#)
- remove cloud_key
 - racli Command, [15-48](#)
- remove cloud_location
 - racli Command, [15-48](#)
- remove cloud_user
 - racli Command, [15-48](#)
- remove db_user
 - racli Command, [15-49](#)
- remove grid_home
 - racli Command, [15-49](#)
- remove host
 - racli Command, [15-49](#)
- remove keystore
 - racli Command, [15-50](#)
- remove nfs
 - racli Command, [15-50](#)
- remove oracle_home
 - racli Command, [15-50](#)
- remove osb_pieces
 - racli Command, [15-51](#)
- remove remote_syslog
 - racli Command, [15-51](#)
- remove tag
 - racli Command, [15-51](#)
- remove vpc_user
 - racli Command, [15-51](#)
- removing a disk in error, [13-39](#)
- repair categories, [26-1](#)
- replacement
 - tape drive, [13-12](#)
- replacement parts, [26-1](#)
- replacing
 - disk controller BBUs, [13-46](#)
 - physical disks, [13-34](#)
- rescue flash drive, [13-48](#)
- resetting
 - ILOM, [13-15](#)
- run
 - racli Command, [15-52](#)
- run check
 - racli Command, [15-52](#)
- run diagnostics
 - racli Command, [15-53](#)
- run preinstall
 - racli Command, [15-54](#)

S

- safety guidelines, [2-1](#)
- SCAN addresses, [3-2](#), [3-13](#)
- Security, [10-5](#)
- serial numbers
 - Cisco switch, [11-4](#), [11-5](#)
 - rack components, [11-2](#)
- servers
 - LED color codes, [26-2](#), [26-4](#)
- setting
 - battery learn cycle, [13-33](#)
 - Cisco switch time zone, [8-14](#)
 - InfiniBand
 - switch date, [14-9](#)
 - switch time zone, [14-9](#)
 - Subnet Manager Master, [8-35](#)
- shipping weights
 - Oracle Exadata Storage Expansion Rack, [2-3](#)
 - Recovery Appliance, [2-3](#)
- SMTP alerts, [4-1](#)
- SNMP alerts, [4-1](#)
- software installation, [9-3](#)
- software modifications, [1-20](#)

- spine switches
 - rack location, [8-8](#), [8-35](#)
- SSH (Secure Shell), [11-1](#)
- SSH protocol 2 (SSH2), [11-1](#)
- start
 - racli Command, [15-55](#)
- start appliance
 - racli Command, [15-55](#)
- start nfs
 - racli Command, [15-55](#)
- status
 - component LEDs, [26-2](#)
 - physical disks, [13-30](#)
 - racli Command, [15-56](#)
- status aide
 - racli Command, [15-56](#)
- status appliance
 - racli Command, [15-56](#)
- status archived_log_grouping
 - racli Command, [15-56](#)
- status cloud_location
 - racli Command, [15-57](#)
- status compliance
 - racli Command, [15-57](#)
- status disk_group
 - racli Command, [15-57](#)
- status keystore
 - racli Command, [15-58](#)
- status listener
 - racli Command, [15-58](#)
- status nfs
 - racli Command, [15-58](#)
- status okv_endpoint
 - racli Command, [15-59](#)
- status request
 - racli Command, [15-59](#)
- status scan
 - racli Command, [15-59](#)
- status scan_listener
 - racli Command, [15-59](#)
- stop
 - racli Command, [15-60](#)
- stop appliance
 - racli Command, [15-60](#)
- stop nfs
 - racli Command, [15-60](#)
- storage servers
 - adding to a rack, [7-1](#)
 - components, [1-14](#)
 - increasing storage location size, [7-21](#)
 - model, [1-5–1-8](#)
- Subnet Manager Master, [8-8](#), [8-35](#)
- Sun Datacenter InfiniBand Switch 36
 - LEDs, [26-4](#)
 - replacement parts, [26-15](#)
- support accounts, [4-2](#)

- switch cabling, [16-8](#), [17-11](#), [18-8](#)
- switch-to-server cabling, [16-8](#), [17-11](#), [18-8](#), [19-10](#), [20-7](#)
- switch-to-switch cabling, [16-7](#), [17-10](#), [18-7](#), [19-9](#), [20-7](#)

T

- tape backup
 - network connections, [3-10](#)
 - Oracle red stack, [1-18](#)
 - third-party restrictions, [1-18](#)
- tape drive
 - replacement, [13-12](#)
- temperature ranges, [2-12](#)
- time zones, [8-14](#)
- trap destinations (ASR), [4-4](#)

U

- uninstall
 - racli Command, [15-60](#)
- uninstall appliance
 - racli Command, [15-61](#)
- uninstall okv_endpoint
 - racli Command, [15-62](#)
- update
 - racli Command, [15-62](#)
- update aide
 - racli Command, [15-63](#)
- update appliance
 - racli Command, [15-63](#)
- update parameter
 - racli Command, [15-64](#)
- update ssh_keys
 - racli Command, [15-64](#)
- update ssh_known_hosts
 - racli Command, [15-64](#)
- update storage_location
 - racli Command, [15-64](#)
- update support_tools
 - racli Command, [15-65](#)
- USB devices, [1-20](#)
- USB flash drive, [13-48](#)
- utilities
 - cluster verification, [13-18](#)
 - exachk, [7-2](#)
 - ipconf, [13-49](#)
 - verify-topology, [8-33](#)

V

- verify-topology utility, [8-33](#)
- verifying
 - database server configuration, [13-16](#)
 - firmware, [13-35](#), [13-36](#), [26-8](#), [26-11](#)

verifying (*continued*)
 hardware, [26-8](#), [26-11](#)
 InfiniBand Network Fabric topology, [8-33](#)
 virtual drive configuration, [13-17](#)
version
 racli Command, [15-65](#)
VIP addresses, [3-2](#), [3-13](#)
virtual local area networks, [3-8](#)

VLAN tagging, [3-12](#)
VLANs, [3-8](#), [8-14](#)

W

warranty restrictions, [1-19](#)
WorkDir subdirectory, [9-2](#)