

Oracle® Big Data Appliance

Owner's Guide

Release 1 (1.0.3)

E25960-03

April 2012

Oracle Big Data Appliance Owner's Guide, Release 1 (1.0.3)

E25960-03

Copyright © 2012, Oracle and/or its affiliates. All rights reserved.

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

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

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

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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

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

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

This software or hardware and documentation may provide access to or information on 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. 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.

Cloudera, Cloudera CDH, and Cloudera Manager are registered and unregistered trademarks of Cloudera, Inc.

Contents

Preface	xiii
Audience	xiii
Documentation Accessibility	xiii
Related Documentation	xiii
Conventions	xiv
Backus-Naur Form Syntax	xiv
1 Introduction to Oracle Big Data Appliance	
About Oracle Big Data Appliance.....	1-1
Oracle Big Data Appliance Rack.....	1-1
Server Components.....	1-2
Spare Kit Components	1-2
Oracle Big Data Appliance Management Software.....	1-3
Oracle Big Data Appliance Restrictions on Use.....	1-4
2 Site Requirements	
General Environmental Requirements	2-1
Space Requirements	2-2
Receiving and Unpacking Requirements	2-2
Maintenance Access Requirements	2-3
Flooring Requirements.....	2-3
Electrical Power Requirements.....	2-3
PDUs for Oracle Big Data Appliance	2-4
Facility Power Requirements.....	2-7
Circuit Breaker Requirements	2-8
Grounding Guidelines.....	2-8
Temperature and Humidity Requirements	2-8
Ventilation and Cooling Requirements	2-9
Network Connection Requirements.....	2-11
Ensuring the Site is Ready.....	2-11
3 Network Requirements	
Overview of Network Requirements	3-1
Default Network Connections.....	3-1
Minimum Requirements for the Networks.....	3-2

Network Diagram	3-2
Cabling the Client Network	3-3
Factory Network Settings	3-5
Port Assignments	3-6
4 Using the Oracle Big Data Appliance Configuration Utility	
Overview of Oracle Big Data Appliance Configuration Utility	4-1
Using Oracle Big Data Appliance Configuration Utility	4-1
Using Oracle Open Office Calc	4-2
Generating the Configuration Files	4-2
About the Configuration Files.....	4-3
Checking for Network Errors.....	4-4
Network Configuration	4-5
Software Configuration	4-7
5 Installing Oracle Big Data Appliance at the Site	
Reviewing Safety Guidelines	5-1
Unpacking Oracle Big Data Appliance	5-2
Contents of the Shipping Kit	5-3
Unpacking Oracle Big Data Appliance	5-4
Placing Oracle Big Data Appliance in Its Allocated Space	5-5
Moving Oracle Big Data Appliance.....	5-5
Stabilizing Oracle Big Data Appliance.....	5-6
Attaching a Ground Cable (Optional).....	5-9
Powering On the System the First Time	5-10
Inspecting the Rack After It Is in Place	5-10
Connecting Power Cords	5-11
Powering On Oracle Big Data Appliance	5-12
6 Using Oracle Integrated Lights Out Manager	
Oracle ILOM Overview	6-1
Oracle ILOM Interfaces	6-2
Oracle ILOM Users	6-2
Administrative Network Diagram	6-2
ILOM IP Addresses for Oracle Big Data Appliance Rack Components	6-3
Connecting to ILOM Using the Network	6-3
Connecting to the Command-Line Interface.....	6-3
Connecting to the Web GUI.....	6-3
Opening a Remote KVM Session	6-4
Connecting to ILOM Using a Serial Connection	6-4
Connecting to a Server ILOM.....	6-5
Connecting to ILOM of a Sun Network QDR InfiniBand Gateway Switch	6-5
Using ipmitool Commands When SP Network Information Is Lost.....	6-6
7 Using the dcli Utility	
Overview of the dcli Utility	7-1

Setting Up Passwordless SSH	7-1
Basic Use of dcli.....	7-2
dcli Syntax.....	7-3
Return Values	7-5
Examples	7-5

8 Configuring Oracle Big Data Appliance

Configuring the KVM Switch.....	8-1
Configuring the KVM Switch to Access the Servers.....	8-3
Accessing a Server Using the KVM Switch.....	8-3
Configuring the Cisco Ethernet Switch	8-3
Configuring the InfiniBand Switches.....	8-9
Configuring the Power Distribution Units	8-13
Configuring the Threshold Settings for the Power Distribution Units.....	8-14
Checking Oracle Big Data Appliance Servers	8-15
Configuring the Network	8-18
Verifying the Factory Software Image	8-18
Copying the Configuration Files to Oracle Big Data Appliance.....	8-19
Starting the Network Configuration	8-20
Connecting to the Network	8-21
Completing the Network Configuration	8-21
Reinstalling the Base Image.....	8-23
Checking the Health of the Network	8-25
bdacheckib.....	8-25
bdacheckhw	8-26
bdachecknet	8-27
bdachecksw	8-27
bdadiag	8-28
bdaid	8-29
bdaimagevalidate	8-29
bdaredoclientnet.....	8-29
bdaserials.....	8-30
iblinkinfo	8-30
imageinfo.....	8-31
imagehistory	8-31
listlinkup.....	8-31
showvlan	8-32
showvnics.....	8-32

9 Connecting Multiple Oracle Big Data Appliance Racks

Extending a Rack by Adding Another Rack	9-1
Cabling Two Racks Together	9-2
Cabling Several Racks Together.....	9-3

10 Configuring Ethernet Over InfiniBand

Introduction to Virtual NIC (VNIC).....	10-1
---	------

Setting Up Ethernet Over InfiniBand (EoIB).....	10-1
11 Creating Virtual LANs	
Introduction to VLAN	11-1
Tagging Ethernet Connectors With a VLAN Identifier	11-1
Creating VNICs and Associating Them with VLANs	11-2
For Additional Information.....	11-2
12 Setting Up Auto Service Request	
Understanding Auto Service Request.....	12-1
Getting Ready for ASR Installation	12-2
Installing ASR Manager	12-3
Downloading the Software.....	12-3
Installing the ASM Package.....	12-3
Installing the Oracle Service Tags Bundle.....	12-4
Registering ASR Manager.....	12-4
Verifying ASR Manager.....	12-4
Configuring the Trap Destinations.....	12-5
Activating ASR Assets.....	12-5
Validating Auto Service Request on Oracle Big Data Appliance.....	12-6
Troubleshooting ASR.....	12-7
13 Installing the Oracle Big Data Appliance Software	
Using the Mammoth Utility	13-1
Installing the Software on a Single or Primary Rack.....	13-2
Mammoth Utility Syntax	13-2
What If an Error Occurs During the Installation?.....	13-3
Adding a Rack to an Existing Cluster.....	13-3
Mammoth Utility Steps.....	13-5
14 Maintaining Oracle Big Data Appliance	
Monitoring Ambient Temperature of Servers	14-1
Powering On and Off Oracle Big Data Appliance.....	14-2
Non-Emergency Power Procedures	14-2
Emergency Power-Off Considerations	14-3
Cautions and Warnings.....	14-4
Adding Memory to a Server.....	14-4
Maintaining the Physical Disks of Servers.....	14-4
Verifying the Server Configuration.....	14-4
Changing InfiniBand IP Addresses.....	14-6
Maintaining the InfiniBand Network.....	14-7
Backing Up and Restoring Oracle ILOM Switch Settings.....	14-7
Replacing a Failed InfiniBand Switch	14-8
Verifying InfiniBand Network Operation	14-9
Understanding the Network Subnet Manager Master.....	14-11
Changing the NTP Servers.....	14-12

A Replacement Units

Understanding Repair Categories.....	A-1
Compliance with the European Restriction of the Use of Certain Hazardous Substances in Electrical and Electronic Equipment Directive.....	A-2
Disk Controller Batteries	A-2
Replacing Batteries in Disk Controllers	A-3
LED Status Descriptions	A-3
Sun Fire Server LEDs.....	A-3
Sun Datacenter InfiniBand Switch 36 LEDs	A-3
Parts for Sun Fire Servers	A-4
Parts for InfiniBand Switches.....	A-8
Parts for Cisco Ethernet Switch	A-9
Parts for KVM Switch and KMM Tray Components	A-9
Parts for Power Distribution Units	A-10
Cables and Cable Accessories for Oracle Big Data Appliance Rack.....	A-10

B Cabling Tables

Cable Color Coding	B-1
Oracle Big Data Appliance Rack Layout	B-1
Keyboard, Video and Mouse Network Cable Connections	B-2
Oracle Integrated Lights Out Manager Cable Connections.....	B-3
Single-Phase Power Distribution Unit Cable Connections	B-4
Three-Phase Power Distribution Unit Cable Connections	B-5
In-Rack InfiniBand Switch-to-Switch Cable Connections	B-6
In-Rack InfiniBand Switch-to-Server Cable Connections	B-6

C Multi-Rack Cabling Tables

Understanding Multi-Rack Cabling.....	C-1
Key to Table Abbreviations.....	C-3
Two Rack Cabling	C-4
Three Rack Cabling	C-5
Four Rack Cabling.....	C-7
Five Rack Cabling	C-9
Six Rack Cabling	C-12
Seven Rack Cabling.....	C-15
Eight Rack Cabling	C-19

Index

List of Figures

2-1	Low-Voltage, Single-Phase Power Connector	2-5
2-2	Low-Voltage Three-Phase Power Connector	2-6
2-3	High-Voltage Single-Phase Power Connector	2-6
2-4	High-Voltage Three-Phase Power Connector	2-7
2-5	Typical Data Center Configuration for Perforated Floor Tiles	2-11
3-1	Network Diagram for Oracle Big Data Appliance	3-3
4-1	Installation Template	4-4
5-1	Oracle Big Data Appliance in Shipping Crate	5-3
5-2	Carefully Push Oracle Big Data Appliance from Behind	5-5
5-3	Securing Oracle Big Data Appliance Rack Using the Leveling Feet	5-6
5-4	Location of Mounting Brackets on Rack	5-7
5-5	Bottom View of Oracle Big Data Appliance Rack	5-7
5-6	Base Position to Measure	5-8
5-7	Securing Oracle Big Data Appliance Rack Using Mounting Brackets	5-9
5-8	Earth Ground Attachment Bolt Location	5-10
5-9	Location of Power Cables	5-11
5-10	Power Cord Routing Out the Bottom	5-12
5-11	Power Cord Routing Out the Top	5-12
5-12	PDU Switch Locations	5-13
6-1	Administrative Network in Oracle Big Data Appliance	6-2
6-2	Oracle ILOM Web Interface	6-4
14-1	Oracle ILOM Configuration Backup	14-8
B-1	Rack Layout of Oracle Big Data Appliance Rack	B-2
B-2	Sun Network QDR InfiniBand Gateway Switch Ports	B-6
C-1	Connections Between Spine Switch and Leaf Switches	C-2
C-2	Connections Between Sun Fire Servers and Leaf Switches	C-2
C-3	Connecting Switches Across Two Racks	C-3
C-4	Connecting Switches Across Multiple Racks	C-3

List of Tables

1-1	Oracle Big Data Appliance Rack Components	1-1
1-2	Sun Fire X4270 M2 Server Components	1-2
1-3	Bundled Components and Accessories	1-2
1-4	Bundled Cables	1-2
1-5	In-Rack Spare Cables	1-3
2-1	Environmental Requirements for Oracle Big Data Appliance	2-1
2-2	Access Route Requirements for Oracle Big Data Appliance	2-3
2-3	Maintenance Access Requirements for Oracle Big Data Appliance.....	2-3
2-4	Floor Load Requirements for Oracle Big Data Appliance	2-3
2-5	Low-Voltage 15 kVA Single-Phase PDUs	2-4
2-6	Low-Voltage 15 kVA Three-Phase PDUs	2-5
2-7	High-Voltage 15 kVA Single-Phase PDUs).....	2-6
2-8	High-Voltage 15 kVA Three-Phase PDUs	2-7
2-9	Temperature, Humidity and Altitude Requirements.....	2-9
3-1	Default IP Addresses for Oracle Big Data Appliance.....	3-5
3-2	Oracle Big Data Appliance Port Numbers	3-6
3-3	CDH Port Numbers	3-6
4-1	Naming and Locale Information	4-5
4-2	Network Properties	4-6
4-3	Server Properties	4-6
4-4	Switch Properties	4-6
4-5	PDU Properties.....	4-7
4-6	Installed Components	4-7
4-7	Oracle Auto Service Request	4-7
4-8	Users /Groups.....	4-8
4-9	Cloudera Manager Email Alerts	4-8
7-1	dcli Options.....	7-4
8-1	Threshold Values for Single-Phase, Low-Voltage PDU	8-14
8-2	Threshold Values for Three-Phase, Low-Voltage PDU.....	8-15
8-3	Threshold Values for Single-Phase, High-Voltage PDU	8-15
8-4	Threshold Values for Three-Phase, High-Voltage PDU.....	8-15
13-1	Mammoth Utility Options	13-3
A-1	Oracle Big Data Appliance Rack Repair Categories	A-1
A-2	LED Status Descriptions	A-3
A-3	Sun Datacenter InfiniBand Switch 36 LED Status Descriptions	A-4
A-4	Replacement Parts for a Sun Fire X4270 M2 Server.....	A-5
A-5	Replacement Parts for InfiniBand Switches	A-9
A-6	Replacement Parts forCisco Catalyst Ethernet Switch.....	A-9
A-7	Replacement Parts for KVM Switch and KMM Tray Components.....	A-10
A-8	Replacement Parts for Power Distribution Units.....	A-10
A-9	Replacement Cables for Oracle Big Data Appliance Rack.....	A-11
B-1	KVM Cabling for Oracle Big Data Appliance Rack	B-2
B-2	ILOM Cabling.....	B-3
B-3	Single-Phase PDU Cabling	B-4
B-4	Three-Phase PDU Cabling	B-5
B-5	InfiniBand Switch-to-Switch Cabling	B-6
B-6	InfiniBand Switch-to-Server Cabling	B-6
C-1	Leaf Switch Connections for the First Rack in a Two Rack System.....	C-4
C-2	Leaf Switch Connections for the Second Rack in a Two Rack System.....	C-4
C-3	Leaf Switch Connections for the First Rack in a Three Rack System	C-5
C-4	Leaf Switch Connections for the Second Rack in a Three Rack System.....	C-6
C-5	Leaf Switch Connections for the Third Rack in a Three Rack System	C-6
C-6	Leaf Switch Connections for the First Rack in a Four Rack System	C-7

C-7	Leaf Switch Connections for the Second Rack in a Four Rack System	C-8
C-8	Leaf Switch Connections for the Third Rack in a Four Rack System	C-8
C-9	Leaf Switch Connections for the Fourth Rack in a Four Rack System	C-9
C-10	Leaf Switch Connections for the First Rack in a Five Rack System	C-9
C-11	Leaf Switch Connections for the Second Rack in a Five Rack System	C-10
C-12	Leaf Switch Connections for the Third Rack in a Five Rack System	C-10
C-13	Leaf Switch Connections for the Fourth Rack in a Five Rack System	C-11
C-14	Leaf Switch Connections for the Fifth Rack in a Five Rack System.....	C-11
C-15	Leaf Switch Connections for the First Rack in a Six Rack System	C-12
C-16	Leaf Switch Connections for the Second Rack in a Six Rack System	C-13
C-17	Leaf Switch Connections for the Third Rack in a Six Rack System	C-13
C-18	Leaf Switch Connections for the Fourth Rack in a Six Rack System	C-14
C-19	Leaf Switch Connections for the Fifth Rack in a Six Rack System.....	C-14
C-20	Leaf Switch Connections for the Sixth Rack in a Six Rack System	C-15
C-21	Leaf Switch Connections for the First Rack in a Seven Rack System	C-16
C-22	Leaf Switch Connections for the Second Rack in a Seven Rack System	C-16
C-23	Leaf Switch Connections for the Third Rack in a Seven Rack System	C-17
C-24	Leaf Switch Connections for the Fourth Rack in a Seven Rack System	C-17
C-25	Leaf Switch Connections for the Fifth Rack in a Seven Rack System.....	C-18
C-26	Leaf Switch Connections for the Sixth Rack in a Seven Rack System	C-18
C-27	Leaf Switch Connections for the Seventh Rack in a Seven Rack System.....	C-19
C-28	Leaf Switch Connections for the First Rack in a Eight Rack System	C-20
C-29	Leaf Switch Connections for the Second Rack in a Eight Rack System	C-20
C-30	Leaf Switch Connections for the Third Rack in a Eight Rack System	C-21
C-31	Leaf Switch Connections for the Fourth Rack in a Eight Rack System	C-21
C-32	Leaf Switch Connections for the Fifth Rack in a Eight Rack System.....	C-22
C-33	Leaf Switch Connections for the Sixth Rack in an Eight Rack System.....	C-23
C-34	Leaf Switch Connections for the Seventh Rack in a Eight Rack System.....	C-23
C-35	Leaf Switch Connections for the Eighth Rack in a Eight Rack System	C-24

Preface

This guide describes Oracle Big Data Appliance for acquiring, organizing, and analyzing very large data sets. It includes information about hardware operations, site planning and configuration, and physical, electrical, and environmental specifications.

This preface contains the following topics:

- [Audience](#)
- [Documentation Accessibility](#)
- [Related Documentation](#)
- [Conventions](#)
- [Backus-Naur Form Syntax](#)

Audience

This guide is intended for Oracle Big Data Appliance customers and those responsible for data center site planning, installation, configuration, and maintenance of Oracle Big Data Appliance.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documentation

The following Oracle libraries contain hardware information for Oracle Big Data Appliance. Links to these libraries are available through the Big Data library at

http://docs.oracle.com/cd/E27101_01/index.htm

- Sun Fire X4270 M2 Server library:
<http://download.oracle.com/docs/cd/E19245-01/index.html>

- Sun Rack II 1042 and 1242 library:
<http://download.oracle.com/docs/cd/E19844-01/index.html>
- Sun Network QDR InfiniBand Gateway Switch library:
<http://download.oracle.com/docs/cd/E19671-01/index.html>
- Sun Datacenter InfiniBand Switch 36 library:
<http://download.oracle.com/docs/cd/E19197-01/index.html>
- Oracle Integrated Lights Out Manager (ILOM) 3.0 library:
<http://download.oracle.com/docs/cd/E19860-01/index.html>

Conventions

The following text conventions are used in this document:

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

Backus-Naur Form Syntax

The syntax in this reference is presented in a simple variation of Backus-Naur Form (BNF) that uses the following symbols and conventions:

Symbol or Convention	Description
[]	Brackets enclose optional items.
{ }	Braces enclose a choice of items, only one of which is required.
	A vertical bar separates alternatives within brackets or braces.
...	Ellipses indicate that the preceding syntactic element can be repeated.
delimiters	Delimiters other than brackets, braces, and vertical bars must be entered as shown.
boldface	Words appearing in boldface are keywords. They must be typed as shown. (Keywords are case-sensitive in some, but not all, operating systems.) Words that are not in boldface are placeholders for which you must substitute a name or value.

Introduction to Oracle Big Data Appliance

This chapter describes the features and hardware components of Oracle Big Data Appliance. It also includes usage information.

This chapter contains these topics:

- [About Oracle Big Data Appliance](#)
- [Oracle Big Data Appliance Rack](#)
- [Server Components](#)
- [Spare Kit Components](#)
- [Oracle Big Data Appliance Management Software](#)
- [Oracle Big Data Appliance Restrictions on Use](#)

About Oracle Big Data Appliance

Oracle Big Data Appliance is an engineered system of hardware and software optimized to capture and analyze the massive volumes of unstructured data generated by social media feeds, email, weblogs, photographs, smart meters, sensors, and similar devices.

Oracle Big Data Appliance is engineered to work with Oracle Exadata Database Machine and Oracle Exalytics Business Intelligence Machine to provide the most advanced analysis of all data types, with enterprise-class performance, availability, supportability, and security.

The Oracle Linux operating system and Cloudera's Distribution including Apache Hadoop (CDH) underlie all other software components installed on Oracle Big Data Appliance.

See Also: *Oracle Big Data Appliance Software User's Guide*

Oracle Big Data Appliance Rack

Oracle Big Data Appliance consists of a full rack and the components to connect to your network.

[Table 1–1](#) lists the components in an Oracle Big Data Appliance Rack.

Table 1–1 *Oracle Big Data Appliance Rack Components*

Quantity	Description
18	Sun Fire X4270 M2 Servers

Table 1–1 (Cont.) Oracle Big Data Appliance Rack Components

Quantity	Description
1	Sun Rack II 1242 Base
2	NM2-GW Sun Network QDR InfiniBand Gateway Switches
1	NM2-36P Sun Datacenter InfiniBand Switch 36
1	Cisco Catalyst 4948 Ethernet Switch
1	KVM
2	Power Distribution Units

Server Components

Table 1–2 lists the components of each Sun Fire Server.

Table 1–2 Sun Fire X4270 M2 Server Components

Quantity	Description
1	Sun Fire X4270 M2 Server Base
2	Six-Core Intel Xeon X5675 Processors (3.06 GHz)
6	8GB DDR3 2RX4 1333MHz DIMMs (48 GB RAM)
12	3TB 3.5" 7200 RPM drives
1	Ethernet Port for Oracle Integrated Lights Out Manager v3.0 for remote management
1	Dual-port 4X QDR (40 Gb/s) InfiniBand Host Channel Adapter Network Interface Card
1	Host Bus Adapter (Disk Controller) with 8 internal ports and 512 MB battery-backed write cache
2	Redundant power supplies and fans
1	Embedded Gigabit Ethernet Port
1	USB-2 4GB Flash Drive

Spare Kit Components

The Oracle Big Data Appliance spare kit includes disk drives, accessory kits, cables, and documentation. Table 1–3 and Table 1–4 identify the bundled spares.

Table 1–3 Bundled Components and Accessories

Part Number	Quantity	Description
390-0476-03	2	3TB 3.5" 7200 RPM drives
590-896-501B	1	LCD Accessory Kit
630-446-501	1	KVM Switch Accessory Kit
53-2332-xx	1	Cisco Switch Documents and Accessories Kit

Table 1–4 Bundled Cables

Part Number	Length	Quantity	Description
530-4446-01	5 Meter	10	QSFP passive copper cable

Table 1–4 (Cont.) Bundled Cables

Part Number	Length	Quantity	Description
530-4445-01	3 Meter	6	QSFP passive copper cable
350-1519-01	--	3	ASSY, NM2 serial cable sets

Table 1–4 lists the spare parts that are secured to the center section of an Oracle Big Data Appliance Rack.

Table 1–5 In-Rack Spare Cables

Part Number	Length	Quantity	Description
530-4445-01	3 Meter	2	QSFP passive copper cable
530-4437-01	10 Foot	1	Orange Cat 5
530-4435-01	10 Foot	1	Black Cat 5
530-4433-01	10 Foot	1	Red Cat 5
530-4432-01	7 Foot	1	Blue Cat 5

Oracle Big Data Appliance Management Software

The software components are installed on all 18 servers in Oracle Big Data Appliance Rack. Oracle Linux, required drivers, firmware, and hardware verification utilities are factory installed. Among them are the following tools:

- Oracle Integrated Lights Out Manager (Oracle ILOM) consists of preinstalled, dedicated hardware and software that you can use to manage and monitor the servers and switches in Oracle Big Data Appliance Rack. See [Chapter 6](#).
- `setup-root-ssh` sets up passwordless SSH for the `root` user among all the servers in an Oracle Big Data Appliance Rack. See ["Setting Up Passwordless SSH"](#) on page 7-1.
- `dcli` executes commands across a group of servers on Oracle Big Data Appliance and returns the output. See [Chapter 7](#).
- `bdcheckib` checks the private InfiniBand network. See ["Checking the Health of the Network"](#) on page 8-25.
- `bdchecknet` checks the network configuration. See ["Checking the Health of the Network"](#) on page 8-25.
- `iblinkinfo` lists the connections in the InfiniBand network. See ["Checking the Health of the Network"](#) on page 8-25.
- `listlinkup` identifies the active Ethernet ports. See ["Checking the Health of the Network"](#) on page 8-25.
- `showvlan` lists the Virtual LANs configured on a Sun Network QDR InfiniBand Gateway Switch. See ["Checking the Health of the Network"](#) on page 8-25.
- `showvnics` lists the Virtual NICs created on a Sun Network QDR InfiniBand Gateway Switch. See ["Checking the Health of the Network"](#) on page 8-25.

All end-user software is installed on site using the Mammoth Utility. See [Chapter 13](#), ["Installing the Oracle Big Data Appliance Software."](#)

See Also: *Oracle Big Data Appliance Software User's Guide*

Oracle Big Data Appliance Restrictions on Use

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

- Oracle Big Data Appliance hardware cannot be modified or customized, with one exception. The only allowed hardware modification to Oracle Big Data Appliance is to the administrative 48-port Cisco 4948 Ethernet switch.

Customers may choose to do the following:

- Replace the Gigabit Ethernet switch, at the customer's 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 Oracle Big Data 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 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 4948 Ethernet switch in the rack can be turned off and disconnected from the data center network.
- Customers cannot update the firmware directly on Oracle Big Data Appliance servers. The firmware is updated as part of an Oracle Big Data Appliance patch.
- Customers can update the firmware of the other components of Oracle Big Data Appliance:
 - Customers can update the IOS and firmware versions on the Cisco 4948 Gigabit Ethernet switch to meet their data center requirements.
 - Customers can update the firmware of the KVM switch and the KMM (keyboard, monitor, mouse) as needed.
 - Customers can update the firmware of the components of the Oracle Big Data Appliance servers provided they comply with the restrictions documented in My Oracle Support (MOS) Master Note 1434477.1 and its related notes.
 - Customers can update the firmware of the InfiniBand switches provided they comply with the validated versions documented in MOS Master Note 1434477.1 and its related notes.
- Customers may load additional software on Oracle Big Data Appliance servers. Oracle does not support questions or issues with the non-standard modules. If a server fails, and Oracle suspects the failure may have been caused by a non-standard module, then Oracle Support may refer the customer to the vendor of the non-standard module or ask that the issue be reproduced without the non-standard module. Modifying the server operating system other than by applying official patches and upgrades is not supported. InfiniBand-related packages must always be maintained at the officially supported release.
- Customers can use Cloudera Manager to stop and start services, but cannot use it to move services from one server to another. The Hadoop services such as NameNode and JobTracker must remain on the servers where they were installed by the Mammoth Utility.

See Also: *Oracle Big Data Appliance Software User's Guide* for service locations

- A Hadoop cluster must consist of one or more Oracle Big Data Appliance Racks. One Oracle Big Data Appliance Rack cannot be split into multiple Hadoop clusters.
- Customers cannot connect USB devices to Oracle Big Data Appliance servers except as documented in this guide. In those documented situations, the USB device should not draw more than 100 mA of power.
- The network ports on the servers can be used to connect to external non-Sun Fire Servers using iSCSI or NFS. However, the Fibre Channel Over Ethernet (FCoE) protocol is not supported.
- Only switches specified for use in Oracle Big Data Appliance, Oracle Exadata Database Machine, and Oracle Exalogic Elastic Cloud may be connected to the InfiniBand network. Connecting third-party switches and other switches not used in these engineered systems is not supported.

Site Requirements

This chapter describes the site requirements for Oracle Big Data Appliance. Use this chapter when completing the *Oracle Big Data Appliance Site Checklists*.

This chapter contains these topics:

- [General Environmental Requirements](#)
- [Space Requirements](#)
- [Flooring Requirements](#)
- [Electrical Power Requirements](#)
- [Temperature and Humidity Requirements](#)
- [Ventilation and Cooling Requirements](#)
- [Network Connection Requirements](#)
- [Ensuring the Site is Ready](#)

See Also: *Oracle Big Data Appliance Site Checklists*

General Environmental Requirements

[Table 2–1](#) shows the general environmental. The other sections in this chapter provide detailed information.

Table 2–1 *Environmental Requirements for Oracle Big Data Appliance*

Environmental Component	Oracle Big Data Appliance
Weight	909.5 kg (2005 lbs) See Also: " Flooring Requirements " on page 3
Acoustic levels	8.3 B
Power	Maximum: 12.0 kW (12.25 kVA) Typical: 8.4 kW (8.6 kVA) ¹ See Also: " Electrical Power Requirements " on page 3
Cooling	Maximum: 40,971 BTU/hour (41,807 kJ/hour) Typical: 28,680 BTU/hour (29,265 kJ/hour) See Also: " Temperature and Humidity Requirements " on page 8 and " Ventilation and Cooling Requirements " on page 9

Table 2–1 (Cont.) Environmental Requirements for Oracle Big Data Appliance

Environmental Component	Oracle Big Data Appliance
Air flow front-to-back (subject to actual data center environment)	Maximum: Approximately 1,886 CFM Typical: Approximately 1,340 CFM See Also: " Temperature and Humidity Requirements " on page 8, and " Ventilation and Cooling Requirements " on page 9
IP Addresses	24 for Ethernet network 18 for InfiniBand network See Also: Chapter 3, "Network Requirements."
Network drops	Minimum of 2 network drops See Also: Chapter 3, "Network Requirements."
External connectivity	2 x 1 GbE Ethernet ports 18 x 10 GbE Ethernet ports See Also: Chapter 3, "Network Requirements."

¹ Varies based on application load.

Space Requirements

The space requirements for Oracle Big Data Appliance are as follows:

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

The minimum ceiling height for the cabinet is 2300 mm (90 inches), measured from the true floor or raised floor, whichever is higher. An additional 914 mm (36 inches) is for top clearance. The space above the cabinet and its surroundings must not restrict the movement of cool air between the air conditioner and the cabinet, or the movement of hot air coming out of the top of the cabinet.

Receiving and Unpacking Requirements

Before Oracle Big Data Appliance arrives, ensure that the receiving area is large enough for the package. The following are the package specifications:

- Shipping height: 2159 mm (85 inches)
- Shipping width: 1219 mm (48 inches)
- Shipping depth: 1575 mm (62 inches)
- Shipping weight: 993 kgs. (2189 lbs.)
- Net weight: 909.5 kgs. (2005 lbs.)

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 not, you must provide a standard forklift or other means to unload the rack. You can also request that the rack be shipped in a truck with a lift gate.

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

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

Table 2–2 Access Route Requirements for Oracle Big Data Appliance

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)

Maintenance Access Requirements

The maintenance area must be large enough for Oracle Big Data Appliance and have the required access space. For example, the required space to remove the side panels is 675.64 mm (26.6 inches). [Table 2–3](#) lists the maintenance access requirements.

Open tiles are required for electrical access.

Table 2–3 Maintenance Access Requirements for Oracle Big Data Appliance

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

Flooring Requirements

Oracle recommends that Oracle Big Data Appliance be installed on raised flooring. The site floor and the raised flooring must be able to support the total weight.

[Table 2–4](#) lists the floor load requirements.

Table 2–4 Floor Load Requirements for Oracle Big Data Appliance

Description	Requirement
Maximum allowable weight of installed rack equipment	997.9 kg (2200 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)	1004.71 kg (2215 lbs)

Electrical Power Requirements

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

- Fluctuations caused by brownouts
- Wide and rapid variations in input voltage levels or in input power frequency

- Electrical storms
- Faults in the distribution system, such as defective wiring

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

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

PDUs for Oracle Big Data Appliance

The type of PDU depends on the location where Oracle Big Data Appliance is installed:

- North America, South America, Japan, and Taiwan use low-voltage PDUs.
- Europe, the Middle East and Africa (EMEA) use high-voltage PDUs.
- Asia Pacific (APAC) except for Japan and Taiwan use high-voltage PDUs.

Refer to the appropriate sections for your location:

- [Low-Voltage 15 kVA Single-Phase PDUs](#)
- [Low-Voltage 15 kVA Three-Phase PDUs](#)
- [High-Voltage 15 kVA Single-Phase PDUs](#)
- [High-Voltage 15 kVA Three-Phase PDUs](#)

Low-Voltage 15 kVA Single-Phase PDUs

[Table 2–5](#) lists the requirements for low-voltage one-phase PDUs for North America, South America, Japan, and Taiwan. There are two PDUs per rack.

Table 2–5 Low-Voltage 15 kVA Single-Phase PDUs

Option	Requirement per PDU
Number of inputs	3 x 30 A single phase
Voltage	200 - 240 VAC
Frequency	50/60 Hz
Current	24 A maximum per input
Power rating	15 kVA
Output current	72 A (3 x 24 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection ¹	20A
Data center receptacle	15kVA, with three 30A/250V 2-pole/3-wire NEMA L6-30P plugs

¹ UL489 2-pole circuit breaker

The following are needed to connect Oracle Big Data Appliance to a low-voltage single-phase power source:

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

Table 2–1 shows the low-voltage, single-phase PDU power connector for North America, South America, Japan and Taiwan.

Figure 2–1 Low-Voltage, Single-Phase Power Connector



Low-Voltage 15 kVA Three-Phase PDUs

Table 2–6 lists the requirements for low-voltage three-phase PDUs for North America, South America, Japan, and Taiwan. There are two PDUs per rack.

Table 2–6 Low-Voltage 15 kVA Three-Phase PDUs

Option	Requirement per PDU
Number of inputs	1 x 60 A three phase 4-wire
Voltage	190 - 220 VAC
Frequency	50/60 Hz
Current	40 A maximum per phase
Power rating	14.4 kVA
Output current	69 A (3 x 23 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection ¹	20A
Data center receptacle	15kVA IEC 60309 60A 4-pin 250VAC three phase IP67

¹ UL489 2-pole circuit breaker

The following are needed to connect Oracle Big Data Appliance to a low-voltage three-phase power source:

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

Table 2–2 shows the low-voltage three-phase PDU power connector for North America, South America, Japan and Taiwan.

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



High-Voltage 15 kVA Single-Phase PDUs

Table 2–7 lists the requirements for 15 kVA high-voltage one-phase PDUs for Europe, the Middle East, Africa, and Asia Pacific, except Japan and Taiwan. There are two PDUs per rack.

Table 2–7 High-Voltage 15 kVA Single-Phase PDUs)

Option	Requirement per PDU
Number of inputs	3 x 25 A single phase
Voltage	220-240 VAC
Frequency	50/60 Hz
Current	25 A maximum per input
Power rating	15 kVA
Output current	72 A (3 x 24 A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection ¹	20A
Data center receptacle	15kVA, with three blue 32A/240V splash-proof 2-pole/3-wire IEC 60309 plugs

¹ UL489 1-pole circuit breaker

The following are needed to connect Oracle Big Data Appliance to a high-voltage single-phase power source:

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

Table 2–3 shows the high-voltage single-phase PDU power connector for EMEA and APAC, except for Japan and Taiwan.

Figure 2–3 High-Voltage Single-Phase Power Connector



High-Voltage 15 kVA Three-Phase PDUs

Table 2–8 lists the requirements for high-voltage three-phase PDUs for Europe, the Middle East, Africa, and Asia Pacific, except Japan and Taiwan. There are two PDUs per rack. The following specifications are per PDU.

Table 2–8 High-Voltage 15 kVA Three-Phase PDUs

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

¹ UL489 1-pole circuit breaker

The following are needed to connect Oracle Big Data Appliance to a high-voltage three-phase power source:

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

Table 2–4 shows the high-voltage three-phase PDU power connector for EMEA and APAC, except for Japan and Taiwan.

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



Facility Power Requirements

Electrical work and installations must obey 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.

To prevent catastrophic failures, design the input power sources to ensure adequate power is provided to the PDUs. Use dedicated AC breaker panels for all power circuits that supply power to the PDU. When planning for power distribution requirements, balance the power load between available AC supply branch circuits. In the United States and Canada, ensure that the overall system AC input current load does not exceed 80 percent of the branch circuit AC current rating.

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

Circuit Breaker Requirements

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

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

Note: The owner supplies the circuit breakers. Each power cord requires one circuit breaker.

Grounding Guidelines

The cabinets for Oracle Big Data Appliance are shipped with grounding-type power cords. Always connect the cords to grounded power outlets. Because different grounding methods are used depending on location, check the grounding type and refer to documentation, such as International Electrotechnical Commission (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

Airflow through Oracle Big Data Appliance is from front to back. Refer to [Table 2-1](#) for information on cooling and airflow.

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

Excessive internal temperatures may result in full or partial shutdown of Oracle Big Data Appliance.

[Table 2-9](#) lists the temperature, humidity and altitude requirements for operating and non-operating systems.

Table 2–9 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) maximum	Ambient temperature is reduced by 1 degree Celsius per 300 m above 900 m altitude above sea level

Set conditions to the optimal temperature and humidity ranges to minimize the chance of downtime due to component failure. Operating Oracle Big Data Appliance 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 if the air conditioning system fails.

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

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

Note: Electrostatic discharge (ESD) is easily generated and hard to dissipate in areas of low relative humidity, such as below 35 percent. ESD becomes critical when humidity drops below 30 percent. 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, monitor individual points corresponding to multiple exhaust vents in the main unit and other units in the room. Pay special consideration to humidity when using underfloor ventilation, and monitor each point close to an exhaust vent. Distribution of the temperature and humidity across the entire room is uneven.


Oracle Big Data Appliance is designed to function while mounted in a natural convection air flow. Follow these requirements to meet the environmental specification:

- Ensure the server has adequate air flow.
- 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 914 mm (36 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, provided that cold air from the tiles can flow into the rack. Inadequate cold air flow could result in a higher inlet temperature in the servers due to exhaust air recirculation. Oracle recommends four floor tiles for Oracle Big Data Appliance

[Figure 2-5](#) shows a typical installation of the floor tiles for Oracle Big Data Appliance in a typical data center.

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



Network Connection Requirements

Before installation, network cables must run from your existing network infrastructure to the installation site. The requirements to connect Oracle Big Data Appliance to your existing network infrastructure are as follows:

- Management network connection requirements
 - 1 Ethernet connection for the Ethernet switch in the rack
 - 1 Ethernet connection for the KVM switch in the rack
- Client access network connection requirements
 - 1 Ethernet connection for each server in the rack, for a total of 18 connections.
 - 2 10 GigE connections, one to each Sun Network QDR InfiniBand Gateway Switch

See Also: [Chapter 3, "Network Requirements"](#)

Ensuring the Site is Ready

Before Oracle Big Data Appliance is delivered to the site, perform these tasks to ensure the site is ready:

- [Task 1, "Review Site Requirements"](#)
- [Task 2, "Complete the Oracle Big Data Appliance Configuration Worksheets"](#)
- [Task 3, "Complete the Oracle Big Data Appliance Configuration Utility"](#)
- [Task 4, "Configure the Network"](#)

- [Task 5, "Prepare the Site Based on Requirements"](#)

Task 1 Review Site Requirements

Review the site requirements in this chapter and in the *Oracle Big Data Appliance Site Checklists* to understand the requirements of Oracle Big Data Appliance.

Task 2 Complete the Oracle Big Data Appliance Configuration Worksheets

Provide the network configuration details and select the optional software components to be installed. Complete this document for every set of appliances that will be configured as a single CDH cluster. Give the completed worksheets to your Oracle representative.

Task 3 Complete the Oracle Big Data Appliance Configuration Utility

Complete the spreadsheet from the information in the Oracle Big Data Appliance Configuration Worksheets and generate the Installation Template and configuration files. This step is typically performed by an Oracle representative.

See Also: [Chapter 4, "Using the Oracle Big Data Appliance Configuration Utility"](#)

Task 4 Configure the Network

Configure the existing network using the Installation Template provided by your Oracle representative. This includes registering the networks in the Domain Name System (DNS), assigning IP addresses, and configuring data center switches and firewalls.

See Also:

- [Chapter 3, "Network Requirements"](#)
- [Chapter 4, "Using the Oracle Big Data Appliance Configuration Utility"](#)
- [Chapter 8, "Configuring Oracle Big Data Appliance"](#)

Task 5 Prepare the Site Based on Requirements

Prepare the site based on the requirements, such as installing the network cables and power supplies, before the arrival of Oracle Big Data Appliance:

1. Review the safety guidelines.

See Also: ["Reviewing Safety Guidelines"](#) on page 5-1

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

See Also: ["Space Requirements"](#) on page 2-2

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

See Also: ["Flooring Requirements"](#) on page 2-3

4. Install network cables for Oracle Big Data Appliance.

See Also: ["Network Connection Requirements"](#) on page 2-11

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

See Also:

- ["Electrical Power Requirements"](#) on page 2-3
- The *Sun Rack II Power Distribution Units User's Guide* for the Power Distribution Unit (PDU) power requirements

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

See Also: ["Ventilation and Cooling Requirements"](#) on page 2-9

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

Network Requirements

This chapter describes the network requirements for Oracle Big Data Appliance. This chapter contains these topics:

- [Overview of Network Requirements](#)
- [Cabling the Client Network](#)
- [Factory Network Settings](#)
- [Port Assignments](#)

Overview of Network Requirements

Oracle Big Data Appliance includes 18 servers and the equipment to connect the servers to your network. The network connections allow the servers to be administered remotely and clients to connect to them. Use the information in this section to configure the environment for Oracle Big Data Appliance.

Each server has the following network components and interfaces:

- 1 Dual-port 4X QDR (40 Gb/s) InfiniBand Host Channel Adapter network interface card
- 1 Ethernet port for Oracle Integrated Lights Out Manager v3.0 for remote management
- 1 Gigabit Ethernet port

Default Network Connections

The installation process automatically discovers whether each Sun Network QDR InfiniBand Gateway Switch has at least one 10GigE connection. If they do, then two Virtual Network Interface Cards (VNICs) are configured for each server: one for each switch bonded as bondeth0 in active/passive failover mode. The VNICs are assigned automatically to the available 10GigE connections in round robin fashion. For example, if each switch has three available 10GigE connections, then the VNIC on server 1 is assigned to 10GigE port 1, server 2 to port 2, server 3 to port 3, server 4 to port 1, and so on.

All VNICs are assigned to the same default VLAN. To assign different VLANs to different VNICs, you must delete the initial VNICs and create your own manually.

See Also:

- [Chapter 10, "Configuring Ethernet Over InfiniBand"](#)
- [Chapter 11, "Creating Virtual LANs"](#)

Minimum Requirements for the Networks

Additional configuration, such as defining multiple virtual local area networks (VLANs) or enabling routing, may be required for the switch to operate properly in your environment. If additional configuration is needed, then your network administrator must perform the necessary configuration steps during installation of Oracle Big Data Appliance.

To deploy Oracle Big Data Appliance, ensure that you meet the minimum network requirements. Oracle Big Data Appliance uses three networks. Each network must be on a distinct and separate subnet from the others. These are the network descriptions:

- **Administrative network:** This 1 Gigabit Ethernet (GigE) network connects to your existing administrative network and is used to administer all components of Oracle Big Data Appliance. It connects the servers, ILOM, and switches connected to the Ethernet switch in the rack.

There are three uplinks to the administrative network:

- From the Ethernet switch in the rack
- From the KVM switch in the rack
- From each power distribution unit (PDU).

Network connectivity to the PDUs is required only if the electric current is monitored remotely.

Each server has two network interfaces for administration. One provides administrative access to the operating system through the ETH0 Ethernet interface, and the other provides access to the Integrated Lights Out Manager through the ILOM Ethernet interface. Oracle Big Data Appliance is delivered with the ETH0 and ILOM interfaces connected to the Ethernet switch on the rack. Do not use the ETH0 interface on the servers for client network traffic. Cabling or configuration changes to these interfaces are not permitted.

- **Client access network:** This 10GigE network connects the servers through the gateway switches to your existing client network and is used for client access to the servers. Client applications access the software through this network using Single Client Access Name (SCAN).

There are two Sun Network QDR InfiniBand Gateway Switches in the rack. Each switch supports 1 to 8 connections for client access for a total of up to 16 client network connections. For failover, you must have at least one connection from each switch and scale up according to your requirements for loading data and providing client access.

- **InfiniBand private network:** This network connects the servers using the InfiniBand switches on the rack and the bondib0 interface. This non-routable network is fully contained in Oracle Big Data Appliance and does not connect to your existing network. This network is automatically configured during installation.

Note: All networks must be on distinct and separate subnets from each other.

Network Diagram

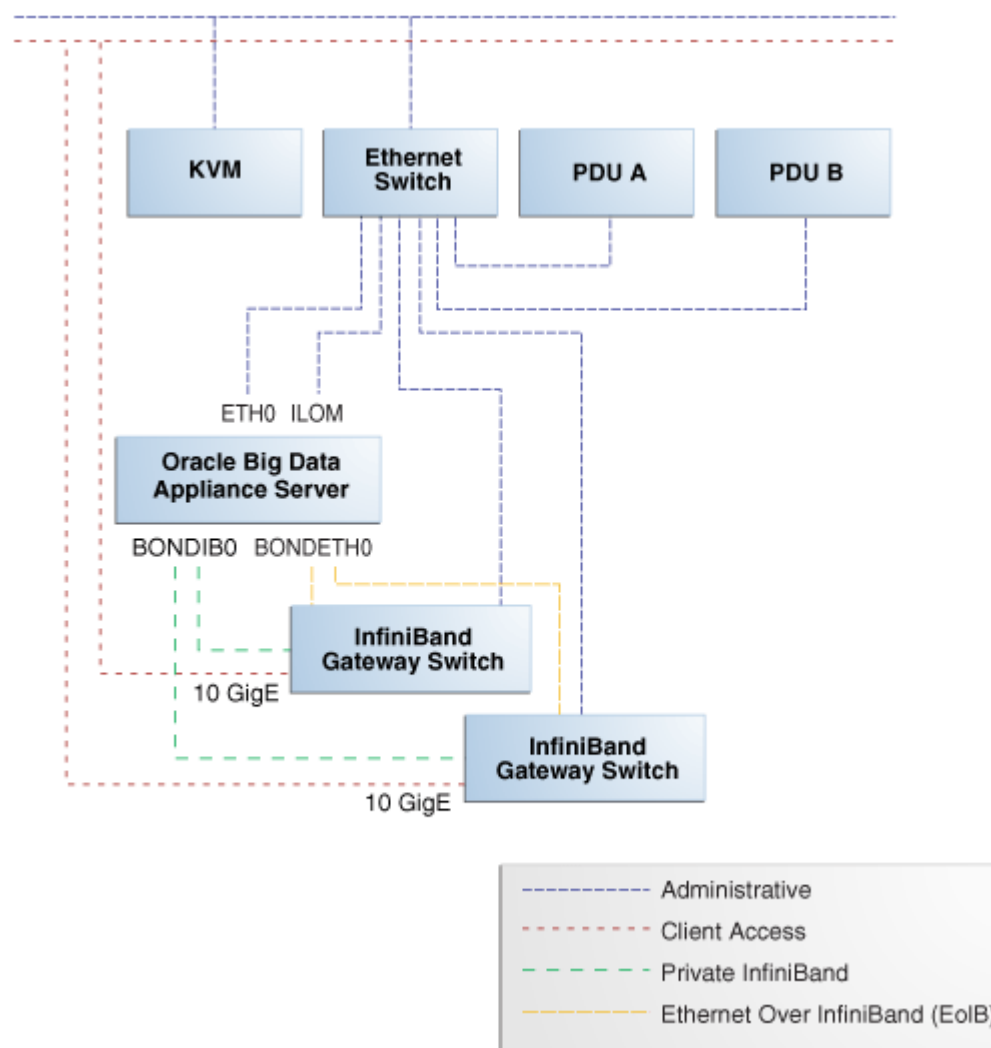
The servers are configured on the network as follows:

- ETH0: Provides access to the operating system using the administrative network.

- **BONDETH0**: Provides access to the server using the client access network. The SCAN addresses are defined on this interface.
- **ILOM**: Provides access to the Integrated Lights Out Manager (ILOM) using the administrative network.

Figure 3–1 shows the network diagram. Use the `preinstall-checkip.sh` script to verify the cable connections when cabling Oracle Big Data Appliance to the existing network. See "Checking for Network Errors" on page 4-4.

Figure 3–1 Network Diagram for Oracle Big Data Appliance



Cabling the Client Network

Each of the two Sun Network QDR InfiniBand Gateway Switches in Oracle Big Data Appliance has eight 10GbE ports. The two switches enable you to create up to 16 10GbE connections for each rack. You can determine how many connections to create based on the bandwidth needed for the client network. For proper functioning, at least one of the eight ports of each gateway switch must have an active connection to the site's 10GbE network.

Physical Ethernet connections are created only between the site network and the gateway switches. The Oracle Big Data Appliance servers are connected only by

InfiniBand to those switches. Each server has two InfiniBand connections, one to each gateway switch, in an active backup mode; only the active InfiniBand connection is used for all InfiniBand traffic to that server. If that connection fails, it immediately fails over to the other connection.

Half of the Oracle Big Data Appliance servers have active connections to one gateway switch, and the other half have active connections to the other gateway switch. Inside Oracle Big Data Appliance, the client network traffic is transported over those InfiniBand connections using the Ethernet over InfiniBand (EoIB) protocol. As a result, each BDA server has two Virtual Network Interfaces (VNICs) that are bonded in the same active backup mode. Each VNIC is assigned a specific port on the gateway switch. If a switch has only one 10GbE connection, then all VNICs for that switch point to the same port. If a switch has multiple connections, then the VNICs are spread across those ports in round-robin fashion.

For example, if you create three 10GbE uplinks from each gateway switch, then the client network traffic from the servers is handled by the switches as shown in the following table:

Server Number	Active Link	Backup Link
1-3	GW Switch 1 Link 1	GW Switch 2 Link 1
4-6	GW Switch 1 Link 2	GW Switch 2 Link 2
...
16-18	GW Switch 2 Link 3	GW Switch 1 Link 3

If you do not want all servers balanced across the available connections, then you can manually delete the VNICs that are automatically created during the network setup process. You can then create your own VNIC to 10GbE port mapping. See [Chapter 10, "Configuring Ethernet Over InfiniBand"](#) and [Chapter 11, "Creating Virtual LANs."](#)

Depending on the number of 10GbE connections, the client network traffic for multiple servers may be sent across the same physical 10GbE connection. The maximum number of connections provides 160 Gbps of client network bandwidth to Oracle Big Data Appliance.

Although you can create up to eight 10GbE connections for each gateway switch, its physical ports are 40 Gbps QSFP ports. Each gateway switch has two of these physical ports reserved for 10GbE connectivity and are equipped with QSFP transceivers that take an optical cable with a MTP/MPO connector. On the site network side are typically 10Gbps ports with SFP+ transceivers that take LC connectors. For these connections, you can use splitter cables that have a single male MTP/MPO connector on one end and four pairs of LC connectors on the other end. Each 10Gbps SFP+ transceiver takes a duplex LC connector for a total of 4 pairs. Thus, you can use a single splitter cable to create up to four 10GbE connections. However, all four duplex LC connectors do not require a connection to the site network.

The splitter cables are not provided with Oracle Big Data Appliance and must be ordered separately from Oracle. They are available in lengths of 10, 20 and 50 meters. Oracle recommends that you order the SFP+ transceivers used at the site end of the 10GbE network from the same manufacturer as the 10GbE switch.

When multiple Oracle Big Data Appliance Racks are connected to form a single Hadoop cluster, Oracle strongly recommends that you spread the 10GbE connections across the switches in different racks. Every gateway switch in every rack must have at least one 10GbE connection. Oracle does not recommend that the 10GbE connections

to the switches in a single rack be used for all the BDA servers in all the racks, although that configuration may be possible.

Factory Network Settings

This initial network configuration is set at the factory for Oracle Big Data Appliance:

- **Gateway:** 192.168.1.254 in all devices as required
- **Subnet Mask:** 255.255.255.0 in all devices as required
- **IP Address Range:** 192.168.1.1 to 192.168.1.211

Table 3–1 lists the default IP addresses for Oracle Big Data Appliance.

Table 3–1 Default IP Addresses for Oracle Big Data Appliance

Host	Administrative IP Addresses	ILOM IP Addresses	InfiniBand Bonded IP Addresses
bda18	192.168.1.18	192.168.1.118	192.168.10.18
bda17	192.168.1.17	192.168.1.117	192.168.10.17
bda16	192.168.1.16	192.168.1.116	192.168.10.16
bda15	192.168.1.15	192.168.1.115	192.168.10.15
bda14	192.168.1.14	192.168.1.114	192.168.10.14
bda13	192.168.1.13	192.168.1.113	192.168.10.13
bda12	192.168.1.12	192.168.1.112	192.168.10.12
bda11	192.168.1.11	192.168.1.111	192.168.10.11
bda10	192.168.1.10	192.168.1.110	192.168.10.10
bdasw-ib3	192.168.1.203	--	--
KVM Tray	--	--	--
KVM Switch	--	--	--
Cisco Switch	192.168.1.200	--	--
bdasw-ib2	192.168.1.202	--	--
bda09	192.168.1.9	192.168.1.109	192.168.10.9
bda08	192.168.1.8	192.168.1.108	192.168.10.8
bda07	192.168.1.7	192.168.1.107	192.168.10.7
bda06	192.168.1.6	192.168.1.106	192.168.10.6
bda05	192.168.1.5	192.168.1.105	192.168.10.5
bda04	192.168.1.4	192.168.1.104	192.168.10.4
bda03	192.168.1.3	192.168.1.103	192.168.10.3
bda02	192.168.1.2	192.168.1.102	192.168.10.2
bda01	192.168.1.1	192.168.1.101	192.168.10.1
bdasw-ib1	196.168.1.201	--	--
PDU A	192.168.1.210	--	--
PDU B	192.168.1.211	--	--

Port Assignments

Table 3–2 identifies the port numbers used by Oracle Big Data Appliance software. Ensure that these ports are free before configuring the network.

Table 3–2 Oracle Big Data Appliance Port Numbers

Port	Used By
2	ssh
80	yumrepos (only during installation)
111	portmap
668	rpc.statd
3306	MySQL Database
5000	Oracle NoSQL Database registration
5001	Oracle NoSQL Database administration
5010-5020	Oracle NoSQL Database processes
6481	xinetd (service tag)
8139	Puppet nodes
8140	Puppet parent
20910	Oracle Data Integrator agent
30920	Automated Service Monitor (ASM)

Table 3–3 lists the ports used by Cloudera's Distribution including Apache Hadoop (CDH). For additional details about these port assignments, go to the Cloudera website at

<http://ccp.cloudera.com/display/CDHDOC/Configuring+Ports+for+CDH3>

Table 3–3 CDH Port Numbers

Port	Component	Service	Access
0	HDFS	Thrift Plugin DataNode	--
0	MapReduce	TaskTracker	Localhost
1004	HDFS	Secure DataNode	External
1006	HDFS	Secure DataNode	External
2181	HBase	HQuorumPeer	--
2181	ZooKeeper	Server	External
2888	HBase	HQuorumPeer	--
2888	ZooKeeper	Server	Internal
3181	Flume	ZooKeeper Server	--
3182	Flume	ZooKeeper Server	--
3183	Flume	ZooKeeper Server	--
3888	HBase	HQuorumPeer	--
3888	ZooKeeper	Server	Internal
8001	Hue	Job Submission Server	Internal

Table 3–3 (Cont.) CDH Port Numbers

Port	Component	Service	Access
8002	Hue	Beeswax Server	Internal
8003	Hue	Beeswax Metastore	Internal
8005	Oozie	Server	Internal
8020	HDFS	NameNode	External
8021	MapReduce	JobTracker	External
8080	HBase	REST Service	External
8088	Hue	Server	External
9083	Hive	Metastore	External
9090	HBase	ThriftServer	External
9290	MapReduce	JobTracker Thrift Plugin	Internal
10000	Hive	HiveServer	External
10090	HDFS	Thrift Plugin NameNode	--
10091	MapReduce	JobTracker Authorization Plugin	Internal
10092	HDFS	Authorization Plugin NameNode	--
10094	Oozie	Server Authorization Plugin	Internal
11000	Oozie	Server	External
16000	Sqoop	Metastore	External
35853	Flume	Node	--
35862	Flume	Node	--
35871	Flume	Master	--
35872	Flume	Master	--
45678	Flume	Master	--
50010	HDFS	DataNode	External
50020	HDFS	DataNode	External
50030	MapReduce	JobTracker	External
50060	MapReduce	TaskTracker	External
50070	HDFS	NameNode	External
50075	HDFS	DataNode	External
50090	HDFS	Secondary NameNode	Internal
50470	HDFS	Secure NameNode	External
50495	HDFS	Secure Secondary NameNode	Internal
57890	Flume	Master	--
60000	HBase	Master	External
60010	HBase	Master	External
60020	HBase	RegionServer	External
60030	HBase	RegionServer	External

Using the Oracle Big Data Appliance Configuration Utility

This chapter describes how to use Oracle Big Data Appliance Configuration Utility. It contains the following topics:

- [Overview of Oracle Big Data Appliance Configuration Utility](#)
- [Using Oracle Big Data Appliance Configuration Utility](#)
- [Network Configuration](#)
- [Software Configuration](#)

Note: An Oracle representative typically completes the Oracle Big Data Appliance Configuration Utility from information provided by the customer in the *Oracle Big Data Appliance Configuration Worksheets*.

Overview of Oracle Big Data Appliance Configuration Utility

The information entered in the Oracle Big Data Appliance Configuration Utility is used to generate the installation and deployment files. After you enter values into the utility, the information is validated. Then the utility generates the files for Oracle Big Data Appliance.

The utility is implemented as a spreadsheet containing these worksheets:

- Welcome
- Network Configuration
- Software Configuration
- Preview

Using Oracle Big Data Appliance Configuration Utility

This section explains how to use the Oracle Big Data Appliance Configuration Utility. It contains these topics:

- [Using Oracle Open Office Calc](#)
- [Generating the Configuration Files](#)
- [About the Configuration Files](#)

Using Oracle Open Office Calc

You must use Oracle Open Office Calc 3.3 to configure the spreadsheet. Do not use other spreadsheet programs like Excel, which can modify the spreadsheet but not generate the files.

To obtain the Oracle Big Data Appliance Configuration Utility, refer to My Oracle Support Note 1434477.1.

To open the Oracle Big Data Appliance Configuration Utility:

1. Download Open Office from this website and follow the instructions for installing it.
<http://www.openoffice.org/>
2. Open Oracle Open Office Calc with `bda_configurator.ods`. If the current settings do not allow macros, you see an error message.
3. To change the macro setting:
 - a. From the Tools menu, choose **Options**.
 - b. In the navigation tree, expand the Oracle Open Office folder and select **Security**.
 - c. In the right pane, click **Macro Security**.
 - d. Either lower the security level or enter a trusted source.
 - e. Save your changes and reload the spreadsheet.
4. To turn off AutoCorrect:
 - a. From the Tools menu, choose **AutoCorrect Options**.
 - b. Select the **Options** tab.
 - c. Clear all check boxes, then click **OK**.AutoCorrect settings can make it difficult for you to enter passwords.

Generating the Configuration Files

This procedure describes how to use Oracle Big Data Appliance Configuration Utility to customize the default configuration settings for your installation. The other sections in this chapter provide detailed information on each field in the spreadsheet.

To generate the configuration files:

1. Open the spreadsheet with Oracle Open Office Calc as described in the previous procedure. The Welcome sheet opens.
2. Click **Next** to edit the Network Configuration sheet.
3. Enter the network configuration settings as specified in the *Oracle Big Data Appliance Configuration Worksheets*.
4. Click **Next**. If your changes are invalid, then you see an error message. You must correct the errors before continuing to the Software Configuration sheet.
To discard your changes, click **Reset**.
5. Enter the software configuration settings as specified in the *Oracle Big Data Appliance Configuration Worksheets*.
6. Click **Preview** to view the configuration settings defined by your entries. To make any changes, either click **Previous** or click the tab for a particular sheet. You must

click Preview on the Software Configuration sheet for the changes to appear on the Preview sheet.

7. When you are satisfied with the configuration, click **Save Configuration** on the Preview sheet to generate the configuration files.

This step does not save the worksheet. To save your changes, click the **Save** icon.

8. Choose a location for the configuration files in the Select Path dialog box. The files are created in a directory named `bda-customer_name/rack_name` in that location, such as `bda-Example Inc/bda1`.
9. If you are configuring multiple racks as a single CDH cluster, then you are done. Otherwise, generate separate configuration files for every set of racks that will be configured as a single CDH cluster.
10. Provide the network administrators and others at the site with a copy of the Configuration Templates to verify that the configuration is correct.

About the Configuration Files

The Oracle Big Data Appliance Configuration Utility generates three files to use when configuring the system:

preinstall-checkip.sh

Runs a series of tests to ensure the specified names and IP addresses for Oracle Big Data Appliance were added correctly to the name server and they do not conflict with the existing network configuration. See ["Checking for Network Errors"](#) on page 4-4 for instructions on running this script.

BdaDeploy.json

Contains the network configuration. Copy this file to a USB drive for copying to Oracle Big Data Appliance. See ["Configuring the Network"](#) on page 8-18.

mammoth-rack_name.params

Contains all information provided in the spreadsheet, including the network configuration, plus port numbers, default user names, and passwords. Copy this file to a USB drive for copying to Oracle Big Data Appliance. See ["Using the Mammoth Utility"](#) on page 13-1.

bda-install-template.pdf

Reproduces the Preview page of the spreadsheet in a printable format. This Installation Template is for the network administrator and others to verify the settings and make any last minute corrections.

[Figure 4-1](#) shows part of the first page of the Installation Template.

Figure 4–1 Installation Template

<< Previous		Installation Preview Big Data Appliance Configuration Utility		Reset
Network Configuration				
General				
Cluster Name	bda			
Rack Name	bda1		Primary Appliance	Yes
Country	America			
Timezone	America/New_York			
Domain	example.com			
DNS Servers	198.51.100.45			
NTP Servers	201.0.113.62			
Search Domains	example.com	us.example.com		
Server Network Info				
Type	Netmask	Gateway		
Administrative – eth0	255.255.255.0	203.0.113.1		
Private – bondib0	255.255.255.0			
Client Access – bondet	255.255.255.0	203.0.114.1		

Checking for Network Errors

After providing the host names and IP addresses for the components of Oracle Big Data Appliance, the network administrator must add them to a name server.

Run the `preinstall-checkip.sh` script at these points in the installation process:

- Before Oracle Big Data Appliance arrives at the site but after the network administrator has added the new names and IP addresses to the name server. This step ensures the network configuration parameters are valid and the network is ready to accept Oracle Big Data Appliance connections.
- After Oracle Big Data Appliance arrives but before it is connected to the network to assure that the network has not changed and it is still ready to accept Oracle Big Data Appliance connections.

To check the network for conflicts:

1. Copy the `preinstall-checkip.sh` script to any Linux system with a 10 GB Ethernet connection to the subnet on which Oracle Big Data Appliance will be deployed.
2. Ensure the script is executable:


```
chmod +x preinstall-checkip.sh
```
3. Execute the script:


```
./preinstall-checkip.sh
```
4. If the script identifies errors, open the `bda-checkip.out` file for details. It contains information about every test the script performed. The output file is created in the same directory as the script.
5. Resolve the conflicts if possible before proceeding.

If you are unable to resolve all IP address conflicts between the existing network and Oracle Big Data Appliance, then you can proceed without connecting to the network up to "[Configuring the Network](#)" on page 8-18. The procedures take longer to complete because of time-outs from attempts to connect to configured services such as Domain Name System (DNS) and NTP.

[Example 4-1](#) shows success and error messages from `preinstall-checkip.sh`.

Example 4-1 Output from `preinstall-checkip.sh`

```
Big Data Appliance pre-installation network verification starting ...

Using name server 198.51.100.21 for all DNS lookups

Domain is example.com

Processing DNS name servers : SUCCESS
Processing NTP time servers : SUCCESS
Processing gateways : SUCCESS
Processing factory IP addresses : SUCCESS
Processing public node names (bondeth0) : SUCCESS
Processing private IP addresses (bondib0) : ERROR - see bda-checkip.out for
details
Processing admin node names (eth0) : SUCCESS
Processing ILOM names : SUCCESS
Processing switch and PDU IP addresses : SUCCESS
Processing ASR hostname : SUCCESS

One or more checks report ERROR. Review bda-checkip.out for details
```

Network Configuration

[Table 4-1](#) describes the general rack properties.

Table 4-1 Naming and Locale Information

Spreadsheet Field	Description
Cluster Name	The name of the CDH cluster. Required.
Rack Name	The name of the Oracle Big Data Appliance Rack. By default, the rack name is derived from the cluster name, so that a CDH cluster named <code>bda</code> has racks named <code>bda1</code> , <code>bda2</code> , and so forth.
Server Base Name	Base name for all servers. A two-digit suffix uniquely identifies each individual server. The Rack Name and Server Base Name are used to generate the host names for all network interfaces: <code>eth0</code> , <code>bondib0</code> , <code>bondeth0</code> , and ILOM. For example, a cluster name of <code>bda</code> and a server base name of <code>node</code> results in host names of <code>bda1node01</code> , <code>bda1node02</code> , and so forth.
Admin Access Suffix	Suffix to the basic host name to form the <code>bondeth0</code> host names.
Private Name Suffix	Suffix to the basic host name to form the <code>bondib0</code> host name.
ILOM Name Suffix	Suffix to the basic host name to form the ILOM name.
Switch Base Name	Suffix to the rack name to form the base name for all switches. For example, a rack name of <code>bda</code> and a switch base name of <code>sw</code> results in switch names of <code>bda02sw-ip</code> , <code>bda02sw-ib1</code> , and so forth.
Customer Name	Name of the enterprise. Required.

Table 4–1 (Cont.) Naming and Locale Information

Spreadsheet Field	Description
Domain Name	Name of the domain in which Oracle Big Data Appliance operates. Required.
Region	Region in which the Oracle Big Data Appliance is located. Click the field, then click the arrow for a list of regions.
Time Zone	Time zone for your installation. You must select a region before selecting the time zone. Click the field, then click the arrow for a list of time zones for your region.
Is this the only or primary rack in the cluster?	Yes or no. Required.

Table 4–2 describes the network properties for the entire rack.

Table 4–2 Network Properties

Spreadsheet Field	Description
DNS Servers	Up to four IP addresses for the Domain Name Server.
NTP Servers	Up to four IP addresses for the Network Time Protocol Server.
Search Domains	Up to four domain names in which Oracle Big Data Appliance operates, such as <code>example.com</code> and <code>us.example.com</code> .

Table 4–3 describes the network properties for individual servers in the rack. You specify the IP addresses assigned to server 01. Sequential numbers are assigned automatically to the other 17 servers. Ensure that ranges of 18 IP addresses are free for each type of access. If 18 sequential IP addresses are not available, then you can edit the output files generated by the Oracle Big Data Appliance Configuration Utility.

Table 4–3 Server Properties

Spreadsheet Field	Description
Administrative - eth0	First IP address of 18 from the management network
Private - bondib0	First IP address of 18
Client Access - bondeth0	First IP address of 18 from the client access network
ILOM	First IP address of 18 from the management network

Table 4–4 describes the network properties for all switches. You configure the switches manually, as described in [Chapter 8, "Configuring Oracle Big Data Appliance."](#)

Table 4–4 Switch Properties

Spreadsheet Field	Description
KVM Switch IP	IP address of the KVM switch
Cisco Switch IP	IP address of the Cisco switch
First InfiniBand Switch IP	First IP address of three for the InfiniBand switches

Table 4–5 describes the network properties for the power distribution units (PDUs). You configure the PDUs manually, as described in [Chapter 8, "Configuring Oracle Big Data Appliance."](#)

Table 4–5 PDU Properties

Spreadsheet Field	Description
PDU A	IP address of the first power distribution unit
PDU B	IP address of the second power distribution unit

Software Configuration

[Table 4–6](#) describes the optional software available for Oracle Big Data Appliance. For more information about these components, see the *Oracle Big Data Appliance Software User's Guide*.

Table 4–6 Installed Components

Spreadsheet Field	Description
Are Big Data Connectors licensed?	Oracle Big Data Connectors facilitate data access between data stored in the CDH cluster and Oracle Database. The connectors require a separate license. If you have a license, choose Yes .
Install Oracle Data Integrator Agent?	The agent is a required component of Oracle Data Integrator, which is a graphical tool for performing extraction, transformation, and loading (ETL) of data into Oracle Database. To configure Oracle Data Integrator for use immediately, choose Yes .
Install Oracle NoSQL Database Community Edition?	Oracle NoSQL Database is a distributed key-value database that provides a storage option to Hadoop Distributed File System (HDFS). To configure Oracle NoSQL Database for use, choose Yes .
Total disk space to allocate for Oracle NoSQL Database in Terabytes	If you install Oracle NoSQL Database, you can allocate 54 or 108 TB for its use. Disk space allocated to Oracle NoSQL Database is not available for HDFS.
External backup NFS directory	The address of an NFS shared directory used to store a backup of the name node data. Enter the address in the form <code>nfs-host:/remote/shared/path</code> , where <code>nfs-host</code> is the fully qualified host name and <code>/remote/shared/path</code> is the directory path. The <code>hdfs</code> user must have write permission on this directory. See Step 11, "SetupMountPoints" for details on setting up this directory.

[Table 4–7](#) describes the options for Auto Service Request (ASR). This service monitors the health of Oracle Big Data Appliance hardware. ASR Manager automatically submits a service request to Oracle Support Services when it detects a fault. See [Chapter 14, "Maintaining Oracle Big Data Appliance,"](#) for more information about ASR.

Table 4–7 Oracle Auto Service Request

Spreadsheet Field	Description
Enable Auto Service Request?	Although you can opt out of this service, Oracle strongly recommends that you enable ASR by choosing Yes .
ASR Manager Host	The fully qualified name of a Linux server on the network, where ASR will be installed.
ASR Manager Port	The port number for ASR Manager. The default port is 162.
ASR Server Root Password	Password for <code>root</code> on the ASR Manager Host.

Table 4–8 identifies the users, groups, and passwords that can be set during the software installation. If you leave the passwords blank, then the installation prompts for them. The Oracle IDs must match those of a connected Oracle Exadata Database Machine to support Network File System (NFS) protocol.

Table 4–8 Users /Groups

Spreadsheet Field	Description
Cloudera Manager admin password	The password for the admin user for Cloudera Manager.
root operating system password	The root password on all servers in Oracle Big Data Appliance.
oracle operating system password	The oracle password on all servers in Oracle Big Data Appliance. Oracle applications run under this identity.
oracle user ID	The ID number of the oracle user. It must match the oracle user ID of a connected Oracle Exadata Database Machine.
oinstall group ID	The ID number of the Oracle Inventory Group (oinstall). It must match the oinstall group ID of a connected Oracle Exadata Database Machine.
dba group ID	The ID number of the dba group. It must match the dba group ID of a connected Oracle Exadata Database Machine.
MySQL administration password	The password for the MySQL Database administrative user.
MySQL password for Oracle Data Integrator	The password for the Oracle Data Integrator user in MySQL Database.

Table 4–9 describes the configuration settings for the email server that Cloudera Manager uses to send alerts from the CDH cluster.

Table 4–9 Cloudera Manager Email Alerts

Spreadsheet Field	Description
Email Server (SMTP) Host	The fully qualified name of the existing SMTP server the company uses on its internal network.
Email Server (SMTP) Port	The port used by the email server.
Email Server (SMTP) User Name	User name for Cloudera Manager to use for authentication on the SMTP server.
Email Server (SMTP) Password	Password for the user name.
Email Server uses SSL?	Yes or no.
Email Alert Recipients	One or more email addresses. These users receive the alerts from Cloudera Manager.

Installing Oracle Big Data Appliance at the Site

This chapter describes how to move, install and configure the hardware through to powering on the system. This chapter contains these topics:

- [Reviewing Safety Guidelines](#)
- [Unpacking Oracle Big Data Appliance](#)
- [Placing Oracle Big Data Appliance in Its Allocated Space](#)
- [Powering On the System the First Time](#)

Reviewing Safety Guidelines

Before Oracle Big Data Appliance arrives, review the following safety precautions to ensure the site is safe and ready for delivery. Failing to observe these precautions can result in injury, equipment damage, or malfunction.

- Do not block ventilation openings.
- Do not install Oracle Big Data Appliance in a location that is exposed to direct sunlight or near a device that may become hot.
- Do not install Oracle Big Data Appliance in a location that is exposed to excessive dust, corrosive gases, or air with high salt concentrations.
- Do not install Oracle Big Data Appliance in a location that is exposed to frequent vibrations. Install it on a flat, level surface.
- Use a power outlet that uses proper grounding. When using shared grounding, the grounding resistance must not be greater than 10 ohms. Ensure that your facility administrator or a qualified electrical engineer verifies the grounding method for the building and performs the grounding work.
- Be sure that each grounding wire used for Oracle Big Data Appliance is used exclusively for Oracle Big Data Appliance. Also be sure to observe the precautions, warnings, and notes about handling that appear on labels on the equipment.
- Do not place cables under the equipment or stretch the cables too tightly.
- Do not disconnect power cords from the equipment while its power is on.
- If you cannot reach the connector lock when disconnecting LAN cables, then press the connector lock with a flathead screwdriver to disconnect the cable. You could damage the system board if you force your fingers into the gap rather than using a flathead screwdriver.

- Do not place anything on top of Oracle Big Data Appliance or perform any work directly above it.
- Do not let the room temperature rise sharply, especially in winter. Sudden temperature changes can cause condensation to form inside Oracle Big Data Appliance. Allow for a sufficient warm-up period prior to operation.
- Do not install Oracle Big Data Appliance near a photocopier, air conditioner, welding machine, or any other equipment that generates loud, electronic noises.
- Avoid static electricity at the installation location. Static electricity transferred to Oracle Big Data Appliance can cause malfunctions. Static electricity is often generated on carpets.
- Confirm the supply voltage and frequency match the electrical ratings indicated for Oracle Big Data Appliance.
- Do not insert anything into any Oracle Big Data Appliance opening, unless doing so is part of a documented procedure. Oracle Big Data Appliance contains high-voltage parts. If a metal object or other electrically-conductive object enters an opening in Oracle Big Data Appliance, 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 and available online at <http://docs.oracle.com/cd/E19115-01/mod.dc.d20/816-7190-12/816-7190-12.pdf>
- *Oracle Big Data Appliance Safety and Compliance Guide* and all safety notices printed on the packaging.
- *Sun Rack II Power Distribution Units User's Guide* at <http://docs.oracle.com/cd/E19844-01/>

Unpacking Oracle Big Data Appliance

You should determine the unpacking location during the site planning process. See [Chapter 2, "Site Requirements."](#)

[Figure 5–1](#) shows the Oracle Big Data Appliance shipping crate.

Figure 5–1 Oracle Big Data Appliance in Shipping Crate

See Also: *Sun Rack II User's Guide* for additional information about unpacking the rack. The guide is available at

<http://download.oracle.com/docs/cd/E19844-01/index.html>

Contents of the Shipping Kit

The shipping kit contains the following tools and equipment to install and service Oracle Big Data Appliance:

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

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

Note: The following items are not included in the shipping kit:

- No. 2 Phillips screwdriver
 - Antistatic wrist strap
-
-

Unpacking Oracle Big Data Appliance

The following procedure describes how to unpack Oracle Big Data Appliance from the shipping crate.

Note: Oracle strongly recommends that you use professional movers when unpacking and installing Oracle Big Data Appliance.

See Also: Sun Rack II unpacking instructions included with the packaging

1. Unpack Oracle Big Data Appliance carefully from the packaging and shipping pallet:
 - a. Remove the shipping carton bands.
 - b. Remove the yellow fasteners and carton top.
 - c. Remove the carton sides and inner top.

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

2. Remove the shipping kit.
3. Attach the ramps to the shipping pallet as follows:
 - a. Remove the ramps from the pallet sides.
 - b. Obtain the parts bag from inside the cabinet.
 - c. Adjust the leveling bolts on the ramps and connect the ramps to the pallet wheel track.
4. Carefully roll Oracle Big Data Appliance off the shipping pallet as follows:
 - a. Unfasten the exterior mounting brackets from the pallet.
 - b. Unfasten the interior mounting brackets from the pallet. Use care when removing the mounting brackets from underneath Oracle Big Data Appliance. Access to the inside mounting brackets might be limited.
 - c. Roll the cabinet down the ramps to the level floor. Oracle recommends having three people available to move the rack down the ramp: two people on both sides to help guide the rack and one person in back.
5. Save the mounting brackets that secure the rack to the shipping pallet.

You can use these mounting brackets to permanently secure Oracle Big Data Appliance to the installation site floor. Do not dispose of these brackets, because you cannot order replacement brackets.

6. Recycle the packaging properly. Follow local laws and guidelines to dispose of the material.

Placing Oracle Big Data Appliance in Its Allocated Space

This section describes how to position, stabilize, and ground Oracle Big Data Appliance. This section contains the following topics:

- [Moving Oracle Big Data Appliance](#)
- [Stabilizing Oracle Big Data Appliance](#)
- [Attaching a Ground Cable \(Optional\)](#)


Moving Oracle Big Data Appliance

The following procedure describes how to move Oracle Big Data Appliance:

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

When moving Oracle Big Data Appliance to the installation site, the front casters do not turn; you must steer the unit by moving the rear casters. You can safely maneuver Oracle Big Data Appliance by carefully pushing it.

Figure 5–2 Carefully Push Oracle Big Data Appliance from Behind



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

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

WARNINGS:

- **Never attempt to move Oracle Big Data Appliance by pushing on the side panels. Pushing on the side panels can tip the rack over. This can cause serious personal injury or death, and also damage to the equipment.**
 - **Never tip or rock Oracle Big Data Appliance because the rack can fall over.**
-
-

Stabilizing Oracle Big Data Appliance

After moving Oracle Big Data Appliance to the installation site, stabilize the rack to ensure that it does not move or tip over. You can stabilize the rack permanently by extending the rack leveling feet, using mounting brackets, or both.


Stabilize Oracle Big Data Appliance Rack with Leveling Feet

The rack contains four leveling feet that you can lower to stabilize the rack, even when the rack is permanently secured to the floor.

To adjust the leveling feet:

1. Locate the four leveling feet located at the bottom corners of Oracle Big Data Appliance Rack.
2. Lower the leveling feet to the floor as shown in [Figure 5-3](#) using the SW 12 mm wrench. When lowered correctly, the four leveling feet should support the full weight of Oracle Big Data Appliance Rack.


Figure 5-3 *Securing Oracle Big Data Appliance Rack Using the Leveling Feet*




Stabilize Oracle Big Data Appliance Rack with Mounting Brackets

The rack can be permanently mounted to the installation site floor using the same four mounting brackets that secured the rack to the shipping pallet.

1. Drill the mounting holes in the installation floor. [Figure 5-4](#) shows the location and dimensions of the mounting brackets.

Figure 5-4 Location of Mounting Brackets on Rack

2. Obtain four bolts and washers to mount Oracle Big Data Appliance Rack to the floor. The bolt holes in the mounting brackets have a 10.0 mm diameter. Oracle does not provide mounting bolts because different floors require different bolt types and strengths. Select bolts that are appropriate for your location.
3. Position Oracle Big Data Appliance Rack over the pre-drilled holes. [Figure 5-5](#) shows the bottom view of Oracle Big Data Appliance Rack and the location for the mounting hole and floor cut dimensions.

Figure 5-5 Bottom View of Oracle Big Data Appliance Rack

- 1: Distance from mounting bracket to the edge of the rack is 113 mm (4.45 inches)
- 2: Width between the centers of the mounting hole slots is 374 mm (14.72 inches)


- 3: Distance between mounting bracket to the edge of the rack is 113 mm (4.45 inches)
- 4: Distance between the centers of the front and rear mounting hole slots is 1120 mm (44.1 inches)
- 5: Depth of cable-routing floor cutout is 330 mm (13 inches)
- 6: Distance between the floor cutout and the edge of the rack is 160 mm (6.3 inches)
- 7: Width of cable-routing floor cutout is 280 mm (11 inches)

If you plan to route data or PDU power cords down through the bottom of the rack, then you must cut a hole in the installation floor site. Cut a rectangle below the rear portion of the rack, between the two rear casters and behind the rear RETMA (Radio Electronics Television Manufacturers Association) rails.


Caution: Do not create a hole where the rack casters or leveling feet brackets will be placed.

Figure 5–6 shows the base positions of the rack from the bottom.

Figure 5–6 Base Position to Measure



4. Open the front and rear doors.
5. Install the mounting brackets to the rack as shown in Figure 5–7 using a 6 mm hexagon Allen wrench key.

Figure 5-7 Securing Oracle Big Data Appliance Rack Using Mounting Brackets

6. Using bolts and washers obtained in step 2, permanently mount your system to the floor using the four mounting brackets as shown in [Figure 5-7](#).
7. Firmly tighten all of the bolts that secure the mounting brackets to Oracle Big Data Appliance Rack and to the floor.


Attaching a Ground Cable (Optional)

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

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

To attach a ground cable:

1. Obtain a grounding cable. It is not shipped with the system.
2. Ensure the installation site has properly grounded the power source in the data center. The facility PDU must have earth ground.
3. Ensure all grounding points, such as raised floors and power receptacles, reference the facility ground.
4. Ensure that direct, metal-to-metal contact is made for this installation. During manufacturing, the ground cable attachment area might have been painted or coated.
5. Attach the ground cable to an attachment point located at the bottom rear of the system frame, as shown in [Figure 5-8](#). The attachment point is an adjustable bolt that is inside the rear of the cabinet on the right side.

Figure 5–8 Earth Ground Attachment Bolt Location

Powering On the System the First Time

Before powering on the system for the first time, you must inspect it and connect the power cords. This section contains the following topics:

- [Inspecting the Rack After It Is in Place](#)
- [Connecting Power Cords](#)
- [Powering On Oracle Big Data Appliance](#)

Note: After powering on Oracle Big Data Appliance, configure the system as described in [Chapter 8, "Configuring Oracle Big Data Appliance."](#)

Inspecting the Rack After It Is in Place

This procedure describes how to visually examine the Oracle Big Data Appliance physical system after it is in place, but before power is supplied.

To inspect Oracle Big Data Appliance Rack:

1. Check the rack for damage.
2. Check the rack for loose or missing screws.
3. Check Oracle Big Data Appliance Rack for the ordered configuration. Refer to the Customer Information Sheet (CIS) on the side of the packaging.
4. Check that all cable connections are secure and firmly in place as follows:
 - a. Check the power cables. Ensure that the correct connectors have been supplied for the data center facility power source.
 - b. Check the network data cables.
5. Check the site location tile arrangement for cable access and airflow.
6. Check the data center airflow that leads into the front of Oracle Big Data Appliance Rack.

See Also: ["Ventilation and Cooling Requirements"](#) on page 2-9 for more information

Connecting Power Cords

This procedure describes how to connect power cords to Oracle Big Data Appliance Rack.

To connect power cords to Oracle Big Data Appliance Rack:

1. Open the rear cabinet door.
2. Ensure that the correct power connectors have been supplied.
3. Unfasten the power cord cable ties. The ties are for shipping only and are no longer needed. See [Figure 5-9](#).
4. Route the power cords to the facility receptacles as shown in [Figure 5-10](#) or [Figure 5-11](#). You can route the power cords out the top or the bottom of the rack.
5. Secure the power cords in bundles as shown in [Figure 5-10](#).
6. Ensure the breaker switches are OFF before connecting the power cables.
7. Plug the PDU power cord connectors into the facility receptacles.

Figure 5-9 Location of Power Cables




Figure 5–10 Power Cord Routing Out the Bottom





Figure 5–11 Power Cord Routing Out the Top




Powering On Oracle Big Data Appliance

Note: Oracle Big Data Appliance nodes may take up to 15 minutes to start through the normal BIOS POST tests.

To power on Oracle Big Data Appliance:

1. Switch on the PDU B circuit breakers, one at a time. PDU B is on the right side of the rack when viewed from the rear. The circuit breakers are on the rear of the Oracle Big Data Appliance Rack cabinet as shown in [Figure 5-12](#). Press the ON (I) side of the toggle switch. When the breaker is ON, the switch is flush with the side of the PDU.

Figure 5-12 PDU Switch Locations

2. Verify the expected power LEDs are on. The LEDs are located as follows:
 - Sun Fire Servers: Top LED
 - Cisco switch: Left LED (viewed from front) is green, and the other LEDs are red
 - InfiniBand switches: Right LED (viewed from front) labeled PS1
 - KVM switch: Lower LED B
 - KMM switch: Only supplied by PDU B
3. Connect power to PDU A. Ensure the breaker switches are in the OFF position before connecting the power cables.

See Also: "[Reviewing Safety Guidelines](#)" on page 5-1 for information about single phase PDUs and cabling

4. On single-phase systems, ensure the following inputs are on the same phase:
 - PDU A Input 2 and PDU B Input 0
 - PDU A Input 1 and PDU B Input 1
 - PDU A Input 0 and PDU B Input 2
5. Switch on the PDU A circuit breakers one at a time.
6. Verify the expected power LEDs are on throughout the rack.
7. Perform a visual check of all cable connections in the rack. Do not press every connector to verify connection.

8. Verify the OK LED is blinking standby for all systems. The OK LED blinks every 3 seconds when in standby mode. It is dark for 2 to 3 minutes while the ILOM is booting before going to standby mode.

Each time the system restarts, a file is generated in the root directory with the results of the validation check. The file is named either `BDA_REBOOT_SUCCEEDED` or `BDA_REBOOT_FAILED`, depending on the success or failure of the validation checks.

Using Oracle Integrated Lights Out Manager

Oracle Integrated Lights Out Manager (ILOM) is a preinstalled component of Oracle Big Data Appliance that you can use to monitor the servers and switches. You use ILOM in [Chapter 8, "Configuring Oracle Big Data Appliance."](#)

This chapter contains the following topics:

- [Oracle ILOM Overview](#)
- [Administrative Network Diagram](#)
- [ILOM IP Addresses for Oracle Big Data Appliance Rack Components](#)
- [Connecting to ILOM Using the Network](#)
- [Connecting to ILOM Using a Serial Connection](#)

See Also: Oracle Integrated Lights Out Manager 3.0 Documentation Library at

<http://docs.oracle.com/cd/E19860-01/>

Oracle ILOM Overview

Oracle Integrated Lights Out Manager (Oracle ILOM) provides advanced service processor (SP) hardware and software that you can use to manage and monitor the servers and switches in Oracle Big Data Appliance Rack. Oracle ILOM dedicated hardware and software is preinstalled on these components. ILOM automatically initializes as soon as power is applied.

Oracle ILOM enables you to actively manage and monitor servers in Oracle Big Data Appliance regardless of the operating system state, providing you with a reliable Lights Out Management (LOM) system.

With ILOM, you can proactively:

- Learn about hardware errors and faults as they occur
- Remotely control the power state of a server
- View the graphical and non-graphical consoles
- View the current status of sensors and indicators on the system
- Determine the hardware configuration of your system
- Receive generated alerts about system events in advance

The ILOM service processor runs its own embedded operating system and has a dedicated Ethernet port, which together provide out-of-band management capability. In addition, you can access ILOM from the server operating system (Oracle Linux).

Using ILOM, you can remotely manage Oracle Big Data Appliance as if you were using the local KVM.

Oracle ILOM Interfaces

Oracle ILOM supports two interfaces for accessing its features and functions. You can choose to use a browser-based web interface or a command-line interface.

Web Interface

The web interface enables you use a browser to log in to the SP, then to perform system management and monitoring.

Command-Line Interface

The command-line interface (CLI) enables you to operate ILOM using keyboard commands and adheres to industry-standard DMTF-style CLI and scripting protocols. ILOM supports SSH v2.0 and v3.0 for secure access to the CLI. Using the CLI, you can reuse existing scripts and automate tasks using familiar interfaces.

Oracle ILOM Users

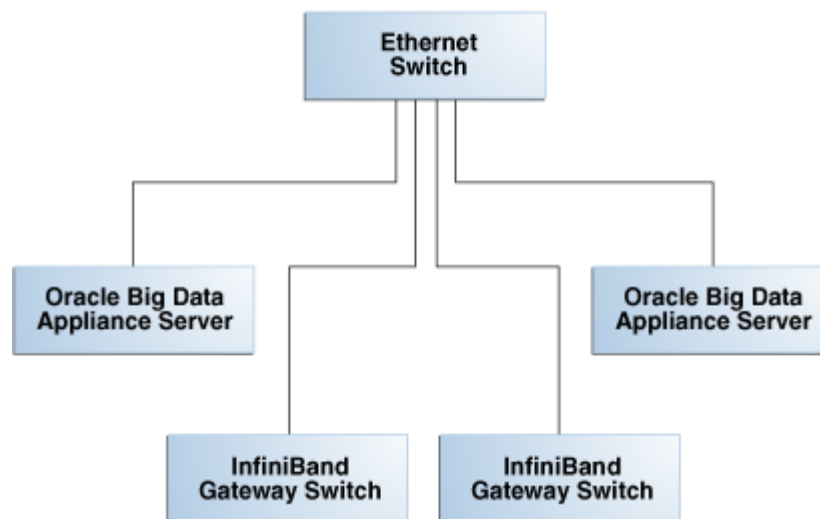
Oracle ILOM on Oracle Big Data Appliance is configured with these two users:

- `ilom-admin`: Provides the write privileges needed to change the system configuration, without the full power of `root`.
- `ilom-operator`: Provides read-only privileges to monitor the system.

Administrative Network Diagram

Figure 6-1 illustrates the management/ILOM network. It shows two of the 18 servers and the two Sun Network QDR InfiniBand Gateway Switches. The Cisco Ethernet management switch is connected to the servers and the InfiniBand switches.

Figure 6-1 Administrative Network in Oracle Big Data Appliance



ILOM IP Addresses for Oracle Big Data Appliance Rack Components

You require at least one static IP address for Service Processor (ILOM access). For the list of default ILOM IP addresses assigned to Oracle Big Data Appliance Rack components at the time of manufacturing, see ["Factory Network Settings"](#) on page 3-5.

After you reconfigure these IP addresses during the initial configuration of the Oracle Big Data Appliance using Oracle Big Data Appliance Configuration Utility, you must use those IP addresses to access ILOM.

Connecting to ILOM Using the Network

You can typically access ILOM using the network over an Ethernet connection. You must know the ILOM Ethernet address. Before system configuration, the address is the factory IP address setting. After system configuration, you can use either the component name or IP address listed in the Configuration Worksheet. You can use either the CLI or the browser interface to access ILOM. Alternatively, you can launch a remote KVM session to access ILOM.

Note: You can use this connection method when ILOM IP addresses can be accessed over the network. Oracle recommends that you use this connection method to access ILOM. See [Chapter 8, "Configuring Oracle Big Data Appliance."](#)

This section discusses the following topics:

- [Connecting to the Command-Line Interface](#)
- [Connecting to the Web GUI](#)
- [Opening a Remote KVM Session](#)

Connecting to the Command-Line Interface

1. Open an SSH client, such as PuTTY.
2. In the **Host Name (or IP address)** field, enter the ILOM name or IP address of the component to connect to. For example, before configuration, you might enter the factory default IP address of 192.168.1.203 for a Sun Network QDR InfiniBand Gateway Switch. After configuration, you might enter the new IP address or a name such as bda1sw-ib2.
3. Ensure that **SSH** is chosen as the **Connection Type**.
4. Type the user name and the password, when prompted. The default user name is `root`, and the default password is `welcome1`.
The CLI command prompt (`#`) is displayed.
5. Enter `help` for a list of commands.

Connecting to the Web GUI

1. Open a web browser and navigate to the ILOM name or IP address of an Oracle Big Data Appliance server or switch.
The login screen is displayed.

2. Enter the user name and the password. The administrative user name is `ilom-admin`, and the default password is `welcome1`.

Note: Use the `passwd` command to change the password.

3. Click **Log In**.

The web GUI is displayed, as shown in [Figure 6–2](#).

Figure 6–2 Oracle ILOM Web Interface

The screenshot shows the Oracle Integrated Lights Out Manager (ILOM) web interface. The header includes the Oracle logo and the text 'Oracle® Integrated Lights Out Manager'. A navigation menu contains tabs for System Information, System Monitoring, Configuration, User Management, Maintenance, and Switch/Fabric Monitoring Tools. Below the menu, there are sub-tabs for SUN DCS GW Firmware Versions, ILOM Versions, Session Time-Out, Components, and Identification Information. The main content area is titled 'SUN DCS GW Firmware Versions' and contains a table with the following data:

Property	Value
GW FW version	2.0.5-2
Build time	Nov 29 2011 16:05:05
Last FW upgrade	2011-12-18 02:32:59 +0000
FW upgrade status	SUCCESS

Opening a Remote KVM Session

To access the ILOM consoles for Oracle Big Data Appliance components that are connected to ILOM:

1. Ensure that pop-up blockers are disabled in your browser before you launch the remote console.
2. Type the ILOM IP address in the address bar of a web browser.
3. Log in to ILOM using the administrative user name and password (`ilom-admin` and `welcome1`). After login, the ILOM home page is displayed.
4. Click the **Remote Control** tab, and click **Launch Remote Console**.

The remote console is displayed.

Connecting to ILOM Using a Serial Connection

You can connect to ILOM using a serial connection if you are unable to access ILOM using the network due to any of the following problems:

- Misconfiguration of the network
- Misconfiguration of ILOM IP addresses
- Misconfiguration of Cisco Ethernet switch ports
- Sideband configuration issues

Connecting to a Server ILOM

To connect to ILOM of an Oracle Big Data Appliance server using a serial connection:

1. Attach a serial cable from a terminal or a PC running terminal emulation software to the SER MGT port of an Oracle Big Data Appliance server. The cable should be 15 feet or shorter.
2. Verify that your terminal or laptop is operational.
3. Configure the terminal device or the terminal emulation software to use the following settings:
 - 8N1: eight data bits, no parity, one stop bit
 - 9600 baud (default, but can be set to any standard rate up to 115200)
 - Disable software flow control (XON/XOFF)
 - Disable hardware control
4. Verify that power is supplied to either PSU.

If there is power applied to either PSU, then ILOM will be functional regardless of the power state of the servers.

5. Press Enter on the terminal device. A connection between the terminal device and the ILOM is established.

The ILOM login prompt is displayed.

6. Log in to the CLI using the administrative user name and the password (ilom-admin and welcome1).

The ILOM default command prompt is displayed.

Connecting to ILOM of a Sun Network QDR InfiniBand Gateway Switch

To connect to ILOM of a Sun Network QDR InfiniBand Gateway Switch using a serial connection:

1. Attach a USB-to-Serial connector to the USB port of the switch.
2. Verify that your terminal or laptop is operational.
3. Configure the terminal device or the terminal emulation software to use the following settings:
 - 8N1: eight data bits, no parity, one stop bit
 - 115200 baud
 - Disable software flow control (XON/XOFF)
 - Disable hardware control

4. Press the Return or Enter key on the serial device several times to synchronize the connection.

You may see text similar to the following:

```
where nm2name is the host name of the management controller. The name  
might be the word hostname.
```

Even if you do not see the text, go to Step 5.

5. Type root for the login name followed by the root password of welcome1. The # prompt is displayed.

Note: If you do not see this output or prompt, there is a problem with the serial configuration, the USB-to-Serial connector, or the CLI.

Using ipmitool Commands When SP Network Information Is Lost

If the service processor (SP) network information gets lost, you can use an SSH shell to connect to the ETH0 interface of the operating system on a server and issue `ipmitool` commands to reconfigure network settings. For example, this command resets the ILOM root password:

```
ipmitool sunoem cli 'set /SP/users/rootpassword=somepasswd' password
```

For a list of `ipmitool` commands, issue the command without arguments:

```
ipmitool
```

Note: For more information about `ipmitool` commands and options, see:

<http://ipmitool.sourceforge.net/manpage.html>

Using the dcli Utility

The `dcli` utility executes commands across a group of servers on Oracle Big Data Appliance and returns the output. You use `dcli` in [Chapter 8, "Configuring Oracle Big Data Appliance."](#)

This chapter contains the following topics:

- [Overview of the dcli Utility](#)
- [dcli Syntax](#)
- [Return Values](#)
- [Examples](#)

Overview of the dcli Utility

The `dcli` utility executes commands on multiple Oracle Big Data Appliance servers in parallel, using the InfiniBand (`bondib0`) interface to make the connections. You can run the utility from any server.

Setting Up Passwordless SSH

The `dcli` utility requires a passwordless Secure Shell (SSH) between the local server and all target servers. You run the `dcli` utility on the local server, and the commands specified in `dcli` execute on the target servers.

Two scripts facilitate the use of SSH on Oracle Big Data Appliance: `setup-root-ssh` and `remove-root-ssh`. These utilities accept two options that are also used by `dcli`:

- `-C`: Targets all the servers in a Hadoop cluster.
- `-g`: Targets a user-defined set of servers.

See [Table 7-1, "dcli Options"](#) for details about these options.

To set up passwordless SSH for root:

1. Connect to a Oracle Big Data Appliance server using PuTTY or a similar utility. Select an SSH connection type.
2. Log in as `root`.
3. Set up passwordless SSH for `root` across the rack:

```
setup-root-ssh
```

Or, to set up passwordless SSH across a Hadoop cluster of multiple racks:

```
setup-root-ssh -C
```

You see the message "ssh key added" from each server.

4. You can now run any ssh command on any server in the rack without entering a password. In addition to `dcli` commands, you can use `scp` to copy files between servers.
5. To remove passwordless SSH from root:

```
remove-root-ssh
```

Basic Use of dcli

This topic identifies some of the basic options to the `dcli` command. For a complete list of options, see [Table 7-1, "dcli Options"](#).

Getting Help

To see the `dcli` help page, issue the `dcli` command with the `-h` or `--help` options. You can see a description of the commands by issuing the `dcli` command with no options.

Identifying the Target Servers

You can identify the servers where you want the commands to run either in the command line or in a file. For a list of default target servers, use the `-t` option. To change the target servers for the current command, use the `-c` or `-g` options described in [Table 7-1, "dcli Options"](#).

The `/opt/oracle/bda` directory contains two files for executing commands on multiple servers:

- `rack-hosts-infiniband` is the default target group of servers for the `dcli`, `setup-root-ssh`, and `remove-root-ssh` utilities. The file initially contains the default factory IP addresses. The network configuration process changes this file to the custom IP addresses identified in the Oracle Big Data Appliance Configuration Worksheets.
- `cluster-hosts-infiniband` contains the names of all servers in the Hadoop cluster created by the Mammoth Utility. A cluster can span one or more Oracle Big Data Appliance racks.

You can manually create additional files with groups of servers that you want to manage together. For example, you might manage servers 5 to 18 together, because they have no special functions like servers 1 to 4.

Specifying the Commands

You typically specify a command for execution on the target servers on the command line. However, you can also create command file for a series of commands that you often issue together or for commands with complex syntax. See the `-x` option.

You can also copy files to the target servers without executing them by using the `-f` option.

Controlling the Output Levels

You can request more information with the `-v` option or less information with the `-n` option. You can also limit the number of returned lines with the `--maxlines` option, or replace matching strings with the `-r` option.

Following are examples of various output levels using a simple example: the Linux `date` command.

Note: The output from only one server (node07) is shown for comparison. The syntax shown in these examples issue the `date` command on all 18 servers.

This is the default output, which lists the server followed by the output.

```
# dcli date
bdalnode07-adm.us.oracle.com: Tue Feb 14 10:22:31 PST 2012
```

The minimal output returns OK for completed execution:

```
# dcli -n date
OK: ['bdalnode07.us.oracle.com']
```

Verbose output provides extensive information about the settings under which the command ran:

```
dcli -v date
options.nodes: None
options.destfile: None
options.file: None
options.group: dcserverns
options.maxLines: 100000
options.listNegatives: False
options.pushKey: False
options.regexp: None
options.sshOptions: None
options.scpOptions: None
options.dropKey: False
options.serializeOps: False
options.userID: root
options.verbosity 1
options.vmsstatOps None
options.execfile: None
argv: ['/opt/oracle/bda/bin/dcli', '-g', 'dcserverns', '-v', 'date']
Success connecting to nodes: ['bdalnode07.us.oracle.com']
...entering thread for bdalnode07.us.oracle.com:
execute: /usr/bin/ssh -l root bdalnode07.us.oracle.com 'date'
...exiting thread for bdalnode07.us.oracle.com status: 0
bdalnode07.us.oracle.com: Tue Feb 14 10:24:43 PST 2012]
```

dcli Syntax

```
dcli [options] [command]
```

Parameters

options

Options described in [Table 7-1](#). You can omit all options to run a command on all servers.

command

Any command that runs from the operating system prompt. If the command contains punctuation marks or special characters, then enclose the command in double quotes.

The backslash (\) is the escape character. Precede the following special characters with a backslash on the command line to prevent interpretation by the shell. The backslash is not needed in a command file. See the `-x` option.

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

Table 7-1 dcli Options

Option	Description
-c <i>nodes</i>	Specifies a comma-separated list of Oracle Big Data Appliance servers where the command is executed.
-C	Uses the list of servers in <code>/opt/oracle/bda/cluster-rack-infiniband</code> as the target. See "Identifying the Target Servers" on page 7-2.
-d <i>destfile</i>	Specifies a target directory or file name for the -f option.
-f <i>file</i>	Specifies files to be copied to the user's home directory on the target servers. The files are not executed. See the -l option.
-g <i>groupfile</i>	Specifies a file containing a list of Oracle Big Data Appliance servers where the command is executed. Either server names or IP addresses can be used in the file.
-h, --help	Displays a description of the commands.
-k	Pushes the ssh key to each server's <code>/root/.ssh/authorized_keys</code> file. See "Setting Up Passwordless SSH" on page 7-1 for an easier alternative.
-l <i>userid</i>	Identifies the user ID for logging in to another server. The default ID is <code>root</code> .
--maxlines= <i>maxlines</i>	Identifies the maximum lines of output displayed from a command executed on multiple servers. The default is 10,000 lines.
-n	Abbreviates the output for non-error messages. Only the server name is displayed when a server returns normal output (return code 0). You cannot use the -n and -r options together.
-r <i>regexp</i>	Replaces the output with the server name for lines that match the specified regular expression.
-s <i>sshoptions</i>	Specifies a string of options that are passed to SSH.
--scp= <i>scpoptions</i>	Specifies a string of options that are passed to Secure Copy (SCP), when these options are different from <i>sshoptions</i> .
--serial	Serializes execution over the servers. The default is parallel execution.
-t	Lists the target servers.
--unkey	Drops the keys from the <code>authorized_key</code> files of the target servers.
-v	Displays the verbose version of all messages.
--version	Displays the dcli version number.

Table 7-1 (Cont.) dcli Options

Option	Description
<code>--vmstat=VMSTATOPS</code>	<p>Displays the syntax of the Linux Virtual Memory Statistics utility (<code>vmstat</code>). This command returns process, virtual memory, disk, trap, and CPU activity information.</p> <p>To issue a <code>vmstat</code> command, enclose its options in quotes. For example:</p> <pre>--vmstat="-a 3 5"</pre> <p>See your Linux documentation for more information about <code>vmstat</code>.</p>
<code>-x <i>execfile</i></code>	Specifies a command file to be copied to the user's home directory and executed on the target servers. See the <code>-l</code> option.

Return Values

- 0: The command ran successfully on all servers.
- 1: One or more servers were inaccessible or remote execution returned a nonzero value. You get a message listing the unresponsive servers. Execution continues on the other servers.
- 2: A local error prevented the command from executing.

If you interrupt the local `dcli` process, then the remote commands may continue without returning their output or status.

Examples

Following are examples of the `dcli` utility.

This example returns the default list of target servers:

```
# dcli -t
Target nodes: ['bdalnode01-adm.us.oracle.com', 'bdalnode02-adm.us.oracle.com',
'bdalnode03-adm.us.oracle.com', 'bdalnode04-adm.us.oracle.com',
'bdalnode05-adm.us.oracle.com', 'bdalnode06-adm.us.oracle.com',
'bdalnode07-adm.us.oracle.com', 'bdalnode08-adm.us.oracle.com',
'bdalnode09-adm.us.oracle.com']
```

The next example checks the temperature of all servers:

```
# dcli 'ipmitool sunoem cli "show /SYS/T_AMB" | grep value'

bdalnode01-adm.us.oracle.com: value = 22.000 degree C
bdalnode02-adm.us.oracle.com: value = 22.000 degree C
bdalnode03-adm.us.oracle.com: value = 22.000 degree C
bdalnode04-adm.us.oracle.com: value = 23.000 degree C
.
.
.
```

Configuring Oracle Big Data Appliance

This chapter describes how to configure the system, accounts, and software for Oracle Big Data Appliance. Many of the procedures in this chapter use Oracle Integrated Lights Out Manager (ILOM) and the `dcli` utility.

This chapter contains the following topics:

- [Configuring the KVM Switch](#)
- [Configuring the Cisco Ethernet Switch](#)
- [Configuring the InfiniBand Switches](#)
- [Configuring the Power Distribution Units](#)
- [Checking Oracle Big Data Appliance Servers](#)
- [Configuring the Network](#)
- [Reinstalling the Base Image](#)
- [Checking the Health of the Network](#)

See Also:

- [Chapter 6, "Using Oracle Integrated Lights Out Manager"](#)
- [Chapter 7, "Using the dcli Utility"](#)

Note:

- You must have the Installation Template to complete these procedures. It identifies the IP addresses and names of the servers and switches required by the manual configuration steps.
 - Oracle recommends that an Oracle Field Engineer and an Oracle Service Representative perform the configuration steps described in this chapter.
-
-

Configuring the KVM Switch

This topic contains procedures for configuring the KVM switch:

- [Starting the KVM Switch](#)
- [Connecting the KVM to the Management Network](#)
- [Checking the KVM Firmware Version](#)

Starting the KVM Switch

1. Ensure that all connected components are powered off.
2. Pull the KVM tray out from the front of the rack, and open it using the handle.
3. Touch the touch pad.
4. Toggle between the host and KVM interface by pressing the Ctrl key on the left side twice, similar to double-clicking with a mouse.
5. Select **Target Devices** from the Unit View of the user interface. The number of sessions shown should be 18.

The sessions are numbered from the bottom of the rack to the top.

6. If 18 sessions are not shown:
 - a. Select **IQ Adaptors** under the Ports heading.
 - b. Click the **Port** table heading to sort the sessions by port number.
 - c. Note any missing sessions, so you can fix them later.
 - d. Return to the Target Devices screen.

Connecting the KVM to the Management Network

1. Select **Local** from User Accounts.
2. Click **Admin** under Users.
3. Set the password for the Admin account to `welcome1`, then click **Save**. Do not modify any other parameters.
4. From Appliance Settings, choose **Network**, then **IPv4**. The Network Information screen appears.
5. Enter values for Address, Subnet, and Gateway.
6. From Appliance Settings, choose **DNS** to display the DNS Information screen.
7. Enter the IP addresses of the DNS servers, then click **Save**.
8. Under Appliance Settings, choose **General**.
9. Connect the KVM LAN1 Ethernet port to the management network.
10. To verify the port has been configured correctly, ensure the MAC address on the Network Information screen matches the label next to the LAN1/LAN2 ports at the rear of the KVM switch.
11. Select **Overview** from Appliance.
12. Enter a name for the KVM switch, then click **Save**.
13. To reboot the KVM switch, click **Reboot** under Overview and **Yes** to confirm.

Checking the KVM Firmware Version

1. Examine the firmware version of the switch by selecting **Versions** from Appliance Settings. There are two version numbers, Application and Boot:

```
Required version is:  
Application 1.10.2.17762  
Boot 1.9.16473
```

The recommended firmware version is 1.10.2 or later. To upgrade the firmware, continue with this procedure. Otherwise, you are done.

2. Download the firmware from this website to a USB flash drive:

<http://www.avocent.com/Pages/GenericTwoColumn.aspx?id=12541>

3. Plug the flash drive into the KVM USB port and open a browser session.
4. Log in to the KVM as Admin with password welcome1.
5. Select **Overview** from Appliance.
6. Select **Upgrade Firmware** from the Tools list.
7. Select the connection method, such as **http** or **ftp**.
8. Enter the file name of the downloaded firmware.
9. Click **Upgrade**.

The upgrade process takes 5-10 minutes, including an automatic reboot.

10. Confirm the firmware version.

Configuring the KVM Switch to Access the Servers

To configure the KVM switch to access the servers:

1. Select **Target Devices** from Unit View.
2. Power on the server. The power button is on the front panel.
3. Click the system name in the Name column.
4. Click **Overview** and overwrite the name with the Oracle standard naming format of customer prefix, node type, and number. For example, bda1node03 identifies the third server from the bottom of the bda1 rack.
5. Press **Save**.
6. Repeat steps 2 through 5 for each server in the rack. Each server boots through BIOS, and boots the operating system with the default factory IP configuration.

Accessing a Server Using the KVM Switch

To access a server using the KVM switch:

1. Select **Target Devices** from the Unit View.
2. Click the system name in the Name column.
3. Click the KVM session.

Configuring the Cisco Ethernet Switch

The Cisco Catalyst 4948 Ethernet switch supplied with Oracle Big Data Appliance is minimally configured during installation. These procedures configure the Cisco Ethernet switch into one large Virtual LAN. To create multiple VLANs, see [Chapter 11, "Creating Virtual LANs."](#)

This configuration disables IP routing and sets the following:

- Host name
- IP address
- Subnet mask
- Default gateway
- Domain name

- Name server
- NTP server
- Time
- Time zone

Prerequisites

- Do not connect the Cisco Ethernet switch until the network administrator has verified the running configuration and made any necessary changes.
- Do not connect the Cisco Ethernet switch to the customer network until the IP addresses on all components have been configured in Oracle Big Data Appliance. This sequence prevents any duplicate IP address conflicts, which are possible due to the default addresses set in the components when shipped.
- Configure the Cisco Ethernet switch with the network administrator.

See Also:

- [Chapter 3, "Network Requirements"](#)
- *Oracle Big Data Appliance Site Checklists*

Configuring the Cisco Ethernet Switch

1. Connect a serial cable from the Cisco switch console to a laptop or similar device. An RJ45 to DB9 serial cable is included with the Cisco documentation package.
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 enable mode using the following command. The default password is welcome1.

```
Switch> enable
Password:
```

4. Configure the network for a single VLAN. The following is an example of the configuration:

```
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# *Jan 23 15:54:00.506: %SYS-5-CONFIG_I:Configured from console by
console
Switch# write memory
Building configuration...
Compressed configuration from 2474 bytes to 1066 bytes [OK ]
```

5. If the network does not require IP routing on the switch, then disable the default IP routing setting and configure the default gateway. This method is preferred. Consult the network administrator if in doubt.

```
Switch# configure terminal
```



```

Enter configuration commands,one per line. End with CNTL/Z.
Switch(config)# no ip routing
Switch(config)# ip default-gateway 10.17.7.1
Switch(config)# end
*Jan 23 15:54:00.506: %SYS-5-CONFIG_I:Configured from console by console
Switch# write memory
Building configuration...
CCompressed configuration from 3600 bytes to 1305 bytes[OK]]

```

6. If the network requires IP routing on the switch, then keep the default IP routing setting and configure the default gateway as follows:

```

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
*Jan 23 15:55:02.506: %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 using the standard Oracle Big Data Appliance naming convention of *rack_namesw-ip*. This example uses the name *bdalsw-ip*.

```

Switch# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
Switch(config)# hostname bdalsw-ip
bdalsw-ip(config)# end
*Jan 23 15:57:50.886: %SYS-5-CONFIG_I: Configured from console by console
bdalsw-ip# write memory
Building configuration...
Compressed configuration from 3604 bytes to 1308 bytes[OK]
bdalsw-ip#
The system host name appears in the prompt.

```

Setting Up Telnet Access

1. Set the password for telnet access if necessary; it should be set out of the box:

```

bdalsw-ip # configure terminal
Enter configuration commands,one per line. End with CNTL/Z.
bdalsw-ip(config)# enable password welcomel
bdalsw-ip(config)# enable secret welcomel
The enable secret you have chosen is the same as your enable password.
This is not recommended. Re-enter the enable secret.
bdalsw-ip(config)# end
bdalsw-ip# write memory
*Jan 23 15:57:50.886: %SYS-5-CONFIG_I:Configured from console by console
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes [OK ]

```

2. Set up telnet access. In this example, the first login output shows the password is not set and telnet access is disabled. If the login command returns nothing, then the password is set and telnet access is available.

```

Switch# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
bdalsw-ip(config)# line vty 0 15
bdalsw-ip(config-line)# login
%Login disabled on line 1,until 'password'is set
%Login disabled on line 2,until 'password'is set
%Login disabled on line 3,until 'password'is set

```

```

...
bdal-sw-ip(config-line)# password welcome1
bdal-sw-ip(config-line)# login
bdal-sw-ip(config-line)# end
bdal-sw-ip# write memory
*Jan 23 15:58:53.630: %SYS-5-CONFIG_I: Configured from console by consolemory
Building configuration...
Compressed configuration from 3604 bytes to 1308 bytes[OK]

```

3. To disable telnet access and prevent remote access, follow this example:

```

Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdal-sw-ip(config)# line vty 0 15
bdal-sw-ip(config-line)# no password
bdal-sw-ip(config-line)# login
%Login disabled on line 1, until 'password' is set
%Login disabled on line 2, until 'password' is set
%Login disabled on line 3, until 'password' is set
...
bdal-sw-ip(config-line)# end
bdal-sw-ip# write memory
*Jan 23 15:58:53.630: %SYS-5-CONFIG_I: Configured from console by consolemory
Building configuration...
Compressed configuration from 3604 bytes to 1308 bytes[OK]

```

Identifying the DNS Servers

- Configure up to three Domain Name System (DNS) servers, replacing the values shown here with valid ones for the site:

```

bdal-sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdal-sw-ip(config)# ip domain-name us.example.com
bdal-sw-ip(config)# ip name-server 10.7.7.3
bdal-sw-ip(config)# ip name-server 172.28.5.5
bdal-sw-ip(config)# ip name-server 10.8.160.1
bdal-sw-ip(config)# end
*Jan 23 16:01:35.010: %SYS-5-CONFIG_I: Configured from console by console
bdal-sw-ip# write memory
Building configuration...
Compressed configuration from 3662 bytes to 1348 bytes[OK]

```

Setting the Clock and Time Zone

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

To set the local time and time zone, ordering is important. The following is an example of setting local time to US Eastern time:

```

bdal-sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdal-sw-ip(config)# clock timezone EST -5
bdal-sw-ip(config)# clock summer-time EDT recurring
bdal-sw-ip(config)# end
bdal-sw-ip# clock set 15:00:00 January 23 2012
bdal-sw-ip# write memory
Building configuration...
Compressed configuration from 3778 bytes to 1433 bytes[OK]
bdal-sw-ip# show clock
15:00:18.819 EST Mon Jan 23 2012

```

Following are descriptions of the commands for setting the clock and time zone:

- To use UTC, issue this command:

```
no clock timezone global configuration
```

- To use a time zone:

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

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

- To set summer time hours:

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

Where *zone* is the time zone to be displayed when summer time is in effect, *week* is the week of the month (1 to 5 or last), *day* is the day of the week, *month* is the month, *hh:mm* is the time in 24-hour format, and *offset* is the number of minutes to add during summer time. The default offset is 60 minutes.

- To manually set the clock to any time:

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

Where *hh:mm:ss* is the time in 24-hour format, *day* is the day of the month, *month* is the month, and *year* is the year. The time specified is relative to the configured time zone.

See Also: *Cisco IOS Configuration Fundamentals Command Reference* at

http://www.cisco.com/en/US/docs/ios/12_2/configfun/command/reference/frf012.html

Configuring the NTP Servers

- Configure up to two NTP servers. The following example shows the NTP server synchronized to local time when the Cisco switch is connected to the network and has access to NTP.

```
bdal-sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdal-sw-ip(config)# ntp server 10.196.23.254 prefer
bdal-sw-ip(config)# ntp server 192.168.9.19
bdal-sw-ip(config)# end
Jan 23 20:00:41.235: %SYS-5-CONFIG_I:Configured from console by console
bdal-sw-ip# write memory
Building configuration...
Compressed configuration from 3870 bytes to 1487 bytes [OK ]
bdal-sw-ip# show ntp status
output varies by network
bdal-sw-ip# show clock
15:00:57.919 EST Mon Jan 23 2012
```

Disabling the Spanning Tree

1. Ask the network administrator whether the network requires the spanning tree to be enabled before connecting the Cisco Ethernet switch.

Spanning tree is enabled by default on switch-to-switch uplink port 48. If this is correct, then you can skip this procedure.

2. If the port must be disabled, then issue these commands:

```
bdal-sw-ip#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#no spanning-tree vlan 1
Switch(config)#end
Jan 23 20:01:15.083: %SYS-5-CONFIG_I: Configured from console by console
bdal-sw-ip#write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes[OK]
```

3. To verify disabling of the spanning tree:

```
bdal-sw-ip#show spanning-tree vlan 1
Spanning tree instance(s) for vlan 1 does not exist.
```

Verifying the Configuration

1. Verify the configuration using the following command:

```
bdal-sw-ip# show running-config
```

The following is an example of the output:

```
Building configuration...
Current configuration :2654 bytes
!
version 12.2
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
.
.
.
```

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, issue these commands:

```
bdal-sw-ip# no ip default-gateway 10.7.7.1
bdal-sw-ip# end
bdal-sw-ip# write memory
```

2. Save the current configuration by issuing this command:

```
bdal-sw-ip# copy running-config startup-config
```

3. Exit from the session with this command:

```
bdal-sw-ip#exit
bdal-sw-ip con0 is now available
```

4. Disconnect the cable from the Cisco console.
5. To check the configuration, attach a laptop to port 48 and ping the IP address of the internal management network. Do not connect the Cisco Ethernet switch to the management network.

Configuring the InfiniBand Switches

Oracle Big Data Appliance has two Sun Network QDR InfiniBand Gateway leaf switches and one Sun Datacenter InfiniBand Switch 36 spine switch. To configure these switches, follow these procedures for each one:

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

Configuring an InfiniBand Switch

1. Connect to an InfiniBand gateway leaf switch using a serial or an Ethernet connection
 - For an Ethernet connection to 192.168.1.202 or 192.168.1.203, see ["Connecting to ILOM Using the Network"](#) on page 6-3.
 - For a serial connection, see ["Connecting to ILOM Using a Serial Connection"](#) on page 6-4.

2. Log in as `ilom-admin` with password `welcome1`.

The switch has a Linux-like operating system and an Oracle ILOM interface that is used for configuration.

3. Change to the `/SP/network` directory.

```
cd /SP/network
```

4. Issue these commands to configure the switch:

```
set pendingipaddress=ip_address
set pendingipnetmask=ip_netmask
set pendingipgateway=ip_gateway
set pendingipdiscovery=static
set commitpending=true
```

Where *ip_address*, *ip_netmask*, and *ip_gateway* represent the appropriate settings on your network.

5. Issue a `show` command to view the changes. If any values are wrong, reissue the set commands ending with `set commitpending=true`.

```
-> show
```

```
/SP/network
```

```
Targets:
```

```
interconnect
ipv6
test
```

```
Properties:
```

```
commitpending = (Cannot show property)
dhcp_ser_ip = none
ipaddress = 10.135.42.24
ipdiscovery = static
ipgateway = 10.135.40.1
ipnetmask = 255.255.255.0
macaddress = 00:21:28:E7:B3:34
managementport = SYS/SP/NET0
outofbandmacaddress = 00:21:28:E7:B3:33
pendingipaddress = 10.135.42.23
```

```

pendingipdiscovery = static
pendingipgateway = 10.135.42.1
pendingipnetmask = 255.255.248.0
pendingmanagementport = /SYS/SP/NET0
sidebandmacaddress = 00:21:28:E7:B3:35
state = enabled

```

Commands:

```

cd
set
show

```

->

6. Set and verify the switch host name, replacing *hostname* with the valid name of the switch, such as `bdal1sw-ib2`. Do not include the domain name.

```

-> set /SP hostname=hostname
-> show /SP hostname

```

/SP

Properties:

```
hostname = bdal1sw-ib2
```

7. Set the DNS server name and the domain name:

```

-> set /SP/clients/dns auto_dns=enabled
-> set /SP/clients/dns nameserver=ip_address
-> set /SP/clients/dns searchpath=domain_name

```

Where *ip_address* is one to three comma-separated IP addresses of the name servers in the preferred search order, and *domain_name* is the full DNS domain name, such as `us.example.com`.

8. Verify the settings:

```

-> show /SP/clients/dns
/SP/clients/dns
Targets:

Properties:
auto_dns = enabled
nameserver = 10.196.23.245, 172.32.202.15
retries = 1
searchpath = us.example.com
timeout = 5

```

Commands:

```

cd
set
show

```

Setting the Time Zone on an InfiniBand Switch

1. Check the current time setting:

```
-> show /SP/clock
```

If the setting is not accurate, continue with these steps.

2. Set the time zone, replacing *zone_identifier* with the time zone on the Configuration Template, such as `America/New_York`:

```
-> set /SP/clock timezone=zone_identifier
```

3. Check the current time setting:

```
-> show /SP/clock
```

If the setting is not accurate, continue with these steps.

4. Set the SP clock manually, replacing *MMDDHHmmCCyy* with the month, day, hour, minute, century, and year.

```
-> set datetime=MMddHHmmCCyy
```

5. Check the current time setting:

```
-> show /SP/clock
```

6. Configure the Network Time Protocol (NTP), replacing *ip_address* with the server address. Server 1 is the primary NTP server and 2 is the secondary server.

```
-> set /SP/clients/ntp/server/1 address=ip_address
```

```
-> set /SP/clients/ntp/server/2 address=ip_address
```

7. Enable the NTP servers:

```
-> set /SP/clock usentpserver=enabled
```

Note: Properly synchronized clocks are required for the Mammoth Utility software installation to succeed. If NTP is not used on the network, then configure the first server in Oracle Big Data Appliance Rack as an NTP server.

8. Verify the settings:

```
-> show /SP/clients/ntp/server/1
```

```
-> show /SP/clients/ntp/server/2
```

```
-> show /SP/clock
```

Checking the Health of an InfiniBand Switch

1. Open the Fabric Management shell:

```
-> show /SYS/Fabric_Mgmt
```

The prompt changes from -> to FabMan@hostname->

2. Check the firmware version, which should be 2.0.5-2 or later. See My Oracle Support Note 1434477.1 for the current version.

```
FabMan@bdalsw-02->version
SUN DCS gw version: 2.0.5-2
Build time: Nov 29 2011 16:05:05
FPGA version: 0x34
SP board info:
Manufacturing Date: 2011.05.31
Serial Number: "NCD6Q0126"
Hardware Revision: 0x0006
Firmware Revision: 0x0000
BIOS version: SUN0R100
BIOS date: 06/22/2010
FabMan@bdalsw-02->
```

3. Check the overall health of the switch and correct any issues:

```
FabMan@bdalsw-ib2-> showunhealthy
OK - No unhealthy sensors
```

4. Check the environment. Ensure that all tests return OK and PASSED, and correct any issues before continuing. This example shows a problem with PSU1 caused by a loose power cord. See the line starting with WARNING PSU.

```
FabMan@bdalsw-ib2-> env_test
Environment test started:
Starting Environment Daemon test:
Environment daemon running
Environment Daemon test returned OK
Starting Voltage test
Voltage ECB OK
Measured 3.3V Main = 3.25
Measured 3.3V Standby = 3.37 V
Measured 12V = 11.97 V
Measured 5V = 4.99 V
Measured VBAT = 3.09 V
Measured 1.0V = 1.01 V
Measured I4 1.2V = 1.22 V
Measured 2.5V = 2.52 V
Measured V1P2 DIG = 1.19 V
Measured V1P2 ANG = 1.18 V
Measured 1.2V BridgeX = 1.22 V
Measured 1.8V = 1.78 V
Measured 1.2V Standby = 1.20 V
Voltage test returned OK
Starting PSU test:
PSU 0 present OK
WARNING PSU 1 present AC Loss
PSU test returned 1 faults
Starting Temperature test:
Back temperature 30
Front temperature 29
SP temperature 36
Switch temperature 52,
.
.
.
```

5. Verify a priority setting of **5** for the InfiniBand Gateway leaf switches or **8** for the InfiniBand Switch 36 spine switch:

```
FabMan@bdalsw-ib2-> setsmpriority list
Current SM settings:
smpriority 5
controlled_handover TRUE
subnet_prefix 0xfe80000000000000
```

If `smpriority` is correct, then you can skip the next step.

6. To correct the priority setting:

- a. Stop the InfiniBand Subnet Manager:

```
FabMan@bdalsw-ib2-> disableism
```

- b. Set the priority to **5** for the InfiniBand Gateway leaf switches or **8** for the InfiniBand Switch 36 spine switch. This example is for a leaf switch:


```
FabMan@bda1sw-ib2-> setpriority 5
```

- c. Restart the InfiniBand Subnet Manager:

```
FabMan@bda1sw-ib2-> enablesm
```

7. If you are connecting this Oracle Big Data Appliance rack to an Oracle Exadata Database Machine or Oracle Exalogic Machine rack:
 - a. Verify that the Exadata InfiniBand switches and the Exalogic spine switch are running firmware version 1.3.3_2 or later.
 - b. On systems running earlier firmware versions, disable the subnet manager. Log in to the switch as `root` and run the `disablesm` command as described previously.
8. Exit the Fabric Manager shell:


```
FabMan@bda1sw-ib2-> exit
->
```
9. Exit the ILOM shell:


```
-> exit
```
10. Log in to the switch as `root` and restart it to ensure all changes take effect:


```
reboot
```
11. Repeat these steps for the other InfiniBand switches.

Configuring the Power Distribution Units

The power distribution units (PDUs) are configured with a static IP address to connect to the network for monitoring. Ensure you have the following before connecting the PDU to the network:

- Static IP address
- Subnet mask
- Default gateway

Connecting the PDU to the Network

1. Use a web browser to access the PDU metering unit by entering the factory default IP address for the unit. The address of PDU A is 192.168.1.210, and the address of PDU B is 192.168.1.211.
The Current Measurement page opens.
2. Click **Net Configuration** in the upper left of the page.
3. Log in as the `admin` user on the PDU metering unit. The default password is `admin`. Change this password after configuring the network.
4. Confirm the DHCP Enabled option is not selected.
5. Enter the following network settings for the PDU metering unit:
 - IP address
 - Subnet mask address
 - Default gateway

6. Click **Submit** to set the network settings and reset the PDU metering unit.

Verifying the PDU Firmware Version

1. Select **Module Info**. If the output displays a firmware version of 1.04 or later, then you are done. Otherwise, continue this procedure to update the firmware version.
2. Download the latest firmware version from My Oracle Support:
 - a. Log in at <http://support.oracle.com>.
 - b. Select the **Patches & Updates** tab.
 - c. For Patch Search, click **Product or Family (Advanced)**.
 - d. For Product, select Sun Rack II PDU.
 - e. For Release, select Sun Rack II PDU 1.0.4.
 - f. Click **Search** to see the Patch Search Results page.
 - g. Click the patch name, such as **12871297**.
 - h. Download the file.
3. Unzip the file on your local system.
4. Return to the PDU metering unit Network Configuration page
5. Scroll down to Firmware Update.
6. Click **Browse**, select the `MKAPP_V1.04.DL` file, and click **Submit**.
7. Click **Browse**, select the `HTML_V1.04.DL` file, and click **Submit**.
8. Click **Module Info** to verify the version number.
9. Click **Net Configuration**, then **Logout**.

Configuring the Threshold Settings for the Power Distribution Units

The PDU current can be monitored directly. Configure the threshold settings to monitor the PDUs. The configurable threshold values for each metering unit module and phase are Info low, Pre Warning, and Alarm.

See Also: *Sun Rack II Power Distribution Units User's Guide* for information about configuring and monitoring PDUs at

<http://docs.oracle.com/cd/E19844-01/index.html>

Table 8–1 lists the threshold values for Oracle Big Data Appliance Rack using a single-phase, low-voltage PDU.

Table 8–1 Threshold Values for Single-Phase, Low-Voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	18	23
A	Module 1, phase 2	0	22	24
A	Module 1, phase 3	0	18	23
B	Module 1, phase 1	0	18	23
B	Module 1, phase 2	0	22	24
B	Module 1, phase 3	0	18	23

Table 8–2 lists the threshold values for Oracle Big Data Appliance Rack using a three-phase, low-voltage PDU.

Table 8–2 Threshold Values for Three-Phase, Low-Voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	32	40
A and B	Module 1, phase 2	0	34	43
A and B	Module 1, phase 3	0	33	42

Table 8–3 lists the threshold values for Oracle Big Data Appliance Rack using a single-phase, high-voltage PDU.

Table 8–3 Threshold Values for Single-Phase, High-Voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A	Module 1, phase 1	0	16	20
A	Module 1, phase 2	0	20	21
A	Module 1, phase 3	0	16	20
B	Module 1, phase 1	0	16	20
B	Module 1, phase 2	0	20	21
B	Module 1, phase 3	0	16	20

Table 8–4 lists the threshold values for Oracle Big Data Appliance Rack using a three-phase, high-voltage PDU.

Table 8–4 Threshold Values for Three-Phase, High-Voltage PDU

PDU	Module/Phase	Info Low Threshold	Pre Warning Threshold	Alarm Threshold
A and B	Module 1, phase 1	0	18	21
A and B	Module 1, phase 2	0	18	21
A and B	Module 1, phase 3	0	17	21

Checking Oracle Big Data Appliance Servers

The following procedure describes how to check the servers.

Note:

- When you use the KVM switch and press the ESC key, the BIOS may receive two ESC characters and prompt to exit. Select **CANCEL**.
 - If you must connect to the ILOM serial management port, then the baud rate setting on Sun Fire Servers changes from the default ILOM setting 9600 to 115200 baud, 8 bits, no parity, and 1 stop bit.
-
-

1. Power on all servers by pressing the power button on the front panel of each server. The servers take 5-10 minutes to start through the normal tests.
2. To configure the KVM with the server names:
 - a. Under Unit View, select Target Devices.
 - b. Click the system name in the Name column.
 - c. Click Overview and overwrite the name with the appropriate name from the Installation Template. For example, `bda1node01` identifies the first server (node01) in a rack named bda1. The servers are numbered from bottom to top, as shown in "[Oracle Big Data Appliance Rack Layout](#)" on page B-1.
 - d. Repeat these steps for all 18 servers. Each server restarts with the factory default IP configuration.
3. Connect to a server using either the KVM or a laptop:
 - KVM: Under Unit View, select Target Devices. Left-click the system name, then click the KVM session.
 - Laptop: Open an SSH connection using PuTTY or a similar utility. Enter the default IP address of the server.
4. Log in as the `root` user to the first server. The password is `welcome1`.
5. Verify that the `/opt/oracle/bda/rack-hosts-infiniband` file exists. If not, create the file with the default IP addresses listed one per line. All `dcli` commands are sent by default to the servers listed in this file. See "[Factory Network Settings](#)" on page 3-5.
6. Set up passwordless SSH for `root` by issuing the `setup-root-ssh` command, as described in "[Setting Up Passwordless SSH](#)" on page 7-1.
7. Verify that SSH keys are distributed across the rack:

```
dcli "hostname ; date"
```

8. If prompted for a password, enter `Ctrl+C` several times. This prompt confirms distribution of the keys, so you can continue to the next step. Otherwise, generate the root SSH keys across the rack, replacing `password` with a valid password:

```
setup-root-ssh -p password
```

Issue the `dcli` command in step 7 again to verify the keys.

9. Verify that the InfiniBand ports are up, two on each server (36 total).

```
# dcli ibstatus | grep phys
192.168.10.1: phys state: 5: LinkUp
192.168.10.1: phys state: 5: LinkUp
..
192.168.10.18: phys state: 5: LinkUp
192.168.10.18: phys state: 5: LinkUp
```

10. Verify that the InfiniBand ports are running at 40 Gb/second (4X QDR):

```
# dcli ibstatus | grep rate
192.168.10.1: rate: 40 Gb/sec (4X QDR)
192.168.10.1: rate: 40 Gb/sec (4X QDR)
..
192.168.10.18: rate: 40 Gb/sec (4X QDR)
192.168.10.18: rate: 40 Gb/sec (4X QDR)
```

11. Verify that Oracle ILOM does not detect any faults:

```
# dcli 'ipmitool sunoem cli "show faulty"'
```

The output should appear as follows for each server:

```
bdalnode02-adm.example.com: Connected. Use ^D to exit.
bdalnode02-adm.example.com: -> show faulty
bdalnode02-adm.example.com: Target      | Property      | Value
bdalnode02-adm.example.com:-----+-----+-----
bdalnode02-adm.example.com:
bdalnode02-adm.example.com: -> Session closed
bdalnode02-adm.example.com: Disconnected
```

12. Save the hardware profile output from each system in a file for review, replacing *filename* with a file name of your choice:

```
# dcli bdacheckhw > filename
```

13. Check the hardware profile output file using commands like the following. In these example, the file name is `all-bdahwcheck.out`.

- To verify there are no failures in the hardware profile:

```
grep -v SUCCESS ~/all-bdahwcheck.out
```

- To verify 24 cores:

```
grep cores ~/all-bdahwcheck.out
```

- To verify 48 GB of memory:

```
grep memory ~/all-bdahwcheck.out
```

- To verify six fans:

```
grep fans ~/all-bdahwcheck.out
```

- To verify the status is OK for both power supplies:

```
grep supply ~/all-bdahwcheck.out
```

- To verify disks 0 to 11 are all the same model, online, spun up, and no alert:

```
grep disk ~/all-bdahwcheck.out | grep "model\|status" | more
```

- To verify the host channel adapter model is Mellanox Technologies MT26428 ConnectX VPI PCIe 2.0:

```
grep Host ~/all-bdahwcheck.out | grep model
```

14. Save the RAID configuration in a file, replacing *filename* with a file name of your choice:

```
dcli MegaCli64 -ldinfo -lall -a0 | grep "Virtual Drive\|State" > filename
```

15. Verify that Optimal 12 virtual drives 0 to 11 are listed for each server. In this example, the RAID configuration is stored in a file named `all-ldstate.out`.

```
less ~/all-ldstate.out
```

16. Save the software profile output from each system into a file for review, replacing *filename* with a file name of your choice:

```
dcli bdacheksw > filename
```

17. Verify that the partition setup and software versions are correct. In this example, the software profile is stored in a file named `all-bdaswcheck.out`.

```
less ~/all-bdaswcheck.out
```

18. Verify the system boots in this order: USB, RAID Slot 0, PXE:

```
dcli 'biosconfig -get_boot_order' | grep DEV | more
```

```
<BOOT_DEVICE_PRIORITY>
  <DEVICE_NAME>USB:02.82;01 Unigen PSA4000</DEVICE_NAME>
  <DEVICE_NAME>RAID:Slot0.F0:(Bus 13 Dev 00)PCI RAID Adapter</DEVICE_NAME>
  <DEVICE_NAME>PXE:IBA GE Slot 0100 v1331</DEVICE_NAME>
  <DEVICE_NAME>PXE:IBA GE Slot 0101 v1331</DEVICE_NAME>
  <DEVICE_NAME>PXE:IBA GE Slot 0700 v1331</DEVICE_NAME>
  <DEVICE_NAME>PXE:IBA GE Slot 0701 v1331</DEVICE_NAME>
</BOOT_DEVICE_PRIORITY>
```

Configuring the Network

The Oracle Big Data Appliance Configuration Utility generates the `BdaDeploy.json` file, which is used to configure the administrative network and the private InfiniBand network. See [Chapter 4, "Using the Oracle Big Data Appliance Configuration Utility,"](#) for information about generating this file.

The network configuration consists of these procedures:

1. [Verifying the Factory Software Image](#)
2. [Copying the Configuration Files to Oracle Big Data Appliance](#)
3. [Starting the Network Configuration](#)
4. [Connecting to the Network](#)
5. [Completing the Network Configuration](#)

Verifying the Factory Software Image

To verify that the factory software image is installed correctly and the servers are operating correctly, check that the `BDA_IMAGING_SUCCEEDED` and `BDA_REBOOT_SUCCEEDED` are in the `/root` directory of each server. If you see `BDA_IMAGING_FAILED` or `BDA_REBOOT_FAILED` in the output, then check the `/root/bda_imaging_status` file for that server for more information. Do not proceed with network configuration until all problems are resolved.

The `dcli` utility requires passwordless SSH for root, as described in ["Setting Up Passwordless SSH"](#) on page 7-1.

```
# dcli ls -l /root | grep BDA
IP address BDA_IMAGING_SUCCEEDED
IP address BDA_REBOOT_SUCCEEDED
.
.
.
```

You can also confirm the image version:

```
# dcli imageinfo
Big Data Appliance Image Info
```

```

IMAGE_VERSION           : 1.0.2
IMAGE_CREATION_DATE     : Sun Mar 4 11:39:36 PST 2012
IMAGE_LABEL             : BDA_MAIN_LINUX.X64_120303
KERNEL_VERSION         : 2.6.32-200.21.1.el5uek
BDA_RPM_VERSION        : bda-1.0.2-1
OFA_RPM_VERSION        : ofa-2.6.32-200.21.1.el5uek-1.5.5-4.0.55.4
JDK_VERSION            : jdk-1.6.0_29-fcs
.
.
.

```

Copying the Configuration Files to Oracle Big Data Appliance

To copy the configuration files:

1. Copy the configuration files to a USB flash drive.
2. Use the KVM to open a console session to the first server. The first server is the lowest server in the rack, which is rack position U01. See [Figure B-1, "Rack Layout of Oracle Big Data Appliance Rack"](#)
3. Log in as the `root` user on the first server. The initial password is `welcome1`.
4. Plug the USB drive into the USB port of the first server. The port is on the right front of the server. Information like the following is displayed on the console:

```

# scsi 0:0:0:0: Direct-Access    CBM      USB 2.0
Q: 0 ANSI:2
sd 0:0:0:0: Attached scsi generic sg14 type 0
sd 0:0:0:0: [sdn] 7954432 512-byte logical blocks: (4.07 GB/3.79 GiB)
sd 0:0:0:0: [sdn] Write Protect is off
sd 0:0:0:0: [sdn] Assuming drive cache: write through
sd 0:0:0:0: [sdn] Assuming drive cache: write through
sd 0:0:0:0: [sdn] Assuming drive cache: write through
sd 0:0:0:0: [sdn] Attached SCSI removable disk

```

5. Issue the `showusb` command to locate the USB drive. The command returns with the mapped device or, if no USB drive is connected, with no output.

```

# showusb
/dev/sdn1

```

6. Create a directory on the server:

```

# mkdir /mnt/usb

```

7. Mount the device using the device name given in step 5. The following is an example of the command.

```

# mount -t vfat /dev/sdn1 /mnt/usb

```

8. Verify the location of the file on the USB flash drive:

```

# ls /mnt/usb
BdaDeploy.json
bin
boot
.
.
.

```

9. Copy `BdaDeploy.json` from the USB flash drive to the `/opt/oracle/bda` directory on the server:

```
# cd /mnt/usb
# cp BdaDeploy.json /opt/oracle/bda
```

Tip: If `mammoth-rack_name.params` is also on the drive, you can copy it to `/opt/oracle/BDAMammoth` for use in [Chapter 13, "Installing the Oracle Big Data Appliance Software."](#)

10. Unmount the USB flash drive and remove the device:

```
# umount /mnt/usb
# rmdir /mnt/usb
```

11. Remove the USB flash drive from the server.

Starting the Network Configuration

The `networksetup-one` script sets up the host names and ILOMs for all servers and configures the administrative network and the private InfiniBand network.

Configuring the Network

1. Log in as the root user on the first server. The initial password is `welcome1`.

```
ssh root@192.168.10.1
```

2. Begin the network configuration:

```
cd /opt/oracle/bda/network
./networksetup-one
```

[Example 8-1](#) shows sample output from the script.

Example 8-1 Sample Output from `networksetup-one`

```
# ./networksetup-one
networksetup-one: check syntax and static semantics of
/opt/oracle/bda/BdaDeploy.json
networksetup-one: passed
networksetup-one: ping servers on ship admin network
networksetup-one: passed
networksetup-one: test ssh to servers on ship admin network
hello from node02
hello from node03
.
.
.
networksetup-one: passed
networksetup-one: copy /opt/oracle/bda/BdaDeploy.json to servers
BdaDeploy.json 0% 0 0.0KB/s --:-- ETABdaDeploy.json 100% 4304 4.2KB/s 00:00
BdaDeploy.json 0% 0 0.0KB/s --:-- ETABdaDeploy.json 100% 4304 4.2KB/s 00:00
.
.
.
networksetup-one: passed
networksetup-one: executing network settings on all servers
networksetup-one: wait a few seconds for the network to restart on 192.168.1.2
.
.
.
bda1node02.us.oracle.com BdaUserConfigNetwork: reset network
bda1node03.us.oracle.com BdaUserConfigNetwork: reset network
```



```

bda1node04.us.oracle.com BdaUserConfigNetwork: reset network
.
.
networksetup-one: deploying this server
networksetup-one: network will restart momentarily, pardon our dust
bda1node01.us.oracle.com BdaUserConfigNetwork: reset network
networksetup-one: generate dcli bda host file lists
networksetup-one: ping server ips on admin network
networksetup-one: passed
networksetup-one: passed
networksetup-one: test ssh server ips on admin network
hello from bda1node02.us.oracle.com
hello from bda1node03.us.oracle.com
hello from bda1node04.us.oracle.com
.
.
networksetup-one: passed

```

Connecting to the Network

Before completing the network configuration, you must connect the administrative and client networks to the data center.

To connect Oracle Big Data Appliance to the network:

- Connect the 1GbE administrative network by connecting the Cisco Ethernet switch to the data center.
- Connect the 10GbE client network by connecting the two Sun Network QDR InfiniBand Gateway Switch leaf switches to the data center.

Completing the Network Configuration

The `networksetup-two` script completes some steps started by `networksetup-one` that require a network connection. It also configures the default VLAN and all required VNICs for the 10GbE client network. It then verifies all network connections and displays a message if it discovers any unexpected ones, including those caused by cabling mistakes.

The 10GbE ports of the Sun Network QDR InfiniBand Gateway Switches must be connected to the data center.

To complete the network configuration:

1. Ensure that both the administrative network and the client network are connected to Oracle Big Data Appliance.

Note: This procedure fails if the networks are not connected. See ["Connecting to the Network"](#) on page 8-21.

2. Run the next script to complete the network setup:

```
./networksetup-two
```

[Example 8-2](#) shows sample output from the script.

Example 8–2 Sample Output from networksetup-two

```
# ./networksetup-two
networksetup-two: check syntax and static semantics of
/opt/oracle/bda/BdaDeploy.json
networksetup-two: passed
networksetup-two: ping server ips on admin network
networksetup-two: passed
networksetup-two: test ssh server ips on admin network
hello from bda1node02.us.oracle.com
hello from bda1node03.us.oracle.com
hello from bda1node04.us.oracle.com
.
.
.
networksetup-two: passed
networksetup-two: run connected network post script on each server
networksetup-two: post network setup for 10.133.42.253
networksetup-two: post network setup for 10.133.42.254
networksetup-two: post network setup for 10.133.43.1
.
.
.
networksetup-two: post network setup for this node
networksetup-two: ping admin servers by name on admin network
networksetup-two: passed
networksetup-two: verify infiniband topology
networksetup-two: passed
networksetup-two: start setup client network (10gigE over Infiniband)
networksetup-two: ping both gtw leaf switches
networksetup-two: passed
networksetup-two: verify existence of gateway ports
networksetup-two: passed
networksetup-two: ping server ips on admin network
networksetup-two: passed
networksetup-two: ping servers by name on admin network
networksetup-two: passed
networksetup-two: test ssh server ips on admin network
hello from bda1node02.us.oracle.com
hello from bda1node03.us.oracle.com
.
.
.
networksetup-two: passed
networksetup-two: check existence of default vlan for port 0A-ETH-1 on bda1sw-ib2
networksetup-two: no default vlan for port, create it
spawn ssh root@10.133.43.36 createvlan 0A-ETH-1 -vlan -1 -pkey default
networksetup-two: verify default vlan for port 0A-ETH-1 for bda1sw-ib2
.
.
.
networksetup-two: passed
networksetup-two: apply eoib on each server
networksetup-two: wait a few seconds for the network to restart on 10.133.42.253
networksetup-two: wait a few seconds for the network to restart on 10.133.42.254
.
.
.
check and delete vNIC for bda1node02 eth9 on switch bda1sw-ib2
check and delete vNIC for bda1node02 eth9 on switch bda1sw-ib3
create vNIC eth9 bda1node02 using switch bda1sw-ib3
```

```

vNIC created
check and delete vNIC for bda1node02 eth8 on switch bda1sw-ib2
.
.
.
networksetup-two: ping server ips on client network
networksetup-two: passed
networksetup-two: test ssh server ips on client network
hello from bda1node02.us.oracle.com
hello from bda1node03.us.oracle.com
.
.
.
networksetup-two: passed
networksetup-two: end setup client network

```

Reinstalling the Base Image

Oracle Big Data Appliance is factory-installed with the operating system and various utilities, as described in ["Oracle Big Data Appliance Management Software"](#) on page 1-3. You may need to reinstall this base image if, for example, you want to return Oracle Big Data Appliance to its original state, or you want to upgrade the base image to a more recent version before using the Mammoth Utility to install the Oracle Big Data Appliance software.

Following is the procedure for reimaging an entire rack.

Caution: If you reinstall the base image, then all files on that server are erased.

To reinstall the base image on all servers in a rack:

1. If the Oracle Big Data Appliance software was installed previously on the rack, then save the `/opt/oracle/BDAMammoth/mammoth-rack_name.params` file to a safe place outside Oracle Big Data Appliance.
2. Download the tarball with the correct version of the base image and copy it to node01 (bottom server). See My Oracle Support Master Note 1434477.1 for the download location. You can copy the tarball to any directory, such as `/tmp`.

The name of the tarball is in this format:

```
BDABaseImage-version.tar
```

For example, `BDABaseImage-1.0.3.tar`.

Tip: You can also download `BDAMammoth-version.run` to a safe location. It contains the version of the Mammoth Utility that you must run to install the end-user software after reimaging.

3. Establish an SSH connection to node01 and log in as `root`.
4. To reimage to the custom network settings, locate the `BdaDeploy.json` configuration file and verify that it reflects the intended network configuration. Edit the file and copy it to `/opt/oracle/bda/` as needed.

Or, to reimage to the factory defaults, locate `BdaShip.json`.

5. Ensure that passwordless SSH is set up:

```
dcli hostname
```

This command should run without errors and return the host names of all 18 Oracle Big Data Appliance servers. If not, then follow the steps in ["Setting Up Passwordless SSH"](#) on page 7-1. Do not continue until the `dcli hostname` command runs successfully on all servers.

6. Verify that at least 4GB are available in the root (/) partition of node01:

```
[# df -h /
Filesystem      Size  Used Avail Use% Mounted on
/dev/md2        161G  8.2G  145G   6% /
```

7. Extract all files from the base image tarball, for example:

```
# tar -xvf BDABaseImage-1.0.3.tar
BDABaseImage-1.0.3/
BDABaseImage-1.0.3/README.baseimage
BDABaseImage-1.0.3/BDABaseImage-1.0.3.iso
BDABaseImage-1.0.3/BDABaseImage-1.0.3.md5sum
BDABaseImage-1.0.3/reimagerack
BDABaseImage-1.0.3/makebdimage
```

8. Change to the `BDABaseImage-version` directory created in the previous step, for example:

```
cd BDABaseImage-1.0.3
```

9. To reimage Oracle Big Data Appliance to the customer network settings specified in `/opt/oracle/bda/BdaDeploy.json`:

```
./reimagerack
```

Or

To reimage an appliance that still has the factory settings:

- a. Ensure that `/opt/oracle/bda/BdaDeploy.json` does not exist.
- b. Issue the `./reimagerack` command.

Or

To restore the factory network settings on a rack configured with custom network settings:

- a. Copy `/opt/oracle/bda/BdaDeploy.json` to a safe location outside Oracle Big Data Appliance.
- b. Disconnect the rack from the network.
- c. Reimage the rack:

```
./reimagerack deploy ship
```

The `reimagerack` utility creates an ISO, copies it to the internal USB drive of each server in the rack, reboots each server, and initializes the installation.

10. Run the Mammoth Utility, as described in [Chapter 13, "Installing the Oracle Big Data Appliance Software."](#)

Checking the Health of the Network

Following are commands that you can run at any time to check the health of the Oracle Big Data Appliance network. This section also contains commands that you may need if the health checks fail.

- `bdcheckib`
- `bdcheckhw`
- `bdchecknet`
- `bdchecksw`
- `bdadiag`
- `bdaid`
- `bdimagevalidate`
- `bdaredoclientnet`
- `bdaserials`
- `iblinkinfo`
- `imageinfo`
- `imagehistory`
- `listlinkup`
- `showvlan`
- `showvnic`

bdcheckib

Checks the InfiniBand cabling between the servers and switches of a single rack, when issued with no options. The network must be configured with custom settings as described by `/opt/oracle/bda/BdaDeploy.json`.

Run this command after connecting as `root` to any server.

The `bdcheckib` command has these options:

-s

Same as running without options except that the network must still be configured with the factory default settings. You can use this option as soon as Oracle Big Data Appliance arrives at the site, even before the switches are configured.

-m *json_file*

Verifies that the InfiniBand switch-to-switch cabling among multiple ranks is correct. To create *json_file*, see the `-g` option.

-g

Generates a sample JSON file named `sample-multi-rack.json`. Use this file as an example of the format required by the `-m` option.

This example checks the switch-to-server InfiniBand cables:

```
[root@node01 network]# bdcheckib
LINK bda1sw-ib3.15A ... bda1node02.HCA-1.2 UP
LINK bda1sw-ib3.15B ... bda1node01.HCA-1.2 UP
LINK bda1sw-ib3.14A ... bda1node04.HCA-1.2 UP
LINK bda1sw-ib3.14B ... bda1node03.HCA-1.2 UP
```

.
.
.

The next example generates the JSON file and shows the output.

```
[root@bda1node01 bda]# bdacheckib -g
[root@bda1node01 bda]# cat sample-multi-rack.json
# This json multirack spec is generated. The array elements are sorted
# alphabetically. A proper arranged json spec representing racks from left to right
# can be used as input to bdacheckib (bdacheckib -m multi-rack.json)
# Note commas separating rack elements are optional.
[
{"SPINE_NAME": "dm01sw-ib1", "LEAF1_NAME": "dm01sw-ib2", "LEAF2_NAME": "dm01sw-ib3"}
{"SPINE_NAME": "bda1sw-ib1", "LEAF1_NAME": "bda1sw-ib2", "LEAF2_NAME": "bda1sw-ib3"}
{"SPINE_NAME": "bda2sw-ib1", "LEAF1_NAME": "bda2sw-ib2", "LEAF2_NAME": "bda2sw-ib3"}
]
```

The last example checks all the racks on the InfiniBand network using the edited JSON file created in the previous example:

```
# bdacheckib -m sample-multi-rack.json
```

```
Verifying rack #1
leaf: dm01sw-ib2
  LINK ... to rack2 UP
  LINK ... to rack2 UP
  LINK ... to rack1 UP
  LINK ... to rack2 UP
  LINK ... to rack3 UP
  LINK ... to rack3 UP
  LINK ... to rack1 UP
  LINK ... to rack1 UP
leaf: dm01sw-ib3
  LINK ... to rack2 UP
  LINK ... to rack2 UP
  LINK ... to rack1 UP
  LINK ... to rack2 UP
  LINK ... to rack1 UP
  LINK ... to rack3 UP
  LINK ... to rack3 UP
  LINK ... to rack1 UP
```

```
Verifying rack #2
leaf: bda1sw-ib2
  LINK ... to rack1 UP
  LINK ... to rack1 UP
.
.
.
```

bdacheckhw

Checks the hardware profile of the server. See ["Checking Oracle Big Data Appliance Servers"](#) on page 8-15 for tips about using this utility.

This example shows the output from the utility:

```
# bdacheckhw
SUCCESS: Correct system model : SUN FIRE X4270 M2 SERVER
SUCCESS: Correct processor info : Intel(R) Xeon(R) CPU X5675 @ 3.07GHz
SUCCESS: Correct number of types of CPU : 1
SUCCESS: Correct number of CPU cores : 24
```

```

SUCCESS: Sufficient GB of memory (>=48): 48
SUCCESS: Correct GB of swap space : 24
SUCCESS: Correct BIOS vendor : American Megatrends Inc.
SUCCESS: Sufficient BIOS version (>=08080102): 08100102
SUCCESS: Recent enough BIOS release date (>=05/23/2011) : 10/11/2011
SUCCESS: Correct ILOM version : 3.0.16.10.a r68533
SUCCESS: Correct number of fans : 6
SUCCESS: Correct fan 0 status : ok
SUCCESS: Correct fan 1 status : ok
.
.
.

```

bdachecknet

Checks whether the network configuration is working properly. Run this command after connecting as root to any server.

This example shows the output from the utility:

```

[root@node01 network]# bdachecknet
bdachecknet: check syntax and static semantics of /opt/oracle/bda/BdaDeploy.json
bdachecknet: passed
bdachecknet: ping test private infiniband ips (bondib0 40gbs)
bdachecknet: passed
bdachecknet: ping test admin ips (eth0 1gbs)
bdachecknet: passed
bdachecknet: ping test client access ips (bondeth0 10gbs Eoib)
bdachecknet: passed
bdachecknet: test admin network resolve and reverse resolve
bdachecknet: passed
bdachecknet: test admin name array matches ip array
bdachecknet: passed
bdachecknet: test client network (eoib) resolve and reverse resolve
bdachecknet: passed
bdachecknet: test client name array matches ip array
bdachecknet: passed
bdachecknet: test ntp servers
bdachecknet: passed
bdachecknet: test arp -a
bdachecknet: passed

```

bdachecksw

Checks the software profile of the server. See "[Checking Oracle Big Data Appliance Servers](#)" on page 8-15 for tips about using this utility.

This example shows the output from the utility:

```

# bdachecksw
SUCCESS: Correct OS disk sda partition info : 1 ext3 raid 2 ext3 raid 3 linux-swap
4 ext3 primary
SUCCESS: Correct OS disk sdb partition info : 1 ext3 raid 2 ext3 raid 3 linux-swap
4 ext3 primary
SUCCESS: Correct data disk sdc partition info : 1 ext3 primary
SUCCESS: Correct data disk sdd partition info : 1 ext3 primary
SUCCESS: Correct data disk sde partition info : 1 ext3 primary
SUCCESS: Correct data disk sdf partition info : 1 ext3 primary
SUCCESS: Correct data disk sdg partition info : 1 ext3 primary
SUCCESS: Correct data disk sdh partition info : 1 ext3 primary
SUCCESS: Correct data disk sdi partition info : 1 ext3 primary

```

```

SUCCESS: Correct data disk sdj partition info : 1 ext3 primary
SUCCESS: Correct data disk sdk partition info : 1 ext3 primary
SUCCESS: Correct data disk sdl partition info : 1 ext3 primary
SUCCESS: Correct software RAID info : /dev/md2 level=raid1 num-devices=2 /dev/md0
level=raid1 num-devices=2
SUCCESS: Correct mounted partitions : /dev/md0 /boot ext3 /dev/md2 / ext3
/dev/sda4 /u01 ext4 /dev/sdb4 /u02 ext4 /dev/sdc1 /u03 ext4 /dev/sdd1 /u04 ext4
/dev/sde1 /u05 ext4 /dev/sdf1 /u06 ext4 /dev/sdg1 /u07 ext4 /dev/sdh1 /u08 ext4
/dev/sdi1 /u09 ext4 /dev/sdj1 /u10 ext4 /dev/sdk1 /u11 ext4 /dev/sdl1 /u12 ext4
SUCCESS: Correct swap partitions : /dev/sdb3 partition /dev/sda3 partition
SUCCESS: Correct Linux kernel version : Linux 2.6.32-200.21.1.el5uek
SUCCESS: Correct Java Virtual Machine version : HotSpot(TM) 64-Bit Server 1.6.0_29
SUCCESS: Correct puppet version : 2.6.11
SUCCESS: Correct MySQL version : 5.5.17
SUCCESS: All required programs are accessible in $PATH
SUCCESS: All required RPMs are installed and valid
SUCCESS: Big Data Appliance software validation checks succeeded

```

bdadiag

Collects diagnostic information about an individual server and returns the name of the compressed file in /tmp where it stored the data. You must be connected to the server as root.

Following are the bdadiag options, which instruct bdadiag to collect additional diagnostics. You can enter the options together on the command line to collect the most information.

hadoop

Collects the CDH cluster logs for Hadoop and the Cloudera Manager logs.

hdfs

Collects the output of a complete Hadoop Distributed File System (HDFS) check.

osw

Collects Oracle OS Watcher logs, which include historical operating system performance and monitoring data.

This example shows the output from the utility:

```

# bdadiag

Big Data Appliance Diagnostics Collection Tool v1.0.3

Checking installed rpms

Generating diagnostics tarball and removing temp directory

=====
Done. The report files are bzip2 compressed in /tmp/bda1node09_bdadiag_2012_04_10_
14_08.tar.bz2
=====

```

The logs are organized in subdirectories, including the following:

```

asr
ilom
install
messages
raid
sysconfig

```


bdaid

Returns information about an individual server. If you need to contact Oracle Support about an issue with Cloudera's Distribution including Apache Hadoop, you should run this command first. You must be connected to the server as `root`.

This example shows the output from the utility:

```
# bdaid
Server Hostname      : bda1node09
Rack Serial Number  : AK00023713
Server Serial Number : 1137FMM06Y
Cluster Name        : Cluster 1
Appliance Name      : bda1
```

bdaimagevalidate

Validates the hardware and software by running `bdacheckhw`, then `bdachecksw`.

bdaredoclientnet

Re-creates the Virtual NICs for all servers in the rack and spreads them across the available 10 GbE ports. You must run this utility after changing the number of 10 GbE connections to a Sun Network QDR InfiniBand Gateway Switch. The `bdaredoclientnet` utility performs a subset of the tasks done by `networksetup-two` during the initial configuration of Oracle Big Data Appliance:

- Verifies that the administrative network is working, the InfiniBand cabling is correct, and the InfiniBand switches are available.
- Determines how many 10 GbE connections are available and connects them to the InfiniBand Gateway switches.
- Deletes all VNICs and re-creates them.
- Connects to each server and updates the configuration files.
- Restarts the client network and verifies that it can connect to each server using the newly configured client network.

If you have special requirements, such as assigning specific VLANs to different ports, then you must configure them manually. See [Chapter 10, "Configuring Ethernet Over InfiniBand"](#) and [Chapter 11, "Creating Virtual LANs."](#)

To re-create the VNICs in a rack:

1. Verify that `/opt/oracle/bda/BdaDeploy.json` exists on all servers and correctly describes the custom network settings. This command identifies files that are missing or have different date stamps:

```
dccli ls -l /opt/oracle/bda/BdaDeploy.json
```

2. Connect to node01 (bottom of rack) using either the administrative network or the KVM. This utility shuts down the client network, so you cannot use it in this procedure.
3. Change directories:

```
cd /opt/oracle/bda/network
```

4. Run the utility:

```
bdaredoclientnet
```



```

149 7[ ]==( 4X 10.0 Gbps Active/ LinkUp)==> 141 1[ ] "bdalnode04 BDA 192.168.41.22 HCA-1" ( )
149 8[ ]==( 4X 10.0 Gbps Active/ LinkUp)==> 123 1[ ] "bdalnode03 BDA 192.168.41.21 HCA-1" ( )
149 9[ ]==( 4X 10.0 Gbps Active/ LinkUp)==> 151 1[ ] "bdalnode06 BDA 192.168.41.24 HCA-1" ( )
149 10[ ]==( 4X 10.0 Gbps Active/ LinkUp)==> 112 1[ ] "bdalnode05 BDA 192.168.41.23 HCA-1" ( )
149 11[ ]==( 4X 10.0 Gbps Active/ LinkUp)==> 139 1[ ] "bdalnode07 BDA 192.168.41.25 HCA-1" ( )
149 12[ ]==(          Down/Disabled)==> [ ] "" ( )
149 13[ ]==(          Down/Disabled)==> [ ] "" ( )
149 14[ ]==( 4X 10.0 Gbps Active/ LinkUp)==> 85 9[ ] "SUN DCS 36P QDR dm01sw-ib1 10.133.40.203" ( )
149 15[ ]==(          Down/Disabled)==> [ ] "" ( )
.
.
.

```

imageinfo

Displays information about the Oracle Big Data Appliance operating system image currently running.

This example identifies the 1.0.3 image:

```

# imageinfo
Big Data Appliance Image Info

IMAGE_VERSION           : 1.0.3
IMAGE_CREATION_DATE     : Sun Apr 1 20:00:43 PDT 2012
IMAGE_LABEL             : BDA_1.0.3_LINUX.X64_RELEASE
KERNEL_VERSION          : 2.6.32-200.21.1.e15uek
BDA_RPM_VERSION         : bda-1.0.3-1
OFA_RPM_VERSION         : ofa-2.6.32-200.21.1.e15uek-1.5.5-4.0.55.4
JDK_VERSION             : jdk-1.6.0_29-fcs

```

imagehistory

Displays a history of operating system upgrades.

This example shows that the appliance was imaged with version 1.0.3 with no upgrades:

```

IMAGE_VERSION           : 1.0.3
IMAGE_CREATION_DATE     : Sun Apr 1 20:00:43 PDT 2012
IMAGING_START_DATE      : Wed Apr 4 16:57:59 UTC 2012
IMAGING_END_DATE        : Wed Apr 4 10:45:48 PDT 2012

```

listlinkup

Shows the Ethernet Bridge ports with active links. Run this command after connecting as root to a Sun Network QDR InfiniBand Gateway Switch.

This example shows three active ports (0A-ETH-1, 0A-ETH-3, and 0A-ETH-4) out of the eight available ports on switch bda1sw-ib3:

```

[root@bda1sw-ib3 ~]# listlinkup | grep Bridge
Bridge-0 Port 0A-ETH-1 (Bridge-0-2) up (Enabled)
Bridge-0 Port 0A-ETH-2 (Bridge-0-2) down (Enabled)
Bridge-0 Port 0A-ETH-3 (Bridge-0-1) up (Enabled)
Bridge-0 Port 0A-ETH-4 (Bridge-0-1) up (Enabled)
Bridge-1 Port 1A-ETH-1 (Bridge-1-2) down (Enabled)
Bridge-1 Port 1A-ETH-2 (Bridge-1-2) down (Enabled)
Bridge-1 Port 1A-ETH-3 (Bridge-1-1) down (Enabled)
Bridge-1 Port 1A-ETH-4 (Bridge-1-1) down (Enabled)

```

showvlan

Lists the VLANs configured on the switch. Run this command after connecting as root to a Sun Network QDR InfiniBand Gateway Switch.

This example shows the default VLAN, which has an ID of 0, on switch bda1sw-ib3:

```
[root@bda1sw-ib3 ~]# showvlan
Connector/LAG  VLN  PKEY
-----
0A-ETH-1      0    ffff
0A-ETH-3      0    ffff
0A-ETH-4      0    ffff
```

showvnics

Lists the Virtual Network Interface Cards (VNICs) created for the switch. Run this command after connecting as root to a Sun Network QDR InfiniBand Gateway Switch.

This example shows the VNICs created in a round-robin process for switch bda1sw-ib3:

```
[root@bda1sw-ib3 ~]# showvnics
ID  STATE  FLG  IOA_GUID          NODE          IID  MAC          VLN  PKEY  GW
-----
561 UP      N    0021280001CF4C23  bda1node13  BDA  192.168.41.31  0000 CE:4C:23:85:2B:0A NO  ffff  0A-ETH-1
564 UP      N    0021280001CF4C53  bda1node16  BDA  192.168.41.34  0000 CE:4C:53:85:2B:0D NO  ffff  0A-ETH-1
567 UP      N    0021280001CF4B58  bda1node01  BDA  192.168.41.19  0000 CE:4B:58:85:2A:FC NO  ffff  0A-ETH-1
555 UP      N    0021280001CF2A5C  bda1node07  BDA  192.168.41.25  0000 CE:2A:5C:85:2B:04 NO  ffff  0A-ETH-1
552 UP      N    0021280001CF4C74  bda1node04  BDA  192.168.41.22  0000 CE:4C:74:85:2B:01 NO  ffff  0A-ETH-1
558 UP      N    0021280001CF179B  bda1node10  BDA  192.168.41.28  0000 CE:17:9B:85:2B:07 NO  ffff  0A-ETH-1
.
.
.
```

Connecting Multiple Oracle Big Data Appliance Racks

This chapter describes how combine multiple Oracle Big Data Appliance racks into one super cluster. It contains the following topics:

- [Extending a Rack by Adding Another Rack](#)
- [Cabling Two Racks Together](#)
- [Cabling Several Racks Together](#)

Extending a Rack by Adding Another Rack

When creating a multirack Hadoop cluster or providing access to Oracle Big Data Appliance from an Oracle Exadata Database Machine, you must connect multiple racks to each other. Racks can be cabled together with no downtime.

During the cabling procedure, note the following:

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

See Also:

- [Appendix B, "Cabling Tables"](#) for information about ports and cables for the racks
- *Sun Datacenter InfiniBand Switch 36 User's Guide* for monitoring the subnet manager at
<http://docs.oracle.com/cd/E19197-01/835-0784-05/z4001de62024434.html>

Cabling Two Racks Together

The following procedure describes how to cable two racks together. 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 InfiniBand switch on the active system.
 - b. Use the `getmaster` command to verify 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 step b to ensure the Subnet Manager Master is running on the spine switch.
2. Ensure the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
3. Completely shutdown the new rack (R2).
4. Cable the leaf switch in the new rack according to [Table C-2](#).
5. Power off leaf switch R1 IB2. This causes all the database servers and Sun Fire Servers to failover their InfiniBand traffic to R1 IB3.
6. Disconnect all inter-switch links between R1IB2 and R1 IB3.
7. Cable leaf switch R1 IB2 according to [Table C-1](#).
8. Power on leaf switch R1 IB2.
9. 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.
10. Power off leaf switch R1 IB3. This causes all the database servers and Sun Fire Servers to failover their InfiniBand traffic to R1 IB2.
11. Cable leaf switch R1 IB3 according to [Table C-1](#).
12. Power on leaf switch R1 IB3.
13. 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.

14. Power on all the InfiniBand switches in R2.
 15. 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.
 16. Ensure the Subnet Manager Master is running on R1 IB1 by running the `getmaster` command from any switch.
 17. Power on all servers in R2.
 18. 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.
 19. Ensure Subnet Manager Master is running on one of the spine switches.
- After cabling the racks together, proceed to configure the racks.

Cabling Several Racks Together

The following procedure describes how to cable several racks together. This procedure assumes that the racks are adjacent to each other. In the procedure, the existing racks are R1, R2,... Rn, the new rack is Rn+1, and the Subnet Manager Master is running on R1 IB1.

1. Set the priority of the current, active Subnet Manager Master to 10 on the spine switch, as follows:
 - a. Log in to any InfiniBand switch on the active system.
 - b. Use the `getmaster` command to verify that the Subnet Manager Master is running on the spine switch.
 - c. Log in to the spine switch.
 - d. Use the `disableesm` command to stop Subnet Manager.
 - e. Use the `setsmpriority 10` command to set the priority to 10.
 - f. Use the `enableesm` command to restart Subnet Manager.
 - g. Repeat step b to ensure the Subnet Manager Master is running on the spine switch.
2. Ensure the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
3. Completely shutdown the new rack (Rn+1).
4. Cable the leaf switch in the new rack according to the appropriate table in [Appendix C, "Multi-Rack Cabling Tables"](#). For example, if rack Rn+1 was R4, then use [Table C-9, "Leaf Switch Connections for the Fourth Rack in a Four Rack System"](#).
5. Complete the following procedure for each of the original racks:

- a. Power off leaf switch Rx IB2. This causes all the database servers and Sun Fire Servers to failover their InfiniBand traffic to Rx IB3.
 - b. Cable leaf switch Rx IB2 according to [Appendix C, "Multi-Rack Cabling Tables"](#).
 - c. Power on leaf switch Rx IB2.
 - d. Wait for 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,... Rn.
 - e. Power off leaf switch Rx IB3. This causes all the database servers and Sun Fire Servers to failover their InfiniBand traffic to Rx IB2.
 - f. Cable leaf switch Rx IB3 according to [Appendix C, "Multi-Rack Cabling Tables"](#).
 - g. Power on leaf switch Rx IB3.
 - h. Wait for 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,... Rn.
All racks should now be rewired according to [Appendix C, "Multi-Rack Cabling Tables"](#).
6. Power on all the InfiniBand switches in Rn+1.
 7. 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 $(n+1)*3$ switches for IB1, IB2, and IB3 in racks R1, R2,... Rn+1.
 8. Ensure the Subnet Manager Master is running on R1 IB1 by running the `getmaster` command from any switch.
 9. Power on all servers in Rn+1.
 10. 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.
 11. Ensure Subnet Manager Master is running on one of the spine switches using the `getmaster` command from any switch.
 12. Ensure Subnet Manager is running on every spine switch using 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 Subnet Manager using the `enableesm` 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 `disableesm` command.

Configuring Ethernet Over InfiniBand

This chapter describes how to set up Ethernet over InfiniBand (EoIB) network configuration for Oracle Big Data Appliance.

This chapter discusses the following topics:

- [Introduction to Virtual NIC \(VNIC\)](#)
- [Setting Up Ethernet Over InfiniBand \(EoIB\)](#)

Introduction to Virtual NIC (VNIC)

A Virtual Network Interface Card (VNIC) maps an Ethernet connector on the Sun Network QDR InfiniBand Gateway Switch to a network interface within the server. A connector (0A-ETH-1 to 0A-ETH-4, and 1A-ETH-1 to 1A-ETH-4) hosts a 10 GbE port. On each Sun Network QDR InfiniBand Gateway Switch on Oracle Big Data Appliance, you can have a maximum of eight 10 GbE uplinks. In most scenarios, the number of 10 GbE uplinks is less than 8.

The network interface on the server, which is created by creating a VNIC on the Sun Network QDR InfiniBand Gateway Switch, uses the `ethx` naming convention, such as `eth8` and `eth9`. You can assign multiple VNICs to a server. In that case, they can be different connectors on the same Sun Network QDR InfiniBand Gateway Switch or on different ones. The default configuration maps one connector on each of the two switches and bonds or balances the pair for increased availability.

In Oracle Big Data Appliance, you create VNICs in Sun Network QDR InfiniBand Gateway Switch manually.

Setting Up Ethernet Over InfiniBand (EoIB)

To set up Ethernet over InfiniBand connectivity:

1. Use an SSH client, such as PuTTY, to log in to a Sun Network QDR InfiniBand Gateway Switch as a root. For example, log in to `bda1sw-ib2` as root.
2. At the command prompt, run the following command:

```
# listlinkup | grep Bridge
```

This is an example of the output of this command:

```
Bridge-0 Port 0A-ETH-1 (Bridge-0-2) up (Enabled)
Bridge-0 Port 0A-ETH-2 (Bridge-0-2) down (Enabled)
Bridge-0 Port 0A-ETH-3 (Bridge-0-1) down (Enabled)
Bridge-0 Port 0A-ETH-4 (Bridge-0-1) down (Enabled)
Bridge-1 Port 1A-ETH-1 (Bridge-1-2) down (Enabled)
```

```
Bridge-1 Port 1A-ETH-2 (Bridge-1-2) down (Enabled)
Bridge-1 Port 1A-ETH-3 (Bridge-1-1) down (Enabled)
Bridge-1 Port 1A-ETH-4 (Bridge-1-1) down (Enabled)
```

In this example, the uplink is 0A-ETH-1; all others are down.

3. Determine the GUIDs of a Oracle Big Data Appliance server as follows:
 - a. On the server that requires the VNIC, log in as `root`, and run the `ibstat` command on the command line. For example, log in to `bda1node01` as `root`. This command displays port information, as in the following example output:

```
# ibstat
CA 'mlx4_0'
  CA type: MT26428
  Number of ports: 2
  Firmware version: 2.9.1000
  Hardware version: b0
  Node GUID: 0x0021280001cf4b56
  System image GUID: 0x0021280001cf4b59
  Port 1:
    State: Active
    Physical state: LinkUp
    Rate: 40
    Base lid: 137
    LMC: 0
    SM lid: 147
    Capability mask: 0x02510868
    Port GUID: 0x0021280001cf4b57
    Link layer: IB
  Port 2:
    State: Active
    Physical state: LinkUp
    Rate: 40
    Base lid: 138
    LMC: 0
    SM lid: 147
    Capability mask: 0x02510868
    Port GUID: 0x0021280001cf4b58
    Link layer: IB
```

In the output, information for two ports is displayed. From this example, you must determine which port GUID to use. Before doing so, write down the two Base lid values shown in the output. In the example output, Port 1 has a Base lid value of 137, and Port 2 has a Base lid value of 138.

- b. On the same server, run the following command on the command line to report information about all active links in the InfiniBand fabric:

```
# iblinkinfo.pl -R | grep hostname
```

Where *hostname* is the name of the server. For example, `bda1node01`. You can also specify the bonded IPoIB address of the server.

The following is the example output of this command:

```
148 6[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 138 2[ ] "bda1node01
HCA-1" ( )
149 6[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 137 1[ ] "bda1node01
HCA-1" ( )
```

From this example output, note the switch lid values. Port 2 base lid of 138 is associated with switch lid 148 (the first column in the output). Port 1 base lid of 137 is associated with switch lid 149.

4. Determine which gateway switch is associated with the switch lids (148 and 149) by comparing the first column of the `iblinkinfo` output to the lid value of the `ibswitches` command as follows:

- a. On the server, run the `ibswitches` command. The example output of this command follows.

```
# ibswitches
Switch : 0x002128df0f0ac0a0 ports 36 "SUN IB QDR GW switch bda1sw-ib3
10.133.43.37" enhanced port 0 lid 148 lmc 0
Switch : 0x0021284690eea0a0 ports 36 "SUN DCS 36P QDR bda1sw-ib1
10.133.43.35" enhanced port 0 lid 147 lmc 0
Switch : 0x002128df348ac0a0 ports 36 "SUN IB QDR GW switch bda1sw-ib2
10.133.43.36" enhanced port 0 lid 149 lmc 0
```

- b. In this example output, identify the switches that lid values 148 and 149 are associated with.

Lid 148 is associated with gateway switch `bda1sw-ib3` is associated with and GUID `0x002128df0f0ac0a0`.

Lid 149 is associated with gateway switch `bda1sw-ib2` and GUID `0x002128df348ac0a0`. The following examples use this switch.

5. Define a dummy Media Access Control (MAC) address in the following format:

last three octets from bda1sw-ib2 switch ib GUID : last three octets of the administrative IP of the server in hexadecimal

For example, the `bda1sw-ib2` switch GUID is `00:21:28:df:34:8a:c0:a0` (`0x002128df348ac0a0`). The last three octets are `8a:c0:a0`.

Assume the administrative IP of the server that requires the VNIC is `192.168.1.1`. The last three octets of the IP address is `168.1.1`. Convert each decimal to hexadecimal, and it is `a8:1:1`.

Then you can define the MAC address as follows:

```
8a:c0:a0:a8.1.1
```

Note: Ensure the MAC address is unique. Only even numbers are supported for the most significant byte of the MAC address (unicast). The address used here is an example only.

6. As root, log in to the switch identified in Step 4. Use its IP address or host name to log in. In this example, the switch is `bda1sw-ib2`.

7. Issue this command to create a VNIC:

```
# createvnic 0A-ETH-1 -guid 00:21:28:df:34:8a:c0:a0 -mac
8a:c0:a0:a8.1.1 -pkey default
```

This command does not associate the VNIC with a VLAN. See [Chapter 11, "Creating Virtual LANs,"](#) for information about creating VLANs.

Oracle Big Data Appliance uses a single partition (the default partition).

8. To verify the VNIC, run the `showvnics` command. The following example output is displayed:

```
# showvnic
ID STATE FLG IOA_GUID          NODE          IID  MAC          VLN PKEY
GW
-----
-----
228 UP      N 0021280001CF4B57 bda1node01 0000 CE:4B:57:85:2A:FC NO ffff
0A-ETH-1
```

9. On the server, issue this command to display the list of VNICs available on the compute node:

```
# mlx4_vnic_info --list
eth8
```

10. Repeat the previous steps to another VNIC on the same server. The second VNIC appears on the server as eth9.

Creating a Bonded Interface

You must configure these two Ethernet over InfiniBand (EoIB) interfaces as a bonded interface named `bondeth0`. This procedure requires that you manually create three files:

- `ifcfg-eth8`: Describes the eth8 connection
- `ifcfg-eth9`: Describes the eth9 connection
- `ifcfg-bondeth0`: Bonds the two connections in active-passive failover mode

To create a bonded interface from the EoIB interfaces:

1. Log in as root to the same server where you created the VNICs and change to the `/etc/sysconfig/network-scripts` directory.
2. Using a text editor such as `vi`, create or edit a file named `ifcfg-eth8` with the following contents. Replace `hostname` with the name of the server, and `address` with the dummy MAC address created in Step 5 of the previous procedure.

```
#ifcfg_eth8 for host hostname
DEVICE=eth8
MASTER=bondeth0
SLAVE=yes
ONBOOT=yes
MTU=65520
BOOTPROTO=none
HOTPLUG=no
IPV6INIT=no
HWADDR=address
```

3. Create or edit a second file named `ifcfg-eth9` for eth9, using its dummy MAC address.
4. Create or edit a file named `ifcfg-bondeth0` with the following parameters. Verify that the settings are correct for your network.

```
#ifcfg-bondeth0 for server bda1node01 generated by BDAeioib.g
DEVICE=bondeth0
TYPE=ETHERNET
BROADCAST=203.0.115.255
ONBOOT=yes
#MTU=65520
BOOTPROTO=none
NETMASK=255.255.255.0
```

```
IPV6INIT=no  
NETWORK=203.0.114.0  
IPADDR=203.0.113.10  
BONDING_OPTS="mode=active-backup miimon=100 downdelay=5000 updelay=5000"
```

5. Start the new bondeth0 interface and restart the server for the changes to take effect.

Creating Virtual LANs

This chapter describes how to set up a Virtual LAN (VLAN) on the Ethernet connector of a Sun Network QDR InfiniBand Gateway Switch, and it describes how to assign a Virtual NIC (VNIC) on a server to use that VLAN.

The tasks described in this chapter are optional.

It contains the following topics:

- [Introduction to VLAN](#)
- [Tagging Ethernet Connectors With a VLAN Identifier](#)
- [Creating VNICs and Associating Them with VLANs](#)
- [For Additional Information](#)

See Also: [Chapter 10, "Configuring Ethernet Over InfiniBand"](#)

Introduction to VLAN

The Ethernet standard has a provision to combine multiple broadcast domains, and thus IP subnets, onto a single Ethernet cable using a Virtual LAN (IEEE 802.1Q VLAN) configuration. To use VLANs, both ends of the Ethernet link must be configured to support the defined VLANs. The benefits include a logical division of workload, enforcing security isolation, and splitting traffic across several manageable broadcast domains. VLANs allow traffic separation from the 10 GbE switch to the servers. By design, Ethernet traffic on one VLAN cannot be seen by any host on a different VLAN. You can use an external router to enable communication between two VLANs.

Tagging Ethernet Connectors With a VLAN Identifier

To tag an Ethernet connector on the gateway switch with a VLAN identifier, you must issue the `createvlan` command on the gateway switch that the VLAN will be associated with.

In this process, you map the following:

- Ethernet connector on the Sun Network QDR InfiniBand Gateway Switch (0A-ETH-1 to 0A-ETH-4, and 1A-ETH-1 to 1A-ETH-4)
- VLAN ID (up to 4094)
- InfiniBand partition key (0xFFFF)

Oracle Big Data Appliance uses the default partition and partition key.

You can associate multiple VLANs to a single Ethernet connector. For example, you can associate VLAN identifiers 10 and 11 to the same Ethernet connector 0A-ETH-1.

To do so, run the `createvlan` command, as in the following example:

1. Log in to the gateway switch interface as `root`, and run the following commands:

```
# createvlan 0A-ETH-1 -VLAN 10 -PKEY default
```

Where 0A-ETH-1 is the Ethernet connector on the gateway switch and 10 is the VLAN identifier.

```
# createvlan 0A-ETH-1 -VLAN 11 -PKEY default
```

Where 0A-ETH-1 is the Ethernet connector and 11 is the VLAN identifier.

2. To verify, run the following command:

```
# showvlan
```

The following information is displayed:

Connector/LAG	VLAN	PKEY
0A-ETH-1	0	ffff
1A-ETH-3	0	ffff

Creating VNICs and Associating Them with VLANs

If you plan to associate a VNIC with a VLAN, you should provide a VLAN identifier when mapping the MAC address, partition key, GUID, and Ethernet connector. In Oracle Big Data Appliance, the default partition key (0xFFFF) is used.

To create a VNIC and associate with a VLAN:

1. Take steps 1 through 6 in "[Setting Up Ethernet Over InfiniBand \(EoIB\)](#)" on page 10-1.
2. In step 7, add a `-vlan` parameter to the `createvnic` command to create the VNIC and associate it with a VLAN. This example creates a VNIC associated with VLAN 10.

```
# createvnic 0A-ETH-1 -guid 00:21:28:df:34:8a:c0:a0 -mac
8a:c0:a0:a8.1.1 -vlan 10 -pkey default
```

3. To create a second VNIC using the same Ethernet connector and GUI, repeat the previous `createvnic` command using different values for the `-mac` and `-vlan` parameters.
4. Continue the steps in "[Setting Up Ethernet Over InfiniBand \(EoIB\)](#)" on page 10-1, and repeat them for the second switch.
5. Create a bonded interface as described in "[Creating a Bonded Interface](#)" on page 10-4.

For Additional Information

For information about deleting VLANs, see "Controlling VLANs and VNICs" in the "Administering Gateway Resources" chapter of the *Sun Network QDR InfiniBand Gateway Switch Administration Guide* at

<http://docs.oracle.com/cd/E19671-01/>

Setting Up Auto Service Request

This chapter explains how to install and configure Auto Service Request for Oracle Big Data Appliance. It contains these topics:

- [Understanding Auto Service Request](#)
- [Getting Ready for ASR Installation](#)
- [Installing ASR Manager](#)
- [Registering ASR Manager](#)
- [Verifying ASR Manager](#)
- [Configuring the Trap Destinations](#)
- [Activating ASR Assets](#)
- [Validating Auto Service Request on Oracle Big Data Appliance](#)
- [Troubleshooting ASR](#)

Understanding Auto Service Request

Auto Service Request (ASR) is designed to automatically open service requests when specific Oracle Big Data Appliance hardware faults occur. To enable this feature, the Oracle Big Data Appliance components must be configured to send hardware fault telemetry to the ASR Manager software. This service covers components in Sun Fire Servers. ASR Manager must be installed on a server that has connectivity to Oracle Big Data Appliance and an outbound Internet connection using HTTPS or an HTTPS proxy.

Note: Set up ASR Manager before configuring Oracle Big Data Appliance. Ensure that Auto Service Request is selected in the Oracle Big Data Appliance Configuration Worksheets, so that ASR components are installed and configured on the appliance. See ["Software Configuration"](#) on page 4-7.

When a hardware problem is detected, 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 database administrator is even aware the problem exists.

ASR is applicable only for component faults. Not all component failures are covered, though the most common components such as disk, fan, and power supplies are covered.

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.

Before using Auto Service Request, set up the following:

- Oracle Premier Support for Systems or Oracle Premier Support for Operating Systems
- Technical contact responsible for Oracle Big Data Appliance
- Valid shipping address for Oracle Big Data Appliance parts

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

A service request may 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 customers continue to monitor their systems for faults and call Oracle Support Services if they do not receive notice that a service request has been filed automatically.

See Also:

Oracle Auto Service Request web page at

<http://www.oracle.com/technetwork/server-storage/asr/overview/index.html>

Oracle Auto Service Request User Documentation at

<http://www.oracle.com/technetwork/server-storage/asr/documentation/index.html>

Getting Ready for ASR Installation

Before installing ASR, complete these prerequisites:

1. Create a My Oracle Support (MOS) account at <http://support.oracle.com>.
2. Ensure the following are correctly set up:
 - Oracle Premier Support for Systems or Oracle Limited Warranty
 - Technical contact responsible for the Big Data Appliance
 - Valid shipping address for Big Data Appliance repair parts
3. Identify and designate a system to serve as ASR Manager.

ASR Manager must be installed on a server that has connectivity to Oracle Big Data 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. Ensure that the designated system conforms to the "Hardware and Network Configuration Recommendations for ASR." See the Oracle ASR website:

<http://www.oracle.com/technetwork/server-storage/asr/overview/hardware-recommendations-330108.html>

5. Confirm that JDK 6 (1.6.0_04 or later) is running on the designated ASR Manager system:

```
java -version
```

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

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

6. Obtain root access to the designated ASR Manager system.
7. Identify and verify connectivity to Oracle Big Data Appliance.
8. Ensure connectivity to the Internet using HTTPS.

Installing ASR Manager

This section contains these topics:

- [Downloading the Software](#)
- [Installing the ASM Package](#)
- [Installing the Oracle Service Tags Bundle](#)

Downloading the Software

ASR Manager requires installation of these components:

- ASR Manager Version 3.6 or higher
- Oracle Service Tags

For links to the ASR software, go to the Oracle Auto Service Request Downloads website:

<http://www.oracle.com/technetwork/server-storage/asr/downloads/index.html>

Refer to My Oracle Support note 1185493.1 for the most recent software releases.

To download the Oracle Service Tags bundle, go to this site:

<https://updates.oracle.com/download/12757884.html>

Installing the ASM Package

To install the ASM package:

1. Connect to the ASR Manager system.
2. Check whether ASM version 1.3.1 or later is already installed:

```
rpm -q SUNWsasm
```

If it is not installed, then download the latest version as described in ["Downloading the Software"](#) on page 12-3.

3. As the root user, install the ASM package:

```
rpm -i SUNWsasm.version_number.rpm
```

4. Verify the installation:

```
rpm -Vv SUNWsasm-version_number
```

Installing the Oracle Service Tags Bundle

To install the SUNWswasr package:

1. Connect to the ASR Manager system.
2. Download and unzip the SUNWswasr package.
3. As the root user, install the package:

```
rpm -i SUNWswasr.version_number.rpm
```

4. Update the root login file, such as `.profile` or `.bashrc`, to add the asr software to the PATH variable. These commands are for the bash shell:

```
PATH=$PATH:/opt/SUNWawasr/bin/asr
EXPORT PATH
```

5. Verify the installation:

```
rpm -Vv SUNWswasr
```

If the installation failed, then remove it, restart the system, and reinstall the package.

6. Ensure ASM is running for issuing any other ASR commands:

```
/opt/SUNWasm/bin/sasm status
```

Registering ASR Manager

To register ASR Manager:

1. As the root user on the ASR Manager server, issue this command:

```
asr register
```

2. Enter 1 or 2, depending on your location:
 - Americas or Asia Pacific regions: 1 for `transport.oracle.com`
 - Europe, Middle East, or Africa regions: 2 for `transport.oracle.co.uk`
3. If you are using a proxy server to access the Internet, then enter the proxy server information. If you are not using a proxy server, then enter a hyphen (-).
4. Enter your My Oracle Support (MOS) user name and password when prompted. ASR validates the login. After validation, the registration is complete. Passwords are not stored.

Your MOS email address receives output from ASR reports, notification of ASR problems, and service request (SR) generation.

Verifying ASR Manager

Perform these checks on ASR Manager to ensure it is installed properly:

- As the root user, verify that the ASR installer created three CRON jobs:

```
crontab -l
```

The following are the recommended settings:

- `asr report`: Set once a month on Sunday

- `asr heartbeat`: Set twice daily or at least once daily
- `asr update_rules.sh`: Set once daily by default
- 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
```

Configuring the Trap Destinations

The Oracle Big Data Appliance Mammoth Utility configures the servers automatically. See "[Mammoth Utility Steps](#)" on page 13-5.

Activating ASR Assets

To activate ASR assets:

1. On ASR Manager, verify ASR is activated:

```
asr list_asset -i asset_ip
```

In the preceding command, *asset_ip* is the IP address of a Sun Fire Server host or an ILOM. To list all assets, issue this command:

```
asr list_asset
```

The following is an example of the output. You should see a host name and an ILOM name for each of the 18 servers, for a total of 36 activations. If you see fewer than 36, identify the missing ones and activate them. An example of a host name is `bda1node15` and an ILOM name is `bda1node15-c`.

ADDRESS	HOST_NAME	SERIAL_NUMBER	ASR	PRODUCT_NAME
203.0.114.44	bda1node15	1143FMM023	Enabled	SUN FIRE X4270 M2 SERVER...
203.0.115.139	bda1node15-c	1143FMM073	Enabled	SUN FIRE X4270 M2 SERVER
203.0.114.45	bda1node16	1143FMM021	Enabled	SUN FIRE X4270 M2 SERVER...
203.0.115.140	bda1node16-c	1143FMM063	Enabled	SUN FIRE X4270 M2 SERVER
203.0.114.46	bda1node17	1143FMM021	Enabled	SUN FIRE X4270 M2 SERVER...
203.0.115.141	bda1node17-c	1143FMM46B	Enabled	SUN FIRE X4270 M2 SERVER

If no assets are listed, verify that all steps of the Mammoth Utility have run successfully. See [Chapter 13, "Installing the Oracle Big Data Appliance Software."](#)

2. Confirm end-to-end ASR functionality:

```
asr report
```

This is an example of the output:

```
Successfully submitted request for activation status report.
Activation status report will be sent to email address associated with
MOS Account:bdaadmin@example.com
```

The report is sent to My Oracle Support.

3. Activate the assets in My Oracle Support. Only a Customer User Administrator can perform this step.
 - a. Log in to My Oracle Support.
 - b. Click the **More** tab.
 - c. Select **Settings**.
 - d. Click **Pending ASR Activations**.
 - e. Select the asset to activate, and review the information in the ASR Activation - Asset window. Update any information as needed.
 - f. Click **Approve** to complete activation. The assets should be *Active*.
4. Confirm that the assets are listed as *Active* in My Oracle Support:
<http://support.oracle.com>

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

If no assets are discovered, take these steps:

1. Log in to the first server.
2. (Optional) Activate the ASR Manager host using the following command:

```
asr activate_asset -i host_ip
```

Where *host_ip* is the host IP address. If the server is qualified for ASR and entitled to service, then it can be activated.

3. Activate the Integrated Lights Out Managers (ILOMs):

```
asr activate_asset -i ilom_ip
```

In the preceding command, *ilom_ip* is the ILOM IP address.

4. Activate the Sun Fire Servers:

```
asr activate_asset -i host_eth0_ip
```

Where *host_eth0_ip* is the IP address of the first Sun Fire Server on the administrative network.

5. Repeat these `activate_asset` commands on each Sun Fire Server.

Validating Auto Service Request on Oracle Big Data Appliance

To validate the ASR installation, you must generate test events.

To set up a test event on an ILOM:

1. Log in to the ILOM of any server in Oracle Big Data Appliance.
2. Change to the following directory:

```
cd /SP/alertmgmt/rules/3
```

3. Issue this command:

```
show
```

The following information is displayed:

```
/SP/alertmgmt/rules/3
Targets:

Properties:
  type = snmptrap
  level = minor
  destination = 10.10.10.255 (Use the IP of your own ASR manager!)
  destination_port = 162
  community_or_username = public
  snmp_version = 2c
  testrule = (Cannot show property)

Commands:
  cd
  set
  show
```

4. Ensure the properties have valid values as shown in step 3, not blank or set to zeros. Use the appropriate IP address and port for your site.
5. Set the test trap:


```
set testrule=true
```
6. You should receive an email about the event. Verify that an email was also sent to the address listed in your Customer Service Identifier (CSI) for Oracle Big Data Appliance.

To set up a test in the operating system environment:

1. Issue this command to validate the operating system. The sample output shows 10.10.10.123:162 as the ASR Manager address.

```
# /opt/oracle/bda/compmon**/bda_mon_hw_asr.pl -validate_snmp_subscribers -type
asr
Sending test trap to destination - 10.10.10.123:162
```

2. Verify that an email about the event was sent to you and to the address listed in your Customer Service Identifier (CSI) for Oracle Big Data Appliance.

This test event does not open a case, but makes a verifiable entry in the ASR back end.

You received the emails from both tests, then your validation is complete. If not, then proceed to "[Troubleshooting ASR](#)" on page 12-7.

If the emails state that there are contract issues, contact your installation coordinator for further assistance.

Troubleshooting ASR

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

<http://www.oracle.com/technetwork/server-storage/asr/documentation/index.html>

If you continue to have issues, contact ASR support. See My Oracle Support Note 1352349.1)

Installing the Oracle Big Data Appliance Software

This chapter explains how to use the Mammoth Utility to install the software on Oracle Big Data Appliance. It contains these topics:

- [Using the Mammoth Utility](#)
- [Installing the Software on a Single or Primary Rack](#)
- [Mammoth Utility Syntax](#)
- [What If an Error Occurs During the Installation?](#)
- [Adding a Rack to an Existing Cluster](#)
- [Mammoth Utility Steps](#)

Using the Mammoth Utility

The Mammoth Utility installs and configures the software on Oracle Big Data Appliance using the files generated by the Oracle Big Data Appliance Configuration Utility. At a minimum, Mammoth installs and configures Cloudera's Distribution including Apache Hadoop. This includes all the Hadoop software and Cloudera Manager, which is the tool for administering your Hadoop cluster. Mammoth will optionally install and configure Oracle NoSQL Database and, if you have a license, all components of Oracle Big Data Connectors.

In addition to installing the software across all servers in the rack, the Mammoth Utility creates the required user accounts, starts the correct services, and sets the appropriate configuration parameters. When it is done, you have a fully functional, highly tuned, up and running Hadoop cluster.

You must run the Mammoth Utility once for each rack.

- For one Oracle Big Data Appliance Rack that forms one Hadoop cluster, follow the procedure in ["Installing the Software on a Single or Primary Rack"](#) on page 13-2.
- For multiple racks where each rack forms an independent Hadoop cluster, follow the procedure in ["Installing the Software on a Single or Primary Rack"](#) on page 13-2 for each rack.
- For multiple racks that form a single, multirack Hadoop cluster:
 - Identify the primary rack of the cluster, then follow the procedure in ["Installing the Software on a Single or Primary Rack"](#) on page 13-2.
 - For the other racks of the cluster, follow the procedure in [Adding a Rack to an Existing Cluster](#).

Installing the Software on a Single or Primary Rack

Follow this procedure to install and configure the software on a single Oracle Big Data Appliance Rack or on the primary rack of a multiple-rack cluster.

To install the software:

1. Verify that the Oracle Big Data Appliance Rack is configured according to the custom network settings described in `/opt/oracle/bda/BdaDeploy.json`. If the rack is still configured to the factory default IP addresses, first perform the network configuration steps described in ["Configuring the Network"](#) on page 8-18.
2. Verify that the software is not installed on the rack already. If the software is installed and you want to reinstall it:

Caution: These steps result in the loss of all data stored in HDFS. Consider your options carefully before proceeding.

- Use the `mammoth -u` option, described in ["Mammoth Utility Syntax"](#) on page 13-2.
 - If `mammoth -u` fails, then reimage the entire rack using the `reimagerack` utility. See ["Checking the Health of the Network"](#) on page 8-25.
3. Copy `BDAMammoth-version.run` to any directory on node01 (such as `/tmp`). You can download this file from the same location as the base image file. See the procedures under ["Reinstalling the Base Image"](#) on page 8-23.
 4. Log in to node01 as `root` and decompress the `BDAMammoth-version.run` self-extracting file. This example extracts Mammoth version 1.0.3 in the `/tmp` directory:


```
/tmp/BDAMammoth-1.0.3.run
```
 5. Change directories:


```
cd /opt/oracle/BDAMammoth
```
 6. Copy `mammoth-rack_name.params` to the current directory. See ["About the Configuration Files"](#) on page 4-3.
 7. Run the `mammoth` command with the appropriate option. See [Table 13-1](#). This sample command runs steps 1 and 2 on rack bda2:


```
./mammoth -r 1-2 bda2
```

The Mammoth Utility stores the current configuration in the `/opt/oracle/bda/install/state` directory. *Do not delete the files in this directory.* The Mammoth Utility fails without this information if you need to use it again, such as adding a rack to the cluster.

Mammoth Utility Syntax

You must change to the `/opt/oracle/BDAMammoth` directory to use the Mammoth Utility. It has this syntax:

```
./mammoth option [rack_name]
```

Where `rack_name` is the name of an Oracle Big Data Appliance Rack. You must enter the rack name in the first command exactly as it appears in the configuration file name

(`mammoth-rack_name.params`). Afterward, `rack_name` defaults to the rack specified in a previous `mammoth` command. [Table 13–1](#) lists the options.

You must finish installing one rack before starting the installation of another rack.

Table 13–1 Mammoth Utility Options

Option	Description
-h	Displays command Help including command usage and a list of steps.
-i	Runs all mandatory steps, equivalent to <code>-r 1-18</code> .
-l	List the steps of the Mammoth Utility.
-r <i>n-N</i>	Run steps <i>n</i> through <i>N</i> of the Mammoth Utility while no errors occur
-s <i>n</i>	Runs step <i>n</i> .
-u	Uninstalls all software from all racks in an existing Hadoop cluster. <i>This option results in a loss of all data.</i>
-v	Displays the version number of the Mammoth Utility.

Example 13–1 Mammoth Utility Syntax Examples

This command displays Help for the Mammoth Utility:

```
./mammoth -h
```

This command does a complete install on rack `bda3`:

```
./mammoth -i bda3
```

The next command runs steps 2 through 6 on the rack being set up:

```
./mammoth -r 2-6
```

What If an Error Occurs During the Installation?

Each step generates a detailed log file listing the actions performed on each server and whether the step completed successfully. If an error occurs, the script stops. You can then check the log files in `/opt/oracle/BDAMammoth/bdaconfig/tmp`. The log files are named in this format:

```
pagent-server-yyyyymmddhhmmss.log
```

Where *server* is the name of the server and `yyyyymmddhhmmss` identifies the year, month, day, hour, minute, and second that the file was created.

After fixing the problem, you can rerun all steps or a range of steps. You cannot skip steps or run them out of order.

Adding a Rack to an Existing Cluster

Each multirack cluster has one rack designated as the primary rack. Whether a rack is the primary one is indicated in the Oracle Big Data Appliance Configuration Worksheets and specified in the `mammoth-rack_name.params` file. Each rack of a multirack Hadoop cluster has a separate `mammoth-rack_name.params` file. The following procedure describes how to install the software on additional racks that are part of the same cluster:

1. Install the software on the primary rack of the Hadoop cluster. See ["Installing the Software on a Single or Primary Rack"](#) on page 13-2.

2. Ensure that all racks are running the same software version. See ["About Software Version Differences"](#) on page 13-4.
3. Ensure that all racks that form a single Hadoop cluster are cabled together. See [Chapter 9, "Connecting Multiple Oracle Big Data Appliance Racks."](#)
4. Copy the `mammoth-rack_name.params` files of the non-primary racks to node01 (the bottom server) of the primary rack. Do not copy them to the non-primary racks.
5. Connect as `root` to node01 of the primary rack and change to the BDAMammoth directory:

```
cd /opt/oracle/BDAMammoth
```

Note: Always start Mammoth from the primary rack.

6. For each non-primary rack, issue the `mammoth` command with the appropriate option. See ["Mammoth Utility Syntax"](#) on page 13-2. For example, this command starts the installation on rack bda4:

```
./mammoth -i bda4
```

The primary rack of a multirack Hadoop cluster is configured the same as a single Hadoop cluster. It runs the NameNode, Secondary Name Node, Hue, Hive, and other key services. The other racks of a multirack Hadoop cluster are configured differently. They only run the DataNodes and TaskTrackers.

Oracle Big Data Connectors are installed on all nodes of the non-primary racks although no services run on them. Oracle Data Integrator agent still runs on node03 of the primary rack. You cannot add nodes to an Oracle NoSQL Database cluster after it is set up. However, a logical volume is created on the additional rack for future use when nodes can be added to an Oracle NoSQL Database cluster.

The Mammoth Utility obtains the current configuration from the files stored in `/opt/oracle/bda/install/state`. If those files are missing or if any of the services have been moved manually to run on other nodes, then the Mammoth Utility fails.

A new Oracle Big Data Appliance Rack may be factory-installed with a newer image than the previously installed racks. All racks configured as one Hadoop cluster must have the same image. When all racks have the same image, you can install the software on the new rack.

About Software Version Differences

A new Oracle Big Data Appliance Rack may be factory-installed with a newer base image than the previously installed racks. Use the `imageinfo` utility on any server to get the image version. Only when all racks of a single Hadoop cluster have the same image version can you proceed to install the software on the new rack.

To synchronize the new rack with the rest of the Hadoop cluster, either upgrade the existing cluster to the latest image version or downgrade the image version of the new rack.

To downgrade the image version:

- Reimage the new rack to the older version installed on the cluster. See My Oracle Support Master Note 1434477.1 and its related notes.
- Use the older version of the Oracle Big Data Appliance Configuration Utility to generate the configuration files.
- Use the older version of the Mammoth Utility to install the software.

Mammoth Utility Steps

Following are descriptions of the steps that the Mammoth Utility performs when installing the software.

Step 1 SetupInstall

Validates the configuration files.

Step 2 WriteNodelist

Displays a road map of the planned system, including this information:

- Location of important nodes in the system
- List of partitions for HDFS directories
- List of disks reserved for Oracle NoSQL Database
- Ports used by various components
- Location of services
- User names and initial passwords

This step, and every subsequent step, stores information in a file named `/opt/oracle/bda/environment.pp`. Check the contents of this file now to ensure that the environment generated by the Mammoth Utility appears correct.

This step also generates a file named `passwords.pp`. It contains the passwords for various software components that run under an operating system user identity. Operating system root passwords are not written to disk. The last step of the installation removes `passwords.pp`.

Step 3 SetupSSHroot

Sets up a Secure Shell (SSH) for the root user so you can connect to all addresses on the administrative network without entering a password.

Step 4 UpdateEtcHosts

This step performs several tasks:

Generates `/etc/hosts` from the configuration file and copies it to all servers so they use the InfiniBand connections to communicate internally. The file maps private IP addresses to public host names.

Sets up passwordless SSH for the root user on the InfiniBand network.

Sets up an alias to identify the node where the Mammoth Utility is run as the puppet master node. For example, if you run the Mammoth Utility from `bda1node01` with an IP address `192.168.41.1`, then a list of aliases for that IP address includes `bda1node01-master`. The Mammoth Utility uses Puppet for the software installation; the next step describes Puppet in more detail.

Checks the network timing on all nodes. If the timing checks fail, then there are unresolved names and IP addresses that will prevent the installation from running correctly. Fix these issues before continuing with the installation.

Step 5 SetupPuppet

This step configures puppet agents on all nodes and start them, configures a puppet master on the node where the Mammoth Utility is being run, waits for the agents to submit their certificates, and automates their signing. After this step is completed, Puppet can deploy the software.

Puppet is a distributed configuration management tool that is commonly used for managing Hadoop clusters. The puppet master is a parent service and maintains a Puppet repository. A puppet agent operates on each Hadoop node.

A file named `/etc/puppet/puppet.conf` resides on every server and identifies the location of the puppet master.

Puppet operates in two modes:

- Periodic pull mode in which the puppet agents periodically contact the puppet master and asks for an update, or
- Kick mode in which the puppet master alerts the puppet agents that a configuration update is available, and the agents then ask for the update. Puppet operates in kick mode during the Mammoth Utility installation.

In both modes, the puppet master must trust the agent. To establish this trust, the agent sends a certificate to the puppet master node where the sys admin process signs it. When this transaction is complete, the puppet master sends the new configuration to the agent.

For subsequent steps, you can check the Puppet log files on each server, as described in ["What If an Error Occurs During the Installation?"](#) on page 13-3.

Step 6 PatchFactoryImage

Installs the most recent Oracle Big Data Appliance image and system parameter settings.

Step 7 CopyLicenseFiles

Copies third-party licenses to `/opt/oss/src/OSSLicenses.pdf` on every server, as required by the licensing agreements.

Step 8 CopySoftwareSource

Copies third-party software source code to `/opt/oss/src/` on every server, as required by the licensing agreements.

Step 9 CreateLogicalVolumes

Creates a logical volume if physical disks are allocated to Oracle NoSQL Database. This step varies depending on the amount of disk space allocated to Oracle NoSQL Database during configuration:

- **0 terabytes:** This step does nothing.
- **54 terabytes:** The disk space is allocated across the cluster using one disk on each node. The disk mounted at `/u12` is used for the logical volume.
- **108 terabytes:** The disk space is allocated across the cluster using two disks on each node. The disks mounted at `/u11` and `/u12` are used for the logical volume.

The logical volume is mounted at `/lv1`.

After this step finishes, the Linux file systems table in `/etc/fstab` shows the logical disks instead of the physical disks they represent.

Step 10 CreateUsers

Creates the `hdfs` and `mapred` users, and the `hadoop` group. It also creates the `oracle` user and the `dba` and `oinstall` groups.

The various packages installed in later steps also create users and groups during their installation.

See Also: *Oracle Big Data Appliance Software User's Guide* for more information about users and groups.

Step 11 SetupMountPoints

The NameNode and Secondary Name Node data is copied to multiple places to prevent a loss of this critical information should a failure occur in either the disk or the entire node where they are set up. The data is replicated during normal operation as follows:

- The Name Node and Secondary Name Node data is written to a partition that is mirrored so the loss of a single disk can be tolerated. This mirroring is done at the factory as part of the operating system installation.
- This step creates a directory named `/opt/exportdir` on node04 and mounts it on the Name Node and Secondary Name Node. It also exports `/opt/exportdir` from node04 and mounts it at `/opt/shareddir` on all nodes of the cluster. During operation of Oracle Big Data Appliance, the Name Node and Secondary Name Node data is also written to `/opt/exportdir`.
- Optionally, this step mounts on the Name Node and Secondary Name Node a directory on an external server so that the data is written there also. The external server and directory must be identified for this purpose in the *Oracle Big Data Appliance Configuration Worksheets*. You can examine this configuration setting by looking at the value of `$external_dir_path` in `/opt/oracle/bda/puppet/manifests/environment.pp`.

Mammoth checks for these requirements:

- Under the specified directory path, a subdirectory must exist with the same name as the cluster. This subdirectory must be owned by `root`.
- Under this subdirectory, two subdirectories named `nn` and `snn` must exist and be owned by user `hdfs` and group `hadoop`. The `hdfs` UID must be the same as the `hdfs` UID on Oracle Big Data Appliance, and the `hadoop` GID must be the same as the `hadoop` GID on Oracle Big Data Appliance.

For example, if the NFS directory is specified in `environment.pp` as

```
NFS_DIRECTORY=extfiler:/scratch/bda
```

and the cluster name is specified as

```
CLUSTER_NAME=bda1
```

then:

- The `/scratch/bda/bda1` directory must exist on `EXTFILER` and be owned by `root`.
- The `/scratch/bda/bda1/nn` and `/scratch/bda/bda1/snn` directories must exist on `EXTFILER` and be owned by `hdfs` in group `hadoop`.

Step 12 SetupMySQL

Installs and configures MySQL Database. This step creates the primary database and several databases for use by Cloudera Manager on node03. It also sets up replication of the primary database to a backup database on node02.

When this step is complete, you can open MySQL Database:

```
# mysql -uroot -p password
mysql> show databases
```

Step 13 InstallHadoop

Installs all packages in Cloudera's Distribution including Apache Hadoop (CDH) and Cloudera Manager. It then starts the Cloudera Manager server on node02 and configures the cluster.

Step 14 StartHadoopServices

Starts the agents on all nodes and starts all CDH services. After this step, you have a fully functional Hadoop installation.

Cloudera Manager runs on port 7180 of node02. You can open it in a browser, for example:

```
http://bda1node02.example.com:7180
```

Where bda1node02 is the name of node02 and example.com is the domain. The default user name and password is admin, which is changed in [Step 18, "CleanupInstall"](#).

Step 15 StartHiveService

Starts the Hive service on node03 and copies the Hadoop client configuration to /etc/hadoop/conf on all nodes.

Step 16 InstallBDASoftware

Installs Oracle NoSQL Database Community Edition and the server-side components of Oracle Big Data Connectors, if these options were selected in the Oracle Big Data Appliance Configuration Worksheets. Oracle NoSQL Database must be allocated disk space (54 or 108 TB) and Oracle Big Data Connectors must be licensed separately.

Step 17 SetupASR

Installs and configures Auto Service Request (ASR).

Note: For this step to run successfully, the ASR host system must be up with ASR Manager running and configured properly. See [Chapter 12, "Setting Up Auto Service Request"](#).

This step does the following:

- Installs the required software packages
- Configures the trap destinations
- Starts the monitoring daemon

To activate the assets from ASR Manager, see "[Activating ASR Assets](#)" on page 12-5.

Step 18 CleanupInstall

Performs the following:

- Changes the root password on all nodes (optional).
- Changes the Cloudera Manager password if specified in the Installation Template.
- Deletes temporary files created during the installation.
- Copies log files from all nodes to subdirectories in /opt/oracle/bda/install/log.

- Runs cluster verification checks, including TeraSort, to ensure that everything is working properly. It also generates an install summary. All logs are stored in a subdirectory under `/opt/oracle/bda/install/log` on node01.

Step 19 CleanupSSHroot (Optional)

Removes passwordless SSH for `root` that was set up in [Step 3](#).

Maintaining Oracle Big Data Appliance

This chapter describes how to monitor and maintain Oracle Big Data Appliance. Some of these procedures use the `dcli` utility to execute commands in parallel on all servers.

This chapter contains the following topics:

- [Monitoring Ambient Temperature of Servers](#)
- [Powering On and Off Oracle Big Data Appliance](#)
- [Adding Memory to a Server](#)
- [Maintaining the Physical Disks of Servers](#)
- [Changing InfiniBand IP Addresses](#)
- [Maintaining the InfiniBand Network](#)
- [Changing the NTP Servers](#)

See Also: [Chapter 7, "Using the dcli Utility"](#)

Monitoring Ambient Temperature of Servers

Maintaining environmental temperature conditions within design specification for a Sun Fire Server helps to achieve maximum efficiency and targeted component service lifetimes. The impact of validating the ambient temperature range is minimal. The impact of corrective actions vary depending on the environmental conditions.

Temperatures outside the ambient temperature range of 21 to 23 degrees Celsius (70 to 74 degrees Fahrenheit) affect all components within Oracle Big Data Appliance, possibly causing performance problems and shortened service lifetimes.

To monitor the ambient temperature:

1. Connect to an Oracle Big Data Appliance server as `root`.
2. Set up passwordless SSH for `root` by issuing the `setup-root-ssh` command, as described in "[Setting Up Passwordless SSH](#)" on page 7-1.

3. Check the current temperature:

```
dcli 'ipmitool sunoem cli "show /SYS/T_AMB" | grep value'
```

4. If any temperature reading is outside the operating range, then investigate and correct the problem. See [Table 2-9, "Temperature, Humidity and Altitude Requirements"](#).

The following is an example of the command output:

```
bda1node01-adm.us.oracle.com: value = 22.000 degree C
```

```
bda1node02-adm.us.oracle.com: value = 22.000 degree C
bda1node03-adm.us.oracle.com: value = 22.000 degree C
bda1node04-adm.us.oracle.com: value = 23.000 degree C
.
.
.
```

Powering On and Off Oracle Big Data Appliance

This section includes the following topics:

- [Non-Emergency Power Procedures](#)
- [Emergency Power-Off Considerations](#)
- [Cautions and Warnings](#)

Non-Emergency Power Procedures

This section contains the procedures for powering on and off the components of Oracle Big Data Appliance in an orderly fashion. This section contains the following topics:

- [Powering On Oracle Big Data Appliance](#)
- [Powering On Servers Remotely Using ILOM](#)
- [Powering Off Oracle Big Data Appliance](#)
- [Powering On and Off Network Switches](#)

Powering On Oracle Big Data Appliance

Oracle Big Data Appliance is powered on by either pressing the power button on the front of the servers, or by logging in to the ILOM interface and applying power to the system.

The power on sequence is as follows:

1. Turn on all 12 breakers on both PDUs.
Allow 1 to 2 minutes for the ILOM to start.
2. Power up the servers.

Powering On Servers Remotely Using ILOM

You can power on the servers remotely using the ILOM interface. You can access the ILOM using the web console, the command-line interface (CLI), IPMI, or SNMP. For example, to apply power to server `bda1node01` using IPMI, where `bda1node01-c` is the host name of the ILOM for the server to be powered on, run the following command as root from a server that has `ipmitool` installed:

```
ipmitool -H bda1node01-c -U root chassis power on
```

You are prompted for the password.

See Also: Oracle Integrated Lights Out Manager (ILOM) 3.0 documentation for information about using the ILOM to power on the servers:

<http://docs.oracle.com/cd/E19860-01/index.html>

Powering Off Oracle Big Data Appliance

The power off sequence for Oracle Big Data Appliance is as follows:

1. Power down the servers.
2. Turn off all 12 breakers on both PDUs.

Powering Off the Servers Use the Linux shutdown command to power off or restart the servers. Issue this command as `root` to shut down a server immediately:

```
# shutdown -hP now
```

The following command restarts a server immediately:

```
# shutdown -r now
```

See Also: Linux SHUTDOWN manual page for details

Powering Off Multiple Servers at the Same Time Use the `dcli` utility to run the `shutdown` command on multiple servers at the same time. Do not run the `dcli` utility from a server that will be shut down. Set up passwordless SSH for `root`, as described in ["Setting Up Passwordless SSH"](#) on page 7-1.

The following command shows the syntax of the command:

```
dcli -l root -g group_name shutdown -hP now
```

Where `group_name` is a file that contains a list of servers.

The following example shuts down all Oracle Big Data Appliance servers listed in `server_group`:

```
dcli -l root -g server_group shutdown -hP now
```

See Also: [Chapter 7, "Using the dcli Utility"](#)

Powering On and Off Network Switches

The network switches do not have power switches. They power off when power is removed by turning off a PDU or a breaker in the data center.

Emergency Power-Off Considerations

In an emergency, halt power to Oracle Big Data Appliance immediately. The following emergencies may require powering off Oracle Big Data Appliance:

- Natural disasters such as earthquake, flood, hurricane, tornado or cyclone
- Abnormal noise, smell or smoke coming from the system
- Threat to human safety

Emergency Power-Off Procedure

To perform an emergency power-off procedure for Oracle Big Data Appliance, turn off power at the circuit breaker or pull the emergency power-off switch in the computer room. After the emergency, contact Oracle Support Services to restore power to the system.

Emergency Power-Off Switch

Emergency power-off (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 connection to a site EPO switch or relay. Use of the EPO switch removes power from Oracle Big Data Appliance.

Cautions and Warnings

The following cautions and warnings apply to Oracle Big Data Appliance:

- Do not touch the parts of this product that use high-voltage power. Touching them might result in serious injury.
- Do not power off Oracle Big Data Appliance unless there is an emergency. In that case, follow the [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 hardware components.
- Keep the top, front, and back of the cabinets clear to allow proper airflow and prevent overheating of components.
- Use only the supplied hardware.

Adding Memory to a Server

Oracle Big Data Appliance ships from the factory with 48 GB of memory. Six of the 18 DIMM slots are populated with 8 GB DIMMs. You can use populate six empty slots with 8 GB DIMMs to bring the total memory to either 96 GB (12 x 8 GB) or 144 GB (18 x 8 GB). An upgrade to 144 GB may slightly reduce performance because of lower memory bandwidth; memory frequency drops from 1333 MHz to 800 MHz.

You can add memory to all nodes in the cluster or just to specific nodes that need more memory, such as the NameNode (server 01).

To add memory to a Sun Fire Server:

1. Power down the server.
2. Replace the plastic fillers with the six DIMMs as described in the *Sun Fire X4270 M2 Server Service Manual* at <http://docs.oracle.com/cd/E19245-01>.
3. Power on the server.

Maintaining the Physical Disks of Servers

Repair of the physical disks does not require shutting down Oracle Big Data Appliance Rack. However, individual servers may be taken outside of the cluster temporarily and require downtime.

See Also: ["Parts for Sun Fire Servers"](#) on page A-4 for the repair procedures

Verifying the Server Configuration

The 12 disk drives in each Sun Fire Server are controlled by an LSI MegaRAID SAS 92610-8i disk controller. Oracle recommends verifying the status of the RAID devices to avoid possible performance impact or an outage. The impact of validating the RAID devices is minimal. The impact of corrective actions vary depending on the specific issue uncovered and may range from simple reconfiguration to an outage.

Verifying Disk Controller Configuration

Issue this command to verify the disk controller configuration:

```
# MegaCli64 -AdpAllInfo -a0 | grep "Device Present" -A 8
```

The following is an example of the output from the command. There should be 12 virtual drives, no degraded or offline drives, and 14 physical devices. The 14 devices are the controller and the 12 disk drives.

```

                                Device Present
                                =====
Virtual Drives      : 12
  Degraded          : 0
  Offline           : 0
Physical Devices   : 14
  Disks             : 12
  Critical Disks    : 0
  Failed Disks      : 0

```

If the output is different, then investigate and correct the problem.

Verifying Virtual Drive Configuration

Issue this command to verify the virtual drive configuration:

```
# MegaCli64 -LDInfo -lAll -a0
```

Following is an example of the output for Virtual Drive 0. Ensure that State is Optimal.

```

Adapter 0 -- Virtual Drive Information:
Virtual Drive: 0 (Target Id: 0)
Name          :
RAID Level    : Primary-0, Secondary-0, RAID Level Qualifier-0
Size          : 1.817 TB
Parity Size   : 0
State         : Optimal
Strip Size    : 64 KB
Number Of Drives : 1
Span Depth    : 1
Default Cache Policy: WriteBack, ReadAheadNone, Cached, No Write Cache if Bad BBU
Current Cache Policy: WriteBack, ReadAheadNone, Cached, No Write Cache if Bad BBU
Access Policy  : Read/Write
Disk Cache Policy : Disk's Default
Encryption Type : None

```

Verifying Physical Drive Configuration

Use the following command to verify the physical drive configuration:

```
# MegaCli64 -PDList -a0 | grep Firmware
```

The following is an example of the output from the command. Twelve drives should be Online, Spun Up. If the output is different, then investigate and correct the problem.

```

Firmware state: Online, Spun Up
Device Firmware Level: 061A
Firmware state: Online, Spun Up
Device Firmware Level: 061A
Firmware state: Online, Spun Up
Device Firmware Level: 061A

```

Changing InfiniBand IP Addresses

You may need to change the InfiniBand network information on an existing Oracle Big Data Appliance. The change may support a media server with multiple InfiniBand cards, or keep InfiniBand traffic on a distinct InfiniBand network such as having production, test and QA environments in the same rack.

All InfiniBand addresses must be in the same subnet, with a minimum subnet mask of 255.255.240.0 (or /20). Choose a subnet mask wide enough to accommodate possible future expansion of the Oracle Big Data Appliance and InfiniBand network.

You cannot change the host names after running the Mammoth Utility.

To change the InfiniBand IP addresses:

1. Log in to an Oracle Big Data Appliance server as the `root` user.
2. Change to the `/etc/sysconfig/network-scripts` directory.
3. Copy the `ifcfg-bondib0` file, using a name that does not start with `ifcfg`:

```
cp ifcfg-bondib0 orig_ifcfg-bondib0
```
4. Edit `ifcfg-bondib0` to update the `IPADDR`, `NETMASK`, `NETWORK` and `BROADCAST` fields.

Following is an example of a `ifcfg-bondib0` file:

```
IPADDR=192.166.41.26
NETMASK=255.255.255.0
NETWORK=192.166.40.0
BROADCAST=192.166.47.255
DEVICE=bondib0
USERCTL=no
ONBOOT=yes
TYPE=bonding
BONDING_OPTS="mode=active-backup miimon=100 downdelay=5000 updelay=5000 num_
grat_arp=100"
HOTPLUG=no
IPV6INIT=no
BOOTPROTO=none
MTU=65520
```

5. Restart the server:

```
reboot
```
6. Verify the new InfiniBand IP address information:

```
iblinkinfo | grep reconfigured_host_name
```

See "[Checking the Health of the Network](#)" on page 8-25 for information about the `iblinkinfo` command.

7. Make a backup copy of the `/etc/hosts` file:

```
cp /etc/hosts /etc/orig_hosts
```
8. Edit the file to change the IP addresses.
9. Repeat this procedure on all Oracle Big Data Appliance servers.

Maintaining the InfiniBand Network

The InfiniBand network connects the servers through the bondib0 interface to the InfiniBand switches in the rack. This section describes how to perform maintenance on the InfiniBand switches.

This section contains the following topics:

- [Backing Up and Restoring Oracle ILOM Switch Settings](#)
- [Replacing a Failed InfiniBand Switch](#)
- [Verifying InfiniBand Network Operation](#)
- [Understanding the Network Subnet Manager Master](#)

Backing Up and Restoring Oracle ILOM Switch Settings

Oracle Integrated Lights Out Manager (ILOM) supports remote administration of the Oracle Big Data Appliance servers. This section explains how to back up and restore the Oracle ILOM configuration settings, which are set by the Oracle Big Data Appliance Mammoth Utility

This section contains the following topics:

- [Backing Up Oracle ILOM Configuration Settings](#)
- [Restoring Oracle ILOM Configuration Settings](#)

See Also: Oracle Integrated Lights Out Manager 3.0 documentation at

<http://docs.oracle.com/cd/E19860-01/>

Backing Up Oracle ILOM Configuration Settings

The following procedure describes how to back up the Oracle ILOM configuration settings.

1. Open your browser on any system on the same network as Oracle Big Data Appliance and enter the ILOM address of a server. This example uses the ILOM address of server 08:

```
http://bda1node08-c.example.com
```
2. Log in as the `ilom-admin` user. The default password is `welcome1`.
3. Select the **Maintenance** tab.
4. Select the **Backup/Restore** subtab.
5. Select the Backup operation and the Browser transfer method, as shown in [Figure 14-1](#).
6. Enter a pass phrase. The phrase is used to encrypt sensitive information, such as passwords, in the backup.
7. Click **Run** to start the backup. The results are downloaded to your local system in an XML file named `config_backup.xml`.
8. Save the file in a secure location.
9. Click the **Log Out** button.

Figure 14–1 Oracle ILOM Configuration Backup

Restoring Oracle ILOM Configuration Settings

The following procedure describes how to restore the settings on a switch:

1. Open your browser on any system on the same network as Oracle Big Data Appliance and navigate to an Oracle ILOM on an appliance. This example uses Oracle ILOM on server 8:
<http://bda1node08-c.us.example.com>
2. Log in as the `ilom-admin` user. The default password is `welcome1`.
3. Select the **Maintenance** tab.
4. Select the **Backup/Restore** tab.
5. Select the Restore operation and the Browser transfer method.
6. Click **Choose File** and select the `config_backup.xml` file saved previously in a backup operation.
7. Enter the pass phrase that was set during the backup.
8. Click **Run** to restore the configuration.

Replacing a Failed InfiniBand Switch

Take these steps to replace a Sun Network QDR InfiniBand Gateway Switch or a Sun Datacenter InfiniBand Switch 36.

See Also:

- "In-Rack InfiniBand Switch-to-Server Cable Connections" on page B-6 for information on cabling
- *Sun Network QDR InfiniBand Gateway Switch Installation Guide* at <http://docs.oracle.com/cd/E19671-01/821-1186-11/821-1186-11.pdf>
- *Sun Datacenter InfiniBand Switch 36 User's Guide* at <http://docs.oracle.com/cd/E19197-01/index.html>

To replace a failed InfiniBand switch:

1. Disconnect the cables from the switch. All InfiniBand cables have labels at both ends indicating their locations. If any cables do not have labels, then label them.
2. Power off both power supplies on the switch by removing the power plugs.
3. Remove the switch from the rack.
4. Install the new switch in the rack.
5. Restore the switch settings using the backup, as described in "Backing Up and Restoring Oracle ILOM Switch Settings" on page 14-7.
6. Connect to the switch as `ilom_admin` and open the Fabric Management shell:

```
-> show /SYS/Fabric_Mgmt
```

The prompt changes from `->` to `FabMan@hostname->`

7. Disable the Subnet Manager:


```
FabMan@bdalsw-02-> disablesm
```
8. Connect the cables to the new switch, being careful to connect each cable to the correct port.
9. Verify that there are no errors on any links in the fabric:


```
FabMan@bdalsw-02-> ibdiagnet -c 1000 -r
```
10. Enable the Subnet Manager:


```
FabMan@bdalsw-02-> enablesm
```

Note: If the replaced switch was the Sun Datacenter InfiniBand Switch 36 spine switch, then manually fail the Master Subnet Manager back to the switch by disabling the Subnet Managers on the other switches until the spine switch becomes the master, then re-enable the Subnet Manager on all the other switches.

Verifying InfiniBand Network Operation

If any component in the InfiniBand network has required maintenance, including replacing an InfiniBand Host Channel Adapter (HCA) on a server, an InfiniBand switch, or an InfiniBand cable, or if operation of the InfiniBand network is suspected to be substandard, then verify the InfiniBand network is operating properly. The following procedure describes how to verify network operation:

Note: Use this procedure used any time the InfiniBand network is performing below expectations.

1. Run the `ibdiagnet` command to verify InfiniBand network quality:

```
# ibdiagnet -c 1000
```

Investigate all errors reported by this command. It generates a small amount of network traffic and can run during a normal workload.

2. Report switch port error counters and port configuration information. The `LinkDowned`, `RcvSwRelayErrors`, `XmtDiscards`, and `XmtWait` errors are ignored by this command:

```
# ibqueryerrors.pl -rR -s LinkDowned,RcvSwRelayErrors,XmtDiscards,XmtWait
```

3. Check the status of the hardware:

```
# bdacheckhw
```

The following is an example of the output:

```
[SUCCESS: Correct system model : SUN FIRE X4270 M2 SERVER
[SUCCESS: Correct processor info : Intel(R) Xeon(R) CPU X5675 @ 3.07GHz
[SUCCESS: Correct number of types of CPU : 1
[SUCCESS: Correct number of CPU cores : 24
[SUCCESS: Sufficient GB of memory (>=48): 48
[SUCCESS: Correct GB of swap space : 24
[SUCCESS: Correct BIOS vendor : American Megatrends Inc.
[SUCCESS: Sufficient BIOS version (>=08080102): 08080102
[SUCCESS: Recent enough BIOS release date (>=05/23/2011) : 05/23/2011
[SUCCESS: Correct ILOM version : 3.0.16.10.a r68533
[SUCCESS: Correct number of fans : 6
[SUCCESS: Correct fan 0 status : ok
[SUCCESS: Correct fan 1 status : ok
[SUCCESS: Correct fan 2 status : ok
[SUCCESS: Correct fan 3 status : ok
[SUCCESS: Correct fan 4 status : ok
[SUCCESS: Correct fan 5 status : ok
[SUCCESS: Correct number of power supplies : 2
[Im[34mINFO: Detected Santa Clara Factory, skipping power supply checks
[SUCCESS: Correct disk controller model : LSI MegaRAID SAS 9261-8i
[SUCCESS: Correct disk controller firmware version : 12.12.0-0048
[SUCCESS: Correct disk controller PCI address : 13:00.0
[SUCCESS: Correct disk controller PCI info : 0104: 1000:0079
[SUCCESS: Correct disk controller PCIe slot width : x8
[SUCCESS: Correct disk controller battery type : iBBU08
[SUCCESS: Correct disk controller battery state : Operational
[SUCCESS: Correct number of disks : 12
[SUCCESS: Correct disk 0 model : SEAGATE ST32000SSSUN2.0
[SUCCESS: Sufficient disk 0 firmware (>=61A): 61A
[SUCCESS: Correct disk 1 model : SEAGATE ST32000SSSUN2.0
[SUCCESS: Sufficient disk 1 firmware (>=61A): 61A
.
.
.
[SUCCESS: Correct disk 10 status : Online, Spun Up No alert
[SUCCESS: Correct disk 11 status : Online, Spun Up No alert
[SUCCESS: Correct Host Channel Adapter model : Mellanox Technologies MT26428
ConnectX VPI PCIe 2.0
```

```
[SUCCESS: Correct Host Channel Adapter firmware version : 2.9.1000
[SUCCESS: Correct Host Channel Adapter PCI address : 0d:00.0
[SUCCESS: Correct Host Channel Adapter PCI info : 0c06: 15b3:673c
[SUCCESS: Correct Host Channel Adapter PCIe slot width : x8
[SUCCESS: Big Data Appliance hardware validation checks succeeded
```

4. Check the status of the software:

```
# bdachcksw

[SUCCESS: Correct OS disk sda partition info : 1 ext3 raid 2 ext3 raid 3
linux-swap 4 ext3 primary
[SUCCESS: Correct OS disk sdb partition info : 1 ext3 raid 2 ext3 raid 3
linux-swap 4 ext3 primary
[SUCCESS: Correct data disk sdc partition info : 1 ext3 primary
[SUCCESS: Correct data disk sdd partition info : 1 ext3 primary
[SUCCESS: Correct data disk sde partition info : 1 ext3 primary
[SUCCESS: Correct data disk sdf partition info : 1 ext3 primary
[SUCCESS: Correct data disk sdg partition info : 1 ext3 primary
[SUCCESS: Correct data disk sdh partition info : 1 ext3 primary
[SUCCESS: Correct data disk sdi partition info : 1 ext3 primary
[SUCCESS: Correct data disk sdj partition info : 1 ext3 primary
[SUCCESS: Correct data disk sdk partition info : 1 ext3 primary
[SUCCESS: Correct data disk sdl partition info : 1 ext3 primary
[SUCCESS: Correct software RAID info : /dev/md2 level=raid1 num-devices=2
/dev/md0 level=raid1 num-devices=2
[SUCCESS: Correct mounted partitions : /dev/md0 /boot ext3 /dev/md2 / ext3
/dev/sda4 /u01 ext4 /dev/sdb4 /u02 ext4 /dev/sdc1 /u03 ext4 /dev/sdd1 /u04 ext4
/dev/sde1 /u05 ext4 /dev/sdf1 /u06 ext4 /dev/sdg1 /u07 ext4 /dev/sdh1 /u08 ext4
/dev/sdi1 /u09 ext4 /dev/sdj1 /u10 ext4 /dev/sdk1 /u11 ext4 /dev/sdl1 /u12 ext4
[SUCCESS: Correct swap partitions : /dev/sdb3 partition /dev/sda3 partition
[SUCCESS: Correct Linux kernel version : Linux 2.6.32-200.21.1.el5uek
[SUCCESS: Correct Java Virtual Machine version : HotSpot(TM) 64-Bit Server
1.6.0_29
[SUCCESS: Correct puppet version : 2.6.11
[SUCCESS: Correct MySQL version : 5.5.17
[SUCCESS: All required programs are accessible in $PATH
[SUCCESS: All required RPMs are installed and valid
[SUCCESS: Big Data Appliance software validation checks succeeded
```

Understanding the Network Subnet Manager Master

The Subnet Manager manages all operational characteristics of the InfiniBand network, such as the following:

- Discover the network topology
- Assign a local identifier to all ports connected to the network
- Calculate and program switch forwarding tables
- Monitor changes in the fabric

The InfiniBand network can have multiple Subnet Managers, but only one Subnet Manager is active at a time. The active Subnet Manager is the Master Subnet Manager. The other Subnet Managers are the Standby Subnet Managers. If a Master Subnet Manager is shut down or fails, then a Standby Subnet Manager automatically becomes the Master Subnet Manager.

Each Subnet Manager has a configurable priority. When multiple Subnet Managers are on the InfiniBand network, the Subnet Manager with the highest priority becomes the Master Subnet Manager. On Oracle Big Data Appliance, the Subnet Managers on leaf

switches are configured as priority 5, and the Subnet Managers on spine switches are configured as priority 8.

The following guidelines determine where Subnet Managers run on Oracle Big Data Appliance:

- Only run Subnet Managers on the switches in Oracle Big Data Appliance. Running Subnet Manager on any other device is not supported.
- When the InfiniBand network consists of one, two or three racks cabled together, all switches must run Subnet Manager. The Master Subnet Manager runs on a spine switch.
- When the InfiniBand network consists of four or more racks cabled together, then only spine switches run Subnet Manager. The leaf switches must disable Subnet Manager.

See Also:

- Sun Network QDR InfiniBand Gateway Switch library at <http://docs.oracle.com/cd/E19671-01/>
- Sun Datacenter InfiniBand Switch 36 library at <http://download.oracle.com/docs/cd/E19671-01/index.html>

Changing the NTP Servers

The configuration information for Network Time Protocol (NTP) servers can be changed after the initial setup. The following procedure describes how to change the NTP configuration information for InfiniBand switches, Cisco switches, and Sun Fire Servers. Oracle recommends changing each server individually.

To update the Oracle Big Data Appliance servers:

1. Stop NTP services on the server.
2. Update the `/etc/ntp.conf` file with the IP address of the new NTP server.
3. Repeat these steps for each server.

To update the InfiniBand switches:

1. Log in to the switch as the `ilom-admin` user.
2. Follow the instructions in "[Setting the Time Zone on an InfiniBand Switch](#)" on page 8-10.

To update the Cisco Ethernet switch:

1. Use telnet to connect to the Cisco Ethernet switch.
2. Erase the current setting:

```
# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
(config)# no ntp server current_IPaddress
```
3. Enter the new IP address:

```
# configure terminal
Enter configuration commands,one per line.End with CNTL/Z.
(config)# ntp server new_IPaddress
```
4. Save the current configuration:

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

5. Exit from the session:

```
# exit
```

See Also: ["Configuring the Cisco Ethernet Switch"](#) on page 8-3

Restart Oracle Big Data Appliance after changing the servers and switches.

Replacement Units

This appendix lists the replacement units for Oracle Big Data Appliance Rack. All replacement units are FRUs (field replaceable units), which are installed by trained Oracle field engineers.

This appendix contains the following topics:

- [Understanding Repair Categories](#)
- [Compliance with the European Restriction of the Use of Certain Hazardous Substances in Electrical and Electronic Equipment Directive](#)
- [Disk Controller Batteries](#)
- [LED Status Descriptions](#)
- [Parts for Sun Fire Servers](#)
- [Parts for InfiniBand Switches](#)
- [Parts for Cisco Ethernet Switch](#)
- [Parts for KVM Switch and KMM Tray Components](#)
- [Parts for Power Distribution Units](#)
- [Cables and Cable Accessories for Oracle Big Data Appliance Rack](#)

See Also: [Chapter 14, "Maintaining Oracle Big Data Appliance"](#) for information about the pre-installation and post-installation steps

Understanding Repair Categories

Repair categories describe the potential impact of a repair action. [Table A-1](#) describes the categories, and actions.

Table A-1 Oracle Big Data Appliance Rack Repair Categories

Category	Description	Scenario Examples
Hot Swap (HS)	Repair part is hot swappable and can be replaced without shutting down the host system. Procedures may be needed before and after replacement to protect the data.	<ul style="list-style-type: none"> ■ Disks ■ Fans ■ Power supplies
Infrastructure Repair (IR)	Repair of a connectivity component within Oracle Big Data Appliance Rack. No downtime of the rack is required, however, individual components may require downtime.	<ul style="list-style-type: none"> ■ External cables ■ InfiniBand switch ■ Ethernet switch ■ KVM switch/KMM tray

Table A-1 (Cont.) Oracle Big Data Appliance Rack Repair Categories

Category	Description	Scenario Examples
BDA Server Offline (BDA-O)	<p>Repair of the part requires shutting down one Big Data Appliance server. No downtime of the rack is required, however, individual servers may require downtime and temporary removal from the cluster. Rebalancing the work load and data duplication may temporarily impact performance.</p> <p>If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 14-2 for additional information.</p>	<ul style="list-style-type: none"> ■ System boards ■ PCIe cards ■ Memory ■ Processors

Compliance with the European Restriction of the Use of Certain Hazardous Substances in Electrical and Electronic Equipment Directive

Parts that have the RoHS:Y or RoHS:YL designation are compliant with the European Restriction of the Use of Certain Hazardous Substances in Electrical and Electronic Equipment Directive, as follows:

- RoHS:Y compliance is based on present information and belief, this equipment/part number complies with the European Union RoHS Directive, 2002/95/EC, and amending and guidance documents, and does not contain any restricted substances above the homogeneous materials' maximum concentration values, except for selected exemptions allowed by law (excluding the exemption for lead in solders for servers, storage and storage array systems, network infrastructure equipment for switching, signaling, transmission as well as network management for telecommunications).
- RoHS:YL compliance is based on present information and belief, this equipment/part number complies with the European Union RoHS Directive, 2002/95/EC, and amending and guidance documents, and does not contain any restricted substances above the homogeneous materials' maximum concentration values, except for selected exemptions allowed by law (including the exemption for lead in solders for servers, storage and storage array systems, network infrastructure equipment for switching, signaling, transmission as well as network management for telecommunications).
- If a part description contained in a parts list does not contain RoHS:Y or RoHS:YL in its description, this indicates that the part is not RoHS compliant or the information necessary to confirm its compliance has yet to be provided by the product team.

Disk Controller Batteries

The disk controllers in Sun Fire Servers have battery-backed write cache to accelerate write performance. If the battery charge capacity degrades such 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. This results in reduced write performance but no data loss. Sun Fire Servers generate an alert when battery charge capacity is insufficient, the temperature is high, and when the battery should be replaced.

Battery charge capacity degrades over time, and its life expectancy is inversely proportional to the operating temperature. The worst case life expectancy of the battery in Oracle Big Data Appliance Rack is as follows:

Inlet Ambient Temperature	Battery Lifetime
< 25 degrees Celsius (77 degrees Fahrenheit)	3 years
< 32 degrees Celsius (89.6 degrees Fahrenheit)	2 years

Replacing Batteries in Disk Controllers

Oracle replaces failed batteries at no extra charge if the battery charge capacity in the disk controllers falls below the minimum threshold. As part of Premier Support for Systems, Oracle attempts to proactively replace the batteries in Oracle Big Data Appliance Rack before the end of the estimated lifetime, on a best effort basis.

LED Status Descriptions

This section describes the LEDs in Oracle Big Data Appliance Rack and contains these topics:

- [Sun Fire Server LEDs](#)
- [Sun Datacenter InfiniBand Switch 36 LEDs](#)

Sun Fire Server LEDs

[Table A-2](#) describes the LED color codes on Sun Fire X4270 M2 Servers.

Table A-2 LED Status Descriptions

Component	LED Status
Fan module	<ul style="list-style-type: none"> ■ Fan Status LED is Green: The system is powered on and the fan module is functioning correctly. ■ Fan Status LED is Amber: The fan module is faulty. The front and rear panel Service Action Required LEDs are also lit if the system detects a fan module fault.
Power supply	<ul style="list-style-type: none"> ■ OK to Remove LED is green: The power supply can be removed safely during a hot-swap procedure. ■ Service Action Required LED is amber: The power supply is faulty. The front and rear panel Service Action Required LEDs are also lit if the system detects a power supply fault. ■ AC Present LED is green: Power supply can be removed during a hot-swap procedure.
Servers	<ul style="list-style-type: none"> ■ OK to Remove LED is blue: A storage drive can safely be removed during a hot-swap procedure. ■ Service Action Required LED is amber: The storage drive is faulty. The front and rear panel Service Action Required LEDs are also lit if the system detects a storage drive fault. ■ OK/Activity LED is green: Data is being read from or written to the storage drive.

Sun Datacenter InfiniBand Switch 36 LEDs

[Table A-3](#) describes the color codes of the LEDs on Sun Datacenter InfiniBand Switch 36 switches.

Table A-3 Sun Datacenter InfiniBand Switch 36 LED Status Descriptions

Component	LED Status
Chassis	<ul style="list-style-type: none"> ■ Locator LED is white: It flashes when identifying itself. It is on when there is no function, and off when disabled. ■ Attention LED is amber: There is a fault condition. It flashes when there is no function. ■ OK LED is green: Switch is functioning correctly. It flashes when there is no function.
Link status	<ul style="list-style-type: none"> ■ Link LED is green: It is on when link is established. It is off when link is down, and it flashes when there are symbol errors.
Network management ports	<p>Link speed LED:</p> <ul style="list-style-type: none"> ■ Green indicates 1000BASE-T. ■ Amber indicates 100BASE-T. ■ Off indicates no link. ■ Flashing indicates no function. <p>Activity LED:</p> <ul style="list-style-type: none"> ■ Flashing indicates packet activity. ■ On indicates no function. ■ Off indicates no activity.
Power supply	<ul style="list-style-type: none"> ■ OK LED is green: Indicates 12 VDC is supplied. Flashing indicates no function. ■ Attention LED is amber: There is a fault, and 12 VDC shut down. Flashing indicates no function. ■ AC LED is green: AC power is present and good. Flashing indicates no function.

Parts for Sun Fire Servers

Table A-4 lists the replaceable parts for Sun Fire Servers, which are covered under the Premier Support for Systems warranty.

Note: Oracle Big Data Appliance ships with two spare drives. If a drive is faulty, then the customer can replace the drive before a Customer Service Representative arrives if circumstances warrant immediate action.

Table A-4 Replacement Parts for a Sun Fire X4270 M2 Server

Part Number	Description	Repair Category	Repair Procedure
150-3993	Battery, 3V, 390MAH, LITH, COIN	BDA-O	See section 4.6, "Servicing the Server Battery" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_71528
300-2235	Power supply, A249, 1200 W, 12 V, 2U, RoHS:Y	HS	See section 3.4, "Servicing Power Supplies" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/hotswap.html#50438178_11706
350-1502	Universal rack mount kit	BDA-O	<ol style="list-style-type: none"> 1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 14-2 for additional information. 2. See Chapter 2, "Installing the Server Into a Rack With Slide-Rails" in <i>Sun Fire X4170 M2, and X4270 M2 Servers Installation Manual</i> at http://docs.oracle.com/cd/E19245-01/E27205/z4000c671007752.html
371-4743	4 GB USB 2.0 flash drive, RoHS:Y	BDA-O	Not applicable
7012433	3.06 GHz processor	BDA-O	See section 4.7, "Servicing the CPU" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_23425
371-4966	LV DIMM, 8 GB, DDR3, 1RX4, 1333 MHz	BDA-O	See section 4.7, "Servicing Memory Modules (DIMMs)" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_71526
371-4982	6 Gigabit SAS RAID PCI battery module, RoHS:Y (LION), BBU-08	BDA-O	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 14-2 for additional information. 2. See <i>Sun Storage 6 Gb SAS PCIe RAID HBA, Internal Installation Guide For HBA Models SGX-SAS6-R-INT-Z and SG-SAS6-R-INT-Z (internal only)</i> at http://docs.oracle.com/cd/E19337-01/index.html
375-3696	InfiniBand (CX2) dual port CX2 4xQDR PCI-e	BDA-O	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 14-2 for additional information. 2. See section 4.5, "Servicing PCIe Cards" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_86021

Table A-4 (Cont.) Replacement Parts for a Sun Fire X4270 M2 Server

Part Number	Description	Repair Category	Repair Procedure
375-3701	6 Gb SAS RAID disk PCIe HBA, RoHS:Y	BDA-O	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 14-2 for additional information. 2. See section 4.5, "Servicing PCIe Cards" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_86021 <p>Note: Do not run any controller commands in the service manual when replacing the disk.</p>
530-3927	Cable, PDB, MB, 1U+2U, ribbon	BDA-O	<p>See section 5.10, "Servicing Cables" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_58217</p>
530-4228	Cable, mini-SAS, 36 POS LONG, LY, RoHS:Y	BDA-O	<p>See section 5.10, "Servicing Cables" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_78806</p>
541-2075	Bus_bar, power, 1U+2U, RoHS:Y	BDA-O	<p>See section 5.6, "Servicing the Power Distribution Board" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_85936</p>
541-2884	Assembly, riser, X8, X8, 2U, PCI-E, RoHS:Y	BDA-O	<p>See section 4.4, "Servicing PCIe Risers" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_53252</p>
541-3310	Xeon Heatsink with grease, RoHS:Y, 2U	BDA-O	<p>See section 4.7, "Servicing the CPU" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_23425</p>
541-3356	<p>x8/x8 switched PCIe riser assembly, RoHS:Y with the following:</p> <ul style="list-style-type: none"> ■ x8/x8 switch PCIe riser ■ bracket 	BDA-O	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 14-2 for additional information. 2. See section 4.4, "Servicing PCIe Risers" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_53252

Table A-4 (Cont.) Replacement Parts for a Sun Fire X4270 M2 Server

Part Number	Description	Repair Category	Repair Procedure
541-4081	System board assembly, RoHS:Y	BDA-O	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 14-2 for additional information. 2. See section 4.8, "Servicing the Motherboard" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_98526
541-4124	Fan deck, Hyde, RoHS:Y	BDA-O	See section 5.3, "Servicing the Fan Board" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_71556
541-4222	Assembly, fan, CR, 2U, RoHS:Y	HS	See section 3.3, "Servicing the Fan Modules" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/hotswap.html#50438178_28930
541-4223	Power distribution board, 2U, RoHS:Y	BDA-O	See section 5.6, "Servicing the Power Distribution Board" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_71558
541-4304	Indicator, right, Hyde 12/24, roHS:y	BDA-O	See section 5.5, "Servicing the Front Panel LED Modules" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_71563
541-4305	Assembly, disk, SAS2, XP_DC, Hyde, RoHS:Y	BDA-O	See section 5.2, "Servicing the SAS Expander Board" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_28246
541-4306	Assembly, connector board bracket, Hyde	BDA-O	See section 5.7, "Servicing the Connector Board" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_pgfid-1296581
541-4363	12-slot disk backplane, RoHS:Y	BDA-O	See section 5.4, "Servicing the Disk Backplane" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_71557

Table A-4 (Cont.) Replacement Parts for a Sun Fire X4270 M2 Server

Part Number	Description	Repair Category	Repair Procedure
541-4371	Indicator, left, Hyde 12/24, RoHS:Y	BDA-O	See section 5.5, "Servicing the Front Panel LED Modules" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_71563
7010036	FRU DR, 3TB, 3.5" SAS/7200, CORAL	--	See section 3.2, "Servicing Storage Drives and Boot Drives" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/hotswap.html#50438178_61628 Note: Do not run any controller commands in the service manual when replacing the disk.
560-2952	Cable kit, Hyde, RoHS:Y with the following: <ul style="list-style-type: none"> ■ Disk backplane power cable ■ 24-slot DBP power cable ■ 12/24 slot disk backplane signal cable ■ Fan signal cable ■ Fan power cable 	BDA-O	Refer to the entries for the individual items in the kit

Parts for InfiniBand Switches

Table A-5 lists the replaceable parts for the InfiniBand switches used with Oracle Big Data Appliance Rack, which are covered under the Premier Support for Systems warranty.

See "[Cables and Cable Accessories for Oracle Big Data Appliance Rack](#)" on page A-10 for information about InfiniBand cables.

See Also:

- *Sun Network QDR InfiniBand Gateway Switch Service Manual* at
<http://docs.oracle.com/cd/E19671-01/821-1188-11/821-1188-11.pdf>
- *Sun Datacenter InfiniBand Switch 36 User's Guide* at
<http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html>

Table A-5 Replacement Parts for InfiniBand Switches

Part Number	Description	Repair Category	Repair
300-2143	760 watt power supply, RoHS:Y	HS	See <i>Sun Network QDR InfiniBand Gateway Switch Service Manual</i> or <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i>
350-1312	Fan module for Sun Datacenter InfiniBand Switch 36, RoHS:Y	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i>
371-2210	CR2032 3 v battery, RoHS:Y	IR	See <i>Sun Network QDR InfiniBand Gateway Switch Service Manual</i> or <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i>
541-3495	Sun Datacenter InfiniBand Switch 36, RoHS:YL	IR	<ol style="list-style-type: none"> 1. See "Replacing a Failed InfiniBand Switch" on page 8 2. See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i>
7014378	Assembly, system, NM2-Gateway	IR	See <i>Sun Network QDR InfiniBand Gateway Switch Installation Guide</i> at http://docs.oracle.com/cd/E19671-01/821-1186-11/821-1186-11.pdf
594-6603	10GBPS QSFP short wavelength T	IR	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i>

Parts for Cisco Ethernet Switch

Table A-6 lists replaceable parts for the Cisco Ethernet switch, which are covered under the Premier Support for Systems warranty.

See Also: *Cisco Catalyst 4948E Switch Installation Guide* for repair instructions at

http://www.cisco.com/en/US/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html

Table A-6 Replacement Parts for Cisco Catalyst Ethernet Switch

Part Number	Description	Repair Category	Repair
310-0307	Cooling fan	HS	See Chapter 4 of <i>Cisco Catalyst 4948E Switch Installation Guide</i>
371-4784	Cisco Catalyst 4948 switch, RoHS:Y	IR	See Chapter 3 of <i>Cisco Catalyst 4948E Switch Installation Guide</i>
371-4785	Power supply	HS	See Chapter 4 of <i>Cisco Catalyst 4948E Switch Installation Guide</i>

Parts for KVM Switch and KMM Tray Components

Table A-7 lists the replacement parts for the KVM switch and KMM tray components, which are covered under the Premier Support for Systems warranty.

See Also: *MergePoint Unity Switch Installer/User Guide* for repair instructions at

<http://pcs.mktg.avocent.com/@@content/manual/590883501c.pdf>

Table A-7 Replacement Parts for KVM Switch and KMM Tray Components

Part Number	Description	Repair Category	Repair
371-4778	Assembly, KMM, with Japanese keyboard and mouse module	IR	See Chapter 2 of <i>MergePoint Unity Switch Installer/User Guide</i>
371-4779	Avocent MPU4032DAC-001 32-port KVM switch, RoHS:Y	IR	See Chapter 2 of <i>MergePoint Unity Switch Installer/User Guide</i>
371-4780	Avocent KMM drawer with United States keyboard, RoHS:Y	IR	See Chapter 2 of <i>MergePoint Unity Switch Installer/User Guide</i>
371-4781	Avocent (DSRIQ-USB) DB 15 M to RJ45/USB KVM adapter, RoHS:Y	IR	See Chapter 2 of <i>MergePoint Unity Switch Installer/User Guide</i>

Parts for Power Distribution Units

[Table A-8](#) lists the replacement parts for the power distribution units (PDUs), which are covered under the Premier Support for Systems warranty.

Given the number of cables and connections, Oracle recommends the Oracle Big Data Appliance Rack rack be shut down when replacing power distribution unit parts.

See Also: *Sun Rack II Power Distribution Units User's Guide* for repair instructions at

<http://docs.oracle.com/cd/E19844-01/E23443-01/index.html>

Table A-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
371-3995	Three-phase PDU: 15kVA, IEC309 60A, 4 pin, 250VAC, 3ph IP67, North America, RoHS:Y	IR	See Chapter 3 of <i>Sun Rack II Power Distribution Units User's Guide</i>
371-4000	Three-phase PDU: 15kV,5-Pin, IEC309 32A, 5 pin 230/400V, 3ph IP44, international, RoHS:Y	IR	See Chapter 3 of <i>Sun Rack II Power Distribution Units User's Guide</i>
371-4199	Single-phase PDU: 15kVA, with three 30A/250V 2-pole/3-wire NEMA L6-30P plugs, North America, RoHS:Y	IR	See Chapter 3 of <i>Sun Rack II Power Distribution Units User's Guide</i>
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 Chapter 3 of <i>Sun Rack II Power Distribution Units User's Guide</i>

Cables and Cable Accessories for Oracle Big Data Appliance Rack

[Table A-9](#) lists the replacement cables and cable accessories, which are covered under the Premier Support for Systems warranty.

See Also:

- *Sun Network QDR InfiniBand Gateway Switch Service Manual* at <http://docs.oracle.com/cd/E19671-01/821-1188-11/821-1188-11.pdf>
- *Sun Datacenter InfiniBand Switch 36 User's Guide* at <http://docs.oracle.com/cd/E19197-01/835-0784-05/gentextid-226.html>
- See *Cisco Catalyst 4948E Switch Installation Guide* at http://www.cisco.com/en/US/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html

Table A-9 Replacement Cables for Oracle Big Data Appliance Rack

Part Number	Description	Repair Category	Repair
350-1287	Slim rail, CMA, 1U-2U for Sun Fire Server	BDA-O	<ol style="list-style-type: none"> 1. If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 14-2 for additional information. 2. See "Installing the Cable Management Arm" in <i>Sun Fire X4170, X4270, and X4275 Servers Installation Manual</i> at http://docs.oracle.com/cd/E19477-01/820-5827-13/rack_install.html#50614309_99582
350-1519	Serial cable kit, RoHS:Y with the following: <ul style="list-style-type: none"> ■ USB to DB 9-pin M serial cable ■ DB 9-pin F to DB 9-pin F null modem cable 	HS	See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i>
530-4403	4x QSFP copper QDR InfiniBand cable, 2 m, RoHS:Y	HS	<ol style="list-style-type: none"> 1. See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> 2. See "Maintaining the InfiniBand Network" on page 14-7 <p>Note: There is a loss of redundancy when replacing cables.</p>
530-4404	4x QSFP copper QDR InfiniBand cable, 3 m, RoHS:Y	HS	<ol style="list-style-type: none"> 1. See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> 2. See "Maintaining the InfiniBand Network" on page 14-7
530-4415	4x QSFP copper QDR InfiniBand cable, 5 m, RoHS:Y	HS	<ol style="list-style-type: none"> 1. See <i>Sun Datacenter InfiniBand Switch 36 User's Guide</i> 2. See "Maintaining the InfiniBand Network" on page 14-7
530-4432	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, blue, RoHS:Y	HS	See <i>Cisco Catalyst 4948E Switch Installation Guide</i>
530-4433	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 10 foot, red, RoHS:Y	HS	See <i>Cisco Catalyst 4948E Switch Installation Guide</i>

Table A–9 (Cont.) Replacement Cables for Oracle Big Data Appliance Rack

Part Number	Description	Repair Category	Repair
530-4434	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, red, RoHS:Y	HS	See <i>Cisco Catalyst 4948E Switch Installation Guide</i>
530-4435	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 10 foot, black, RoHS:Y	HS	See <i>Cisco Catalyst 4948E Switch Installation Guide</i>
530-4436	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, black, RoHS:Y	HS	See <i>Cisco Catalyst 4948E Switch Installation Guide</i>
530-4437	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, orange, RoHS:Y	HS	See <i>Cisco Catalyst 4948E Switch Installation Guide</i>
530-4438	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 10 foot, black, RoHS:Y	HS	See <i>Cisco Catalyst 4948E Switch Installation Guide</i>

Cabling Tables

The tables in this appendix show the cable connections for Oracle Big Data Appliance. This appendix contains the following topics:

- [Cable Color Coding](#)
- [Oracle Big Data Appliance Rack Layout](#)
- [Keyboard, Video and Mouse Network Cable Connections](#)
- [Oracle Integrated Lights Out Manager Cable Connections](#)
- [Single-Phase Power Distribution Unit Cable Connections](#)
- [Three-Phase Power Distribution Unit Cable Connections](#)
- [In-Rack InfiniBand Switch-to-Switch Cable Connections](#)
- [In-Rack InfiniBand Switch-to-Server Cable Connections](#)

Cable Color Coding

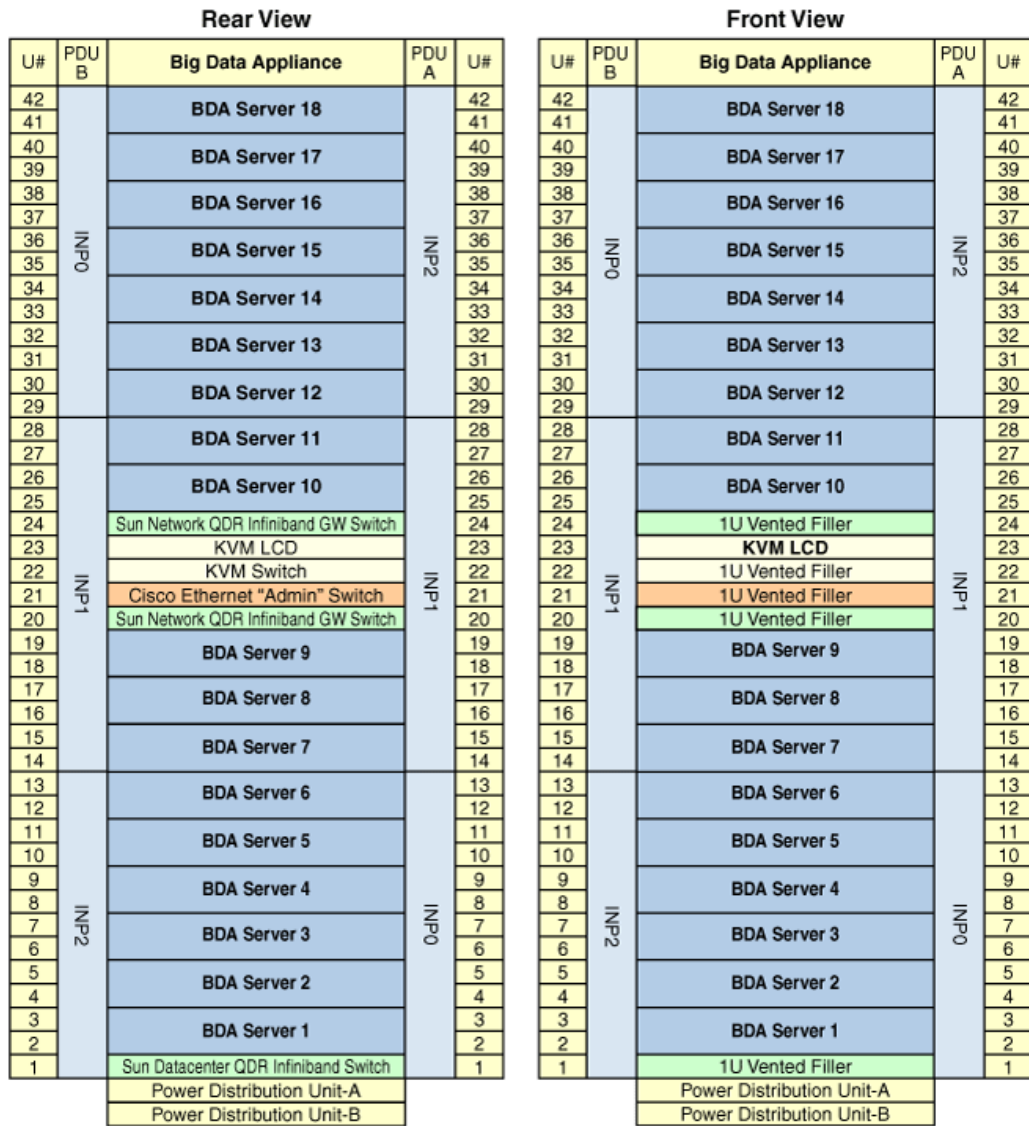
The cables used in Oracle Big Data Appliance are color-coded as follows:

- Black: InfiniBand cables or AC power jumper cables
- Red: Integrated Lights Out Management (ILOM) Ethernet management cables
- Blue: Gigabit Ethernet management (eth0) cables
- Orange: KVM switch to dongle cables

Oracle Big Data Appliance Rack Layout

[Figure B-1](#) shows the front and rear layout of Oracle Big Data Appliance Rack.

Figure B-1 Rack Layout of Oracle Big Data Appliance Rack



Keyboard, Video and Mouse Network Cable Connections

Table B-1 shows the keyboard, video, and mouse (KVM) network cabling. The KVM port on the servers is labeled SER MGT and connects to the KVM switch located in rack unit 22.

Table B-1 KVM Cabling for Oracle Big Data Appliance Rack

From Rack Unit ¹	Port	To KVM Port	Cable Length	Cable Color	Component
U41	Net-0	1	10'	Black	Server Cell 18
U39	Net-0	3	10'	Black	Server Cell 17
U37	Net-0	5	10'	Black	Server Cell 16
U35	Net-0	7	10'	Black	Server Cell 15

Table B-1 (Cont.) KVM Cabling for Oracle Big Data Appliance Rack

From Rack Unit ¹	Port	To KVM Port	Cable Length	Cable Color	Component
U33	Net-0	9	10'	Black	Server Cell 14
U31	Net-0	11	10'	Black	Server Cell 13
U29	Net-0	13	10'	Black	Server Cell 12
U27	Net-0	17	10'	Black	Server Cell 11
U25	Net-0	21	10'	Black	Server Cell 10
U24	Net-0	45	10'	Black	NM2-1B Switch
U20	Net-0	46	10'	Black	NM2-1B Switch
U18	Net-0	25	10'	Black	Server Cell 9
U16	Net-0	29	10'	Black	Server Cell 8
U14	Net-0	31	10'	Black	Server Cell 7
U12	Net-0	33	10'	Black	Server Cell 6
U10	Net-0	35	10'	Black	Server Cell 5
U8	Net-0	37	10'	Black	Server Cell 4
U6	Net-0	39	10'	Black	Server Cell 3
U4	Net-0	41	10'	Black	Server Cell 2
U2	Net-0	43	10'	Black	Server Cell 1
U1	Net-0	47	10'	Black	NM2-1B Switch
PDU-A	Net Mgt	15	1M	White	PDU-A
PDU-B	Net Mgt	19	1M	White	PDU-B
N/A	N/A	48	10'	Blue	Service

¹ Un is the unit location in the rack, where n is the number.

Oracle Integrated Lights Out Manager Cable Connections

Table B-2 shows the cable connections from the servers to the Oracle Integrated Lights Out Manager (ILOM) switch. The ILOM port on the servers is labeled NET MGT and connects to the Cisco Ethernet switch located in rack unit 21. The cables are red and 10 feet long.

Table B-2 ILOM Cabling

From Rack Unit ¹	Ethernet Port	Component
U41	2	Server Cell 18
U39	4	Server Cell 17
U37	6	Server Cell 16
U35	8	Server Cell 15
U33	10	Server Cell 14
U31	12	Server Cell 13
U29	14	Server Cell 12
U27	18	Server Cell 11

Table B–2 (Cont.) ILOM Cabling

From Rack Unit ¹	Ethernet Port	Component
U25	22	Server Cell 10
U18	26	Server Cell 9
U16	30	Server Cell 8
U14	32	Server Cell 7
U12	34	Server Cell 6
U10	36	Server Cell 5
U8	38	Server Cell 4
U6	40	Server Cell 3
U4	42	Server Cell 2
U2	44	Server Cell 1

¹ U*n* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cable Connections

Table B–3 shows the connections for single-phase cabling from each power distribution unit (PDU) to the power supplies in the rack. 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 B–3 Single-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length	Component
U41	G5-6	G0-0	2 meters	Sun Fire Server
U39	G5-3	G0-3	2 meters	Sun Fire Server
U37	G5-0	G0-6	2 meters	Sun Fire Server
U35	G4-6	G1-0	2 meters	Sun Fire Server
U33	G4-4	G1-2	2 meters	Sun Fire Server
U31	G4-2	G1-4	2 meters	Sun Fire Server
U29	G3-6	G2-0	2 meters	Sun Fire Server
U27	G3-5	G2-1	2 meters	Sun Fire Server
U25	G3-3	G2-3	2 meters	Sun Fire Server
U24	G3-1	G2-5	2 meters	NM2
U23	N/A	G3-0	included	KVM
U22	G2-5	G3-1	1 meter	KVM tray
U21	G3-0	G2-6	2 meters	Cisco
U20	G2-4	G3-2	2 meters	NM2
U18	G2-2	G3-4	2 meters	Sun Fire Server
U16	G1-6	G4-0	2 meters	Sun Fire Server
U14	G2-0	G3-6	2 meters	Sun Fire Server
U12	G1-4	G4-2	2 meters	Sun Fire Server

Table B-3 (Cont.) Single-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length	Component
U10	G1-2	G4-4	2 meters	Sun Fire Server
U8	G1-0	G4-6	2 meters	Sun Fire Server
U6	G0-6	G5-0	2 meters	Sun Fire Server
U4	G0-4	G5-2	2 meters	Sun Fire Server
U2	G0-2	G5-4	2 meters	Sun Fire Server
U1	G0-0	G5-6	2 meters	NM2

¹ U n is the unit location in the rack, where n is the number.

Three-Phase Power Distribution Unit Cable Connections

Table B-3 describes three-phase cabling from each power distribution unit (PDU) to the power supplies in the servers. The cables are terminated to PDU-A on the left, are routed to the right to enter CMA, and are bundled in groups of four.

Table B-4 Three-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length	Component
U41	G5-6	G2-0	2 meters	Sun Fire Server
U39	G5-3	G2-3	2 meters	Sun Fire Server
U37	G5-0	G2-6	2 meters	Sun Fire Server
U35	G4-6	G1-0	2 meters	Sun Fire Server
U33	G4-4	G1-2	2 meters	Sun Fire Server
U31	G4-2	G1-4	2 meters	Sun Fire Server
U29	G3-6	G0-0	2 meters	Sun Fire Server
U27	G3-5	G0-1	2 meters	Sun Fire Server
U25	G3-3	G0-3	2 meters	Sun Fire Server
U24	G3-1	G0-5	2 meters	NM2
U23	N/A	G0-3	included	KVM
U22	G2-5	G5-1	1 meter	KVM tray
U21	G3-0	G0-6	2 meters	Cisco
U20	G2-4	G5-2	2 meters	NM2
U18	G2-2	G5-4	2 meters	Sun Fire Server
U16	G4-2	G5-2	2 meters	Sun Fire Server
U14	G2-0	G5-6	2 meters	Sun Fire Server
U12	G1-4	G4-2	2 meters	Sun Fire Server
U10	G1-2	G4-4	2 meters	Sun Fire Server
U8	G1-0	G4-6	2 meters	Sun Fire Server
U6	G0-6	G3-0	2 meters	Sun Fire Server
U4	G0-4	G3-2	2 meters	Sun Fire Server
U2	G0-2	G3-4	2 meters	Sun Fire Server

Table B-4 (Cont.) Three-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length	Component
U1	G0-0	G3-6	2 meters	NM2

¹ U*n* is the unit location in the rack, where *n* is the number.

In-Rack InfiniBand Switch-to-Switch Cable Connections


Table B-5 lists the location, ports, and cables for the InfiniBand connections from switch to switch. Figure B-2 identifies the locations of the ports on a Sun Network QDR InfiniBand Gateway Switch.

Table B-5 InfiniBand Switch-to-Switch Cabling

From InfiniBand Switch Rack Unit ¹	Port	To InfiniBand Switch Rack Unit ¹	Port	Cable Length	Description
U20	9B	U24	9A	2M	Leaf to leaf
U20	10B	U24	10A	2M	Leaf to leaf
U20	11B	U24	11A	2M	Leaf to leaf
U20	8A	U24	8A	2M	Leaf to leaf
U20	9A	U24	9B	2M	Leaf to leaf
U20	10A	U24	10B	2M	Leaf to leaf
U20	11A	U24	11B	2M	Leaf to leaf
U1	0B	U20	8B	3M	Spine to leaf
U1	1B	U24	8B	3M	Spine to leaf

¹ U*n* is the unit location in the rack, where *n* is the number.

Figure B-2 Sun Network QDR InfiniBand Gateway Switch Ports



In-Rack InfiniBand Switch-to-Server Cable Connections

Table B-6 lists the location, ports and cables for the InfiniBand connections from switch to server.

Table B-6 InfiniBand Switch-to-Server Cabling

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit ¹	Port ²	Cable Length
U24	0A	U41	PCIe 3 P1	3 meters

Table B-6 (Cont.) InfiniBand Switch-to-Server Cabling

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit ¹	Port ²	Cable Length
U24	0B	U39	PCIe 3 P1	3 meters
U24	1A	U37	PCIe 3 P1	3 meters
U24	1B	U35	PCIe 3 P1	3 meters
U24	2A	U33	PCIe 3 P1	3 meters
U24	2B	U31	PCIe 3 P1	3 meters
U24	3A	U29	PCIe 3 P1	3 meters
U24	3B	U27	PCIe 3 P1	2 meters
U24	4A	U25	PCIe 3 P1	2 meters
U20	0A	U41	PCIe 3 P2	3 meters
U20	0B	U39	PCIe 3 P2	3 meters
U20	1A	U37	PCIe 3 P2	3 meters
U20	1B	U35	PCIe 3 P2	3 meters
U20	2A	U33	PCIe 3 P2	3 meters
U20	2B	U31	PCIe 3 P2	3 meters
U20	3A	U29	PCIe 3 P2	3 meters
U20	3B	U27	PCIe 3 P2	2 meters
U20	4A	U25	PCIe 3 P2	2 meters
U20	4B	U18	PCIe 3 P2	2 meters
U20	12A	U16	PCIe 3 P2	2 meters
U20	12B	U14	PCIe 3 P1	2 meters
U20	13A	U12	PCIe 3 P1	2 meters
U20	13B	U10	PCIe 3 P1	2 meters
U20	14A	U8	PCIe 3 P1	2 meters
U20	14B	U6	PCIe 3 P1	2 meters
U20	15A	U4	PCIe 3 P1	3 meters
U20	15B	U2	PCIe 3 P1	3 meters
U24	4B	U18	PCIe 3 P1	2 meters
U24	12A	U16	PCIe 3 P1	2 meters
U24	12B	U14	PCIe 3 P2	2 meters
U24	13A	U12	PCIe 3 P2	2 meters
U24	13B	U10	PCIe 3 P2	2 meters
U24	14A	U8	PCIe 3 P2	2 meters
U24	14B	U6	PCIe 3 P2	3 meters
U24	15A	U4	PCIe 3 P2	3 meters
U24	15B	U2	PCIe 3 P2	3 meters

¹ U_n is the unit location in the rack, where n is the number.² P_n is the InfiniBand port, where n is port number.

Multi-Rack Cabling Tables

This appendix contains the tables for multi-rack cabling. It contains the following topics:

- [Understanding Multi-Rack Cabling](#)
- [Key to Table Abbreviations](#)
- [Two Rack Cabling](#)
- [Three Rack Cabling](#)
- [Four Rack Cabling](#)
- [Five Rack Cabling](#)
- [Six Rack Cabling](#)
- [Seven Rack Cabling](#)
- [Eight Rack Cabling](#)

Understanding Multi-Rack Cabling


You can connect multiple Oracle Big Data Appliance Racks or a combination of Oracle Big Data Appliance Racks and Exadata Database Machine Full Racks. Connection to partial racks is not supported.

You need three Sun Network QDR InfiniBand Gateway Switches when cabling multiple racks. Oracle Big Data Appliance Rack has a Sun Datacenter QDR InfiniBand Switch for the spine switch at rack unit 1 (U1). The rack also has two Sun Network QDR InfiniBand Gateway Switches for the leaf switches at rack units 20 and 24 (U20 and U24). See [Figure B-1, "Rack Layout of Oracle Big Data Appliance Rack"](#) on page B-2.

These switches attach to standard Quad Small Form-factor Pluggable (QSFP) connectors at the end of the InfiniBand cables. You can cable together up to eight racks without external InfiniBand switches. The cable lengths in this appendix are for adjacent racks. If they are apart, then longer cables may be required for the connections.


In a single rack, the two leaf switches are connected by seven cables. In addition, each leaf switch has one cable to the spine switch, as shown in [Figure C-1](#). The total for the rack is 45 InfiniBand cables.

Figure C-1 Connections Between Spine Switch and Leaf Switches



The Sun Fire Servers connect to the leaf switches as shown in [Figure C-2](#).

Figure C-2 Connections Between Sun Fire Servers and Leaf Switches



When connecting multiple racks, remove the seven inter-switch connections between leaf switches and the two connections between the leaf switches and the spine switch. From each leaf switch, distribute eight connections over the spine switches in all racks. In multi-rack environments, the leaf switches inside a rack are not directly interconnected, as shown in [Figure C-3](#).


Figure C-3 Connecting Switches Across Two Racks


Figure C-3 shows that 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 Figure C-4.

Figure C-4 Connecting Switches Across Multiple Racks

As shown Figure C-4, 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.

Key to Table Abbreviations

The following abbreviations are used in the tables:

Abbreviation	Description
R_n	Rack n , where n is the number of the rack, such as R1.
IB_n	Unit location in rack, where n is the number, such as IB3.
P_n	InfiniBand port n , where n is port number, such as P8A.

Two Rack Cabling

Table C-1 shows the cable connections for the first spine switch (R1-U1) when cabling two racks.

Note:

- The spine switch is in U1 (IB1) for all racks.
 - The leaf switches are in U20 (IB2) and U24 (IB3).
 - In the following table, the leaf switches are called IB2 and IB3.
-
-

Table C-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

Table C-2 shows the cable connections for the second spine switch (R2-U1) when cabling two racks.

Table C-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

Table C-2 (Cont.) Leaf Switch Connections for the Second Rack in a Two Rack System

Leaf Switch	Connection	Cable Length
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

Three Rack Cabling

Table C-3 shows the cable connections for the first spine switch (R1-U1) when cabling three racks.

Note:

- The spine switch is in U1 (IB1) for all racks.
 - The leaf switches are in U20 (IB2) and U24 (IB3).
 - In the following table, the leaf switches are called IB2 and IB3.
-
-

Table C-3 Leaf Switch Connections for the First Rack in a Three Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-U24-P8A to R1-U1-P3A R1-U24-P8B to R1-U1-P4A R1-U24-P9A to R1-U1-P5A	3 meters
R1 IB3 to Rack 2	R1-U24-P9B to R2-U1-P6A R1-U24-P10A to R2-U1-P7A R1-U24-P10B to R2-U1-P8A	5 meters
R1 U24 to Rack 3	R1-U24-P11A to R3-U1-P9A R1-U24-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

Table C-3 (Cont.) Leaf Switch Connections for the First Rack in a Three Rack System

Leaf Switch	Connection	Cable Length
R1 IB2 to Rack 3	R1-IB2-P11A to R3-U1-P9B R1-IB2-P11B to R3-U1-P10B	5 meters

Table C-4 shows the cable connections for the second spine switch (R2-U1) when cabling three racks.

Table C-4 Leaf Switch Connections for the Second Rack in a Three Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-U24-P8A to R2-U1-P3A R2-U24-P8B to R2-U1-P4A R2-U24-P9A to R2-U1-P5A	3 meters
R2 IB3 to Rack 1	R2-U24-P11A to R1-U1-P9A R2-U24-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-U24-P9B to R3-U1-P6A R2-U24-P10A to R3-U1-P7A R2-U24-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 C-5 shows the cable connections for the third spine switch (R3-U1) when cabling three racks.

Table C-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

Table C-5 (Cont.) Leaf Switch Connections for the Third Rack in a Three Rack System

Leaf Switch	Connection	Cable Length
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

Table C-6 shows the cable connections for the first spine switch (R1-U1) when cabling four racks.

Note:

- The spine switch is in U1 (IB1) for all racks.
 - The leaf switches are in U20 (IB2) and U24 (IB3).
 - In the following table, the leaf switches are called IB2 and IB3.
-
-

Table C-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 C-7 shows the cable connections for the second spine switch (R2-U1) when cabling four racks.

Table C-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 C-8 shows the cable connections for the third spine switch (R3-U1) when cabling four racks.

Table C-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 C-9 shows the cable connections for the fourth spine switch (R4-U1) when cabling four racks.

Table C–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

Table C–10 shows the cable connections for the first spine switch (R1-U1) when cabling five racks.

Note:

- The spine switch is in U1 (IB1) for all racks.
 - The leaf switches are in U20 (IB2) and U24 (IB3).
 - In the following table, the leaf switches are called IB2 and IB3.
-
-

Table C–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

Table C–10 (Cont.) Leaf Switch Connections for the First Rack in a Five Rack System

Leaf Switch	Connection	Cable Length
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B R1-IB2-P9B to R2-U1-P6B	3 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 C–11 shows the cable connections for the second spine switch (R2-U1) when cabling five racks.

Table C–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 C–12 shows the cable connections for the third spine switch (R3-U1) when cabling five racks.

Table C–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

Table C–12 (Cont.) Leaf Switch Connections for the Third Rack in a Five Rack System

Leaf Switch	Connection	Cable Length
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 C–13 shows the cable connections for the fourth spine switch (R4-U1) when cabling five racks.

Table C–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 C–14 shows the cable connections for the fifth spine switch (R5-U1) when cabling five racks.

Table C–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

Table C–14 (Cont.) Leaf Switch Connections for the Fifth Rack in a Five Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 to Rack 1	R5-IB3-P9A to R1-U1-P5A R5-IB3-P9B to R1-U1-P6A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10A to R2-U1-P7A R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B R5-IB2-P8B to R5-U1-P4B	3 meters
R5 IB2 to Rack 1	R5-IB2-P9A to R1-U1-P5B R5-IB2-P9B to R1-U1-P6B	10 meters
R5 IB2 to Rack 2	R5-IB2-P10A to R2-U1-P7B R5-IB3-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters

Six Rack Cabling

Table C–15 shows the cable connections for the first spine switch (R1-U1) when cabling six racks.

Note:

- The spine switch is in U1 (IB1) for all racks.
 - The leaf switches are in U20 (IB2) and U24 (IB3).
 - In the following table, the leaf switches are called IB2 and IB3.
-
-

Table C–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

Table C–15 (Cont.) Leaf Switch Connections for the First Rack in a Six Rack System

Leaf Switch	Connection	Cable Length
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 C–16 shows the cable connections for the second spine switch (R2-U1) when cabling six racks.

Table C–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 C–17 shows the cable connections for the third spine switch (R3-U1) when cabling six racks.

Table C–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	5 meters

Table C–17 (Cont.) Leaf Switch Connections for the Third Rack in a Six Rack System

Leaf Switch	Connection	Cable Length
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	5 meters

Table C–18 shows the cable connections for the fourth spine switch (R4-U1) when cabling six racks.

Table C–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 C–19 shows the cable connections for the fifth spine switch (R5-U1) when cabling six racks.

Table C–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

Table C–19 (Cont.) Leaf Switch Connections for the Fifth Rack in a Six Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P9A to R6-U1-P5A R5-IB3-P9B to R6-U1-P6A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B R5-IB2-P8B to R5-U1-P4B	3 meters
R5 IB2 to Rack 1	R5-IB2-P10A to R1-U1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB3-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P9A to R6-U1-P5B R5-IB2-P9B to R6-U1-P6B	5 meters

Table C–20 shows the cable connections for the sixth spine switch (R6-U1) when cabling six racks.

Table C–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	5 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 2	R6-IB3-P10A to R2-U1-P7B	10 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 3	R6-IB2-P10B to R3-U1-P8B	5 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

Table C–21 shows the cable connections for the first spine switch (R1-U1) when cabling seven racks.

Note:

- The spine switch is in U1 (IB1) for all racks.
- The leaf switches are in U20 (IB2) and U24 (IB3).
- In the following table, the leaf switches are called IB2 and IB3.

Table C–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 C–22 shows the cable connections for the second spine switch (R2-U1) when cabling seven racks.

Table C–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

Table C–22 (Cont.) Leaf Switch Connections for the Second Rack in a Seven Rack

Leaf Switch	Connection	Cable Length
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 C–23 shows the cable connections for the third spine switch (R3-U1) when cabling seven racks.

Table C–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 C–24 shows the cable connections for the fourth spine switch (R4-U1) when cabling seven racks.

Table C–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

Table C–24 (Cont.) Leaf Switch Connections for the Fourth Rack in a Seven Rack

Leaf Switch	Connection	Cable Length
R4 IB3 to Rack 7	R4-IB3-P10A to R7-U1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B R4-IB2-P8B to R4-U1-P4B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-U1-P5B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9B to R6-U1-P6B	5 meters
R4 IB2 to Rack 7	R4-IB2-P10A to R7-U1-P7B	10 meters

Table C–25 shows the cable connections for the fifth spine switch (R5-U1) when cabling seven racks.

Table C–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-IB3-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P9A to R6-U1-P5B	5 meters
R5 IB2 to Rack 7	R5-IB2-P9B to R7-U1-P6B	5 meters

Table C–26 shows the cable connections for the sixth spine switch (R6-U1) when cabling seven racks.

Table C–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

Table C–26 (Cont.) Leaf Switch Connections for the Sixth Rack in a Seven Rack System

Leaf Switch	Connection	Cable Length
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	5 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-U1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-U1-P10A	5 meters
R6 IB3 to Rack 7	R6-IB3-P9A to R7-U1-P5A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-U1-P3B R6-IB2-P8B to R6-U1-P4B	3 meters
R6 IB2 to Rack 1	R6-IB2-P9B to R1-U1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB3-P10A to R2-U1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-U1-P8B	5 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters
R6 IB2 to Rack 7	R6-IB3-P9A to R7-U1-P5B	5 meters

Table C–27 shows the cable connections for the seventh spine switch (R7-U1) when cabling seven racks.

Table C–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-IB3-P9B to R2-U1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-U1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-U1-P8B	10 meters
R7 IB2 to Rack 5	R7-IB2-P11A to R5-U1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB3-P11B to R6-U1-P10B	5 meters

Eight Rack Cabling

Table C–28 shows the cable connections for the first spine switch (R1-U1) when cabling eight racks.

Note:

- The spine switch is in U1 (IB1) for all racks.
- The leaf switches are in U20 (IB2) and U24 (IB3).
- In the following table, the leaf switches are called IB2 and IB3.

Table C–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-P8B	10 meters
R1 IB2 to Rack 8	R1-IB2-P11B to R8-U1-P10B	10 meters

Table C–29 shows the cable connections for the second spine switch (R2-U1) when cabling eight racks.

Table C–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
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

Table C–29 (Cont.) Leaf Switch Connections for the Second Rack in a Eight Rack

Leaf Switch	Connection	Cable Length
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 C–30 shows the cable connections for the third spine switch (R3-U1) when cabling eight racks.

Table C–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	5 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	5 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 C–31 shows the cable connections for the fourth spine switch (R4-U1) when cabling eight racks.

Table C–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

Table C–31 (Cont.) Leaf Switch Connections for the Fourth Rack in a Eight Rack System

Leaf Switch	Connection	Cable Length
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 C–32 shows the cable connections for the fifth spine switch (R5-U1) when cabling eight racks.

Table C–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 C–33 shows the cable connections for the sixth spine switch (R6-U1) when cabling eight racks.

Table C-33 Leaf Switch Connections for the Sixth Rack in an 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	5 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	5 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters
R6 IB2 to Rack 7	R6-IB3-P8B to R7-U1-P4B	5 meters
R6 IB2 to Rack 8	R6-IB2-P9A to R8-U1-P5B	5 meters

Table C-34 shows the cable connections for the seventh spine switch (R7-U1) when cabling eight racks.

Table C-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-IB3-P9B to R2-U1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-U1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-U1-P8B	10 meters
R7 IB2 to Rack 5	R7-IB2-P11A to R5-U1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB3-P11B to R6-U1-P10B	5 meters
R7 IB2 to Rack 8	R7-IB3-P8B to R8-U1-P4B	5 meters

Table C–35 shows the cable connections for the eighth spine switch (R8-U1) when cabling eight racks.

Table C–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	5 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-IB3-P9A to R2-U1-P5B	10 meters
R8 IB2 to Rack 3	R8-IB2-P9B to R3-U1-P6B	10 meters
R8 IB2 to Rack 4	R8-IB2-P10A to R4-U1-P7B	10 meters
R8 IB2 to Rack 5	R8-IB2-P10B to R5-U1-P8B	5 meters
R8 IB2 to Rack 6	R8-IB3-P11A to R6-U1-P9B	5 meters
R8 IB2 to Rack 7	R8-IB3-P1B to R7-U1-P10B	5 meters

Index

A

access route requirements, 2-3
administrative network configuration, 8-18
airflow, 2-8
ambient relative humidity range, 2-9
ambient temperature range, 2-9, 14-1
ASM package, 12-3
ASR
 about
 ASR Manager registration, 12-4
 assets, 12-5
 configuration
ASR Manager, 4-7
Auto Service Request
 See ASR

B

batteries
 life expectancy, A-2
bda_imaging_status file, 8-18
bdacheckhw utility, 8-26
bdacheckib utility, 8-25
bdachecknet command, 8-27
bdachecksw utility, 8-27
BdaDeploy.json file, 4-3
bdadiag utility, 8-28
bdaid utility, 8-29
bdaimagevalidate utility, 8-29
bda-install-template.pdf file, 4-3
bdaredoclientnet utility, 8-29
bdaserials utility, 8-30
Big Data Connectors, 4-7
bondib0, 3-2

C

cabling
 multi-rack, 9-2, C-1
changing IP addresses, 14-6
checking firmware, 8-17
circuit breakers for PDU, 5-13
Cisco Ethernet switch configuration, 8-3
Cisco switch
 configuring default gateway, 8-5

 configuring host name, 8-5
 default port speed, 8-4
 setting password, 8-5
client access network, 3-2
Cloudera Manager, 4-8
Cloudera's Distribution including Apache
 Hadoop, 1-1
commands, remote execution, 1-3, 7-1
configuration files, 8-18
 ntp.conf file, 14-12
Configuration Template, 4-3
configuring
 ASR
copying files, 7-2
createvlan command, 11-1
createvnic command, 10-3

D

dba group, 13-6
dcli utility, 1-3, 7-1
default serial port speed
 Cisco switch, 8-4
disablesm command, 14-9
disk controllers, 14-4

E

earth ground cable, 5-9
electrostatic discharge, 2-9
email configuration, 4-8
environment.pp file, 13-5
Ethernet over InfiniBand configuration, 10-1
Ethernet switch configuration, 8-3

F

Fabric Management, 8-11
factory software image, 8-18
first server location, 8-19

G

ground cable, 5-9
groups, 13-6

H

- hadoop group, 13-6
- hardware checks, 8-15
- hardware diagnostics, 6-1
- hdfs user, 13-6
- humidity range, 2-9

I

- iblinkinfo command, 8-30, 10-2
- ibstat command, 10-2
- ibswitches command, 10-3
- ILOM, 6-1
- ilom-admin user, 6-2
- ilom-operator user, 6-2
- imagehistory utility, 8-31
- imageinfo command, 8-18
- imageinfo utility, 8-31
- InfiniBand configuration, 8-9
- InfiniBand network configuration, 8-18
- InfiniBand private network, 3-2
- InfiniBand switches
 - changing IP addresses, 14-6
- IP Addresses
 - changing for InfiniBand switches, 14-6
- IP addresses
 - changing, 14-6

K

- KVM configuration, 8-1

L

- leveling feet, 5-6
- licenses, third party, 13-6
- lights out management, 6-1
- listlinkup command, 8-31, 10-1
- locating
 - first server, 8-19
 - PDU circuit breakers, 5-13
- logical volume for Oracle NoSQL Database, 13-6

M

- Mammoth Utility, 13-1
 - reasons for failure, 13-4
- mammoth-rack_name.params file, 4-3
- management network, 3-2
- management network configuration, 8-18
- mapred user, 13-6
- MegaCli64 command, 14-5
- minimum ceiling height, 2-2
- mount point
 - logical volume for Oracle NoSQL Database, 13-6
- mounting a device, 8-19
- mounting brackets, 5-4, 5-6

N

- NET0, 3-2
- NET1, 3-2
- NET2, 3-2
- network configuration, 8-18
- network connections, 8-21
- Network Time Protocol, 14-12
- networks
 - client access, 2-11, 3-2
 - InfiniBand, 3-2
 - management, 2-11, 3-2
- networksetup scripts, 8-20, 8-21
- NoSQL Database
 - logical volume, 13-6
- NTP servers, 14-12
- ntp.conf file, 14-12

O

- oinstall group, 13-6
- Oracle Big Data Connectors, 13-4
- Oracle Data Integrator, 4-7, 4-8
- Oracle Exadata Database Machine, 4-8
- Oracle Inventory Group, 4-8
- Oracle Linux, 1-1
- Oracle NoSQL Database, 4-7, 13-6

P

- passwordless SSH, 7-1
- PDU
 - circuit breakers, 5-13
- PDU configuration, 8-13
- PDU thresholds, 8-14
- perforated tiles, 2-10
- power off
 - emergency, 14-3
 - sequence, 14-3
 - switch, 14-3
- power on
 - sequence for servers, 14-2
- powering off, 14-3
- preinstall-checkip script, 4-3, 4-4
- Puppet configuration tool, 13-6

R

- registering ASR Manager, 12-4
- relative humidity, 2-9, 2-10
- remove-root-ssh, 7-2
- repair categories, A-1
- replacing spine switch, 14-9
- RoHS:Y compliance, A-2
- RoHS:YL compliance, A-2

S

- safety guidelines, 5-1
- SCAN addresses, 3-2
- scp file copy utility, 7-2

- server hardware checks, 8-15
- setting
 - Cisco switch time zone, 8-6
- setup-root-ssh, 7-1
- shipping kit, 5-3
- showusb command, 8-19
- showvlan command, 8-32
- showvnic command, 8-32, 10-4
- shutdown command, 14-3
- SMTP configuration, 4-8
- software image, 8-18
- software source code, third party, 13-6
- spanning tree, 8-7
- spares kit, 1-2
- spine switches
 - replacing, 14-9
- stabilizing Big Data Appliance, 5-6
- Sun Online Account, 12-2

T

- temperature range, 2-9
- third-party licenses, 13-6
- third-party software source code, 13-6
- thresholds for PDUs, 8-14
- time zones, 8-6
- top clearance, 2-2

U

- USB flash drive, 8-19
- users, 13-6

V

- verifying
 - database server configuration, 14-4
 - virtual drive configuration, 14-5
- VIP addresses, 3-2
- Virtual LAN creation, 11-1
- virtual local area networks, 3-2
- Virtual NIC configuration, 10-1
- VLANs, 3-2, 8-4

