

Oracle® Big Data Appliance

Owner's Guide

Release 1 (1.1)

E36161-06

March 2013

Describes data center site planning, network configuration, hardware and software installation, and maintenance of Oracle Big Data Appliance.

Copyright © 2012, 2013, 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 END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

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

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

This software or hardware and documentation may provide access to or information 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	xv
Audience.....	xv
Documentation Accessibility	xv
Related Documentation.....	xv
Conventions	xvi
Backus-Naur Form Syntax	xvi
Changes in This Release for Oracle Big Data Appliance Owner's Guide	xvii
Changes in Oracle Big Data Appliance Release 1 (1.1)	xvii
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 Parts 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 That 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
How the Servers Connect to the Gateway Switches	3-3
Using Splitter Cables for Connecting to the 40 Gbps Physical Ports.....	3-4
What About Data Centers Without a 10 GbE Infrastructure?	3-4
Configuring Multiple Racks into a Single Hadoop Cluster	3-5
Factory Network Settings	3-5
Port Assignments	3-6

4 Using the Oracle Big Data Appliance Configuration Utility

Overview of the Oracle Big Data Appliance Configuration Utility	4-1
Using the Oracle Big Data Appliance Configuration Utility	4-1
Using OpenOffice Spreadsheet	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
Removing Oracle Big Data Appliance from the Shipping Crate	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 for the First Time	5-10
Inspecting the Oracle Big Data Appliance Rack After It Is in Place	5-10
Connecting Power Cords	5-10
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
Oracle ILOM IP Addresses for Oracle Big Data Appliance Components	6-3
Connecting to Oracle 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 Oracle ILOM Using a Serial Connection	6-5
Connecting to the Oracle ILOM of a Server	6-5
Connecting to the Oracle ILOM of a Gateway Switch.....	6-5

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
dcli Return Values.....	7-5
dcli Examples	7-5

8 Configuring Oracle Big Data Appliance

Configuring the KVM Switch.....	8-1
Starting the KVM Switch.....	8-2
Connecting the KVM Switch to the Management Network.....	8-2
Checking the KVM Firmware Version.....	8-3
Configuring the KVM Switch to Access the Servers.....	8-3
Accessing a Server by Using the KVM Switch.....	8-3
Configuring the Cisco Ethernet Switch	8-4
Scope of the Configuration	8-4
Prerequisites for Configuring the Ethernet Switch	8-4
Configuring the Ethernet Switch on the Customer Network.....	8-5
Setting Up Telnet Access.....	8-6
Identifying the DNS Servers.....	8-7
Setting the Clock and Time Zone.....	8-7
Configuring the NTP Servers	8-8
Disabling the Spanning Tree	8-9
Verifying the Ethernet Configuration	8-9
Configuring the InfiniBand Leaf and Spine Switches	8-10
Configuring an InfiniBand Switch.....	8-10
Setting the Time Zone and Clock on an InfiniBand Switch.....	8-12
Checking the Health of an InfiniBand Switch.....	8-12
Configuring the Power Distribution Units	8-15
Connecting the PDUs to the Network.....	8-15
Verifying the PDU Firmware Version.....	8-15
Configuring the Threshold Settings for the PDUs	8-16
Configuring the Oracle Big Data Appliance Servers	8-17
Configuring the Network	8-20
Verifying the Factory Software Image	8-20
Copying the Configuration Files to Oracle Big Data Appliance	8-21
Starting the Network Configuration	8-22
Connecting to the Network	8-23
Completing the Network Configuration	8-23
Reinstalling the Base Image.....	8-25

Checking the Health of the Network	8-26
bdacheckcluster	8-27
bdacheckhw	8-27
bdacheckib	8-28
bdachecknet	8-29
bdachecksw	8-30
bdadiag	8-30
bdaid	8-31
bdaimagevalidate	8-31
bdaredoclientnet	8-31
bdaserials	8-32
iblinkinfo	8-33
imagehistory	8-33
imageinfo	8-33
listlinkup	8-34
showvlan	8-34
showvnics	8-34

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 Setting Up Auto Service Request

Understanding Auto Service Request	10-1
Getting Ready for ASR Installation	10-2
Installing ASR Manager	10-3
Downloading the Software	10-3
Installing the Oracle Automated Service Manager Package	10-3
Installing the Oracle Service Tags Bundle	10-4
Registering ASR Manager	10-4
Verifying ASR Manager	10-4
Configuring the Trap Destinations	10-5
Activating ASR Assets	10-5
Validating Auto Service Request on Oracle Big Data Appliance	10-6
Troubleshooting ASR	10-7

11 Installing the Oracle Big Data Appliance Software

Using the Mammoth Utility	11-1
Installing the Software on a Single or Primary Rack	11-2
Upgrading the Software on Oracle Big Data Appliance	11-3
Mammoth Utility Syntax	11-3
What If an Error Occurs During the Installation?	11-4
Adding a Rack to an Existing Cluster	11-4
Mammoth Utility Steps	11-6
Using the Mammoth Reconfiguration Utility	11-10

Changing the Software Configuration.....	11-10
Mammoth Reconfiguration Utility Syntax.....	11-10
12 Maintaining Oracle Big Data Appliance	
Monitoring the Ambient Temperature of Servers.....	12-1
Powering On and Off Oracle Big Data Appliance	12-2
Nonemergency Power Procedures	12-2
Emergency Power-Off Considerations	12-3
Cautions and Warnings.....	12-3
Adding Memory to a Server.....	12-4
Maintaining the Physical Disks of Servers.....	12-4
Verifying the Server Configuration.....	12-4
Replacing a Server Disk.....	12-6
Overview of the Disk Replacement Process.....	12-6
About Disk Drive Identifiers	12-6
Dismounting Partitions Before Replacing a Working Disk	12-7
What If a Server Fails to Restart?	12-9
Replacing a Disk Drive.....	12-10
Identifying the Function of a Disk Drive	12-12
Configuring an Operating System Disk.....	12-13
Configuring an Oracle NoSQL Database Disk	12-19
Configuring an HDFS Disk.....	12-22
Verifying the Disk Configuration	12-23
Changing InfiniBand IP Addresses.....	12-24
Maintaining the InfiniBand Network.....	12-25
Backing Up and Restoring Oracle ILOM Settings.....	12-25
Replacing a Failed InfiniBand Switch	12-26
Verifying InfiniBand Network Operation	12-27
Understanding the Network Subnet Manager Master	12-29
Changing the NTP Servers	12-30
A Error Messages for Oracle Big Data Appliance	
Auto Service Request Alerts: DISKALRT-02001 to DISKALRT-02006	A-1
B Replacement Units	
Understanding Repair Categories.....	B-1
Disk Controller Batteries	B-2
LED Status Descriptions	B-2
Sun Fire Server LEDs.....	B-2
Sun Datacenter InfiniBand Switch 36 LEDs	B-3
Parts for Sun Fire Servers	B-4
Parts for the Sun Datacenter InfiniBand Switch 36.....	B-7
Parts for the Sun Network QDR InfiniBand Gateway Switches.....	B-8
Parts for the Cisco Ethernet Switch	B-9
Parts for the KVM Switch and KMM Tray Components	B-9
Parts for Power Distribution Units	B-10

Cables and Cable Accessories for an Oracle Big Data Appliance Rack	B-10
--	------

C In-Rack Cabling Tables

Cable Color Coding	C-1
Oracle Big Data Appliance Rack Layout	C-1
KVM Network Cable Connections	C-2
Administrative Gigabit Ethernet Cable Connections	C-3
Oracle Integrated Lights Out Manager Cable Connections	C-4
Single-Phase Power Distribution Unit Cable Connections	C-5
Three-Phase Power Distribution Unit Cable Connections	C-5
InfiniBand Switch-to-Switch Cable Connections	C-6
InfiniBand Switch-to-Server Cable Connections	C-7

D Multirack Cabling Tables

Understanding Multirack Cabling	D-1
Cabling the Switches in One Rack	D-1
Cabling the Switches in Two Racks	D-2
Cabling the Switches in Up to Eight Racks	D-3
Key to Cabling Table Abbreviations	D-3
Two-Rack Cabling	D-4
Three-Rack Cabling	D-5
Four-Rack Cabling	D-7
Five-Rack Cabling	D-9
Six-Rack Cabling	D-12
Seven-Rack Cabling	D-15
Eight-Rack Cabling	D-19

Index

List of Examples

4-1	Output from preinstall-checkip.sh	4-5
8-1	Sample Output from networksetup-one	8-22
8-2	Sample Output from networksetup-two	8-23
11-1	Mammoth Utility Syntax Examples	11-4

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 the Shipping Crate.....	5-3
5-2	Carefully Push Oracle Big Data Appliance from Behind.....	5-5
5-3	Stabilizing the Oracle Big Data Appliance Rack Using Leveling Feet	5-6
5-4	Location of Mounting Brackets on the Rack	5-7
5-5	Bottom View of the Oracle Big Data Appliance Rack	5-7
5-6	Base Position of the Rack to Measure (in mm)	5-8
5-7	Securing the Oracle Big Data Appliance Rack Using Mounting Brackets	5-9
5-8	Earth Ground Attachment Bolt Location	5-10
5-9	Location of Power Cord Cable Ties.....	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
12-1	Oracle ILOM Configuration Backup.....	12-26
C-1	Rack Layout of Oracle Big Data Appliance.....	C-2
C-2	Sun Network QDR InfiniBand Gateway Switch Ports	C-7
D-1	Connections Between Spine Switch and Leaf Switches	D-2
D-2	Connections Between the Servers and the Leaf Switches	D-2
D-3	Connecting Switches Across Two Racks	D-3
D-4	Connecting Switches Across Multiple Racks.....	D-3

List of Tables

1-1	Oracle Big Data Appliance Rack Components	1-2
1-2	Sun Fire X4270 M2 Server Components	1-2
1-3	Bundled Components and Accessories	1-2
1-4	Bundled Cables	1-3
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	General Properties for the Rack	4-5
4-2	Network Properties for the Rack	4-6
4-3	Network Properties for Servers	4-6
4-4	Network Properties for Switches	4-6
4-5	Network Properties for Power Distribution Units	4-7
4-6	Installed Components	4-7
4-7	Oracle Auto Service Request	4-8
4-8	Users, Groups, and Passwords	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	8-16
8-2	Threshold Values for Three-Phase, Low-Voltage PDU	8-16
8-3	Threshold Values for Single-Phase, High-Voltage PDU	8-17
8-4	Threshold Values for Three-Phase, High-Voltage PDU	8-17
11-1	Mammoth Utility Options	11-4
12-1	Disk Drive Identifiers	12-7
12-2	Mount Points	12-7
B-1	Oracle Big Data Appliance Repair Categories	B-1
B-2	LED Status Descriptions	B-3
B-3	Sun Datacenter InfiniBand Switch 36 LED Status Descriptions	B-3
B-4	Replacement Parts for a Sun Fire X4270 M2 Server	B-4
B-5	Replacement Parts for the Sun Datacenter InfiniBand Switch 36	B-8
B-6	Replacement Parts for the Sun Network QDR InfiniBand Gateway Switch	B-8
B-7	Replacement Parts for Cisco Catalyst Ethernet Switch	B-9
B-8	Replacement Parts for KVM Switch and KMM Tray Components	B-10
B-9	Replacement Parts for Power Distribution Units	B-10
B-10	Replacement Cables for an Oracle Big Data Appliance Rack	B-11
C-1	KVM Cabling	C-2
C-2	Administrative Gigabit Ethernet Cabling	C-3
C-3	Oracle ILOM Cabling	C-4
C-4	Single-Phase PDU Cabling	C-5
C-5	Three-Phase PDU Cabling	C-6
C-6	InfiniBand Switch-to-Switch Cabling	C-6
C-7	InfiniBand Switch-to-Server Cabling	C-7
D-1	Leaf Switch Connections for the First Rack in a Two-Rack System	D-4
D-2	Leaf Switch Connections for the Second Rack in a Two-Rack System	D-4

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

Preface

This guide describes Oracle Big Data Appliance, which is used 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/E36174_01/index.htm

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

- Sun Rack II 1042 and 1242 library:
<http://docs.oracle.com/cd/E19844-01/>
- Sun Network QDR InfiniBand Gateway Switch library:
http://docs.oracle.com/cd/E26699_01/index.html
- Sun Datacenter InfiniBand Switch 36 library:
http://docs.oracle.com/cd/E26698_01/index.html
- Oracle Integrated Lights Out Manager (ILOM) 3.0 library:
<http://docs.oracle.com/cd/E19860-01/>

Conventions

The following text conventions are used in this document:

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

Changes in This Release for Oracle Big Data Appliance Owner's Guide

This preface lists the changes in *Oracle Big Data Appliance Owner's Guide*.

Changes in Oracle Big Data Appliance Release 1 (1.1)

The following are changes in *Oracle Big Data Appliance Owner's Guide* for Oracle Big Data Appliance Release 1 (1.1).

New Features

- Software upgrade
See "[Upgrading the Software on Oracle Big Data Appliance](#)" on page 11-3..
- Server disk replacement
See "[Replacing a Server Disk](#)" on page 12-6.
- ASR alerts
See "[Auto Service Request Alerts: DISKALRT-02001 to DISKALRT-02006](#)" on page A-1.

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 sections:

- [About Oracle Big Data Appliance](#)
- [Oracle Big Data Appliance Rack](#)
- [Server Components](#)
- [Spare Parts 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, web logs, photographs, smart meters, sensors, and similar devices.

Oracle Big Data Appliance is engineered to work with Oracle Exadata Database Machine and Oracle Exalytics In-Memory 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
1	Sun Rack II 1242 base
2	NM2-GW Sun Network QDR InfiniBand Gateway Switch
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	8 GB DDR3 2RX4 1333 MHz DIMMs (48 GB RAM)
12	3 TB 3.5-inch 7200 RPM drives
1	Ethernet Port for Oracle Integrated Lights Out Manager v3.0 for remote management
1	Dual-port 4X QDR (40 Gbps) 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 4 GB flash drive

Spare Parts Kit Components

The Oracle Big Data Appliance spare parts 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	3 TB 3.5-inch 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 meters	10	QSFP passive copper cable
530-4445-01	3 meters	6	QSFP passive copper cable
350-1519-01	--	3	ASSY, NM2 serial cable sets

Table 1–5 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 meters	2	QSFP passive copper cable
530-4437-01	10 feet	1	Orange Cat 5
530-4435-01	10 feet	1	Black Cat 5
530-4433-01	10 feet	1	Red Cat 5
530-4432-01	7 feet	1	Blue Cat 5

Oracle Big Data Appliance Management Software

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

Note: See ["Checking the Health of the Network"](#) on page 8-26 except where directed otherwise.

- 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 an Oracle Big Data Appliance rack. See [Chapter 6](#).
- The `setup-root-ssh` utility 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.
- The `dcli` utility executes commands across a group of servers on Oracle Big Data Appliance and returns the output. See [Chapter 7](#).
- The `bdacheckcluster` utility checks the health of the cluster.
- The `bdacheckib` utility checks the private InfiniBand network.
- The `bdachecknet` utility checks the network configuration.
- The `iblinkinfo` utility lists the connections in the InfiniBand network.
- The `listlinkup` utility identifies the active Ethernet ports.
- The `showvlan` utility lists the virtual local area networks (VLANs) configured on a Sun Network QDR InfiniBand Gateway Switch.
- The `showvnics` utility lists the virtual network interface cards (VNICs) created on a Sun Network QDR InfiniBand Gateway Switch.

All end-user software is installed on site using the Mammoth Utility. See [Chapter 11](#).

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 can do the following:

- Replace the Gigabit Ethernet switch with an equivalent 1U 48-port Gigabit Ethernet switch that conforms to their internal data center network standards. This replacement must be performed by customers at their own expense and labor, after delivery of Oracle Big Data Appliance. If customers choose 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. Customers must supply the replacement hardware, and make or arrange for this change through other means.
- Remove the Cat 5 cables connected to the Ethernet switch and connect them to the customer network through an external switch or patch panel. Customers 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 Information Center: Oracle Big Data Appliance (ID 1445762.2) and its related notes.
 - Customers can update the firmware of the InfiniBand switches provided they comply with the validated versions documented in Information Center: Oracle Big Data Appliance (ID 1445762.2) and its related notes.
- Customers can load additional software on Oracle Big Data Appliance servers. Oracle does not support questions or issues with the nonstandard modules. If a server fails, and Oracle suspects that the failure may have been caused by a nonstandard module, then Oracle Support may refer the customer to the vendor of the nonstandard module or ask that the issue be reproduced without the nonstandard 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 or the Mammoth Reconfiguration 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 Internet small computer system interface (iSCSI) or network file system (NFS) protocols. 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 can 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 while you complete the *Oracle Big Data Appliance Site Checklists*.

This chapter contains these sections:

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

See Also: *Oracle Big Data Appliance Site Checklists*

General Environmental Requirements

Table 2–1 shows the general environmental requirements for Oracle Big Data Appliance. 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
Net Weight	909.5 kg (2005 lb) See Also: "Flooring Requirements" on page 2-3
Acoustic levels	8.3 Bel
Power	Maximum: 12.0 kW (12.25 kVA) Typical: 8.4 kW (8.6 kVA) See Also: "Electrical Power Requirements" on page 2-3
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 2-8 and "Ventilation and Cooling Requirements" on page 2-9

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 2-8 , and "Ventilation and Cooling Requirements" on page 2-9
IP addresses	24 for Ethernet network 18 for InfiniBand network See Also: Chapter 3, "Network Requirements."
Network drops	Minimum of two network connections See Also: Chapter 3, "Network Requirements."
External connectivity	2 x 1 Gbps Ethernet ports 18 x 10 Gbps 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: 200 cm (79 inches)
- Width: 60 cm with side panels (24 inches)
- Depth: 120 cm (47.5 inches)

The minimum ceiling height for the cabinet is 230 cm (90 inches), measured from the true floor or raised floor, whichever is higher. An additional 92 cm (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: 216 cm (85 inches)
- Shipping width: 122 cm (48 inches)
- Shipping depth: 157.5 cm (62 inches)
- Shipping weight: 993 kg (2189 lb)

If your loading dock meets the height and ramp requirements for a standard freight carrier truck, then you can use a pallet jack to unload the rack. If 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	218.4 cm (86 inches)	204 cm (80.5 inches)
Minimum door width	127 cm (50 inches)	64 cm (25.5 inches)
Minimum elevator depth	162.6 cm (64 inches)	124 cm (49 inches)
Maximum incline	6 degrees	6 degrees
Minimum elevator, pallet jack, and floor loading capacity	1134 kg (2500 lb)	1134 kg (2500 lb)

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 67.6 cm (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	91.4 cm (36 inches)
Front maintenance	91.4 cm (36 inches)
Top maintenance	91.4 cm (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	998 kg (2200 lb)
Maximum allowable weight of installed power distribution units	52 kg (115 lb)
Maximum dynamic load (maximum allowable weight of installed equipment including PDUs)	1050 kg (2315 lb)

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 preinstalled 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) locations, 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 single-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 30A single phase
Voltage	200 to 240 VAC
Frequency	50/60 Hz
Current	24A maximum per input
Power rating	15 kVA
Output current	72A (3 x 24A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection ¹	20A
Data center receptacle	15 kVA, with three 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, 30A at 200 to 240 VAC
- 6 receptacles to connect the PDUs to 6 NEMA L6-30 data center receptacles

[Figure 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 60A three-phase 4-wire
Voltage	190 to 220 VAC
Frequency	50/60 Hz
Current	40A maximum per phase
Power rating	14.4 kVA
Output current	69A (3 x 23A)
Outlets	42 x C13; 6 x C19
Outlet groups	6
Group protection ¹	20A
Data center receptacle	15 kVA IEC 60309 60A 4-pin 250 VAC 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, 60A at 190 to 220 VAC three-phase
- 2 receptacles to connect the PDUs to 2 IEC 60309 60A 4-pin 250 VAC three-phase IP67 data center receptacles

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

¹ UL489 1-pole circuit breaker

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

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

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

¹ UL489 1-pole circuit breaker

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

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

[Figure 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 a qualified electrician to determine what type of power is supplied to the building.

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

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

Circuit Breaker Requirements

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

In addition to 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. See [Table 2-1](#) for information about 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 nonoperating systems.

Table 2-9 Temperature, Humidity, and Altitude Requirements

Condition	Operating Requirement	Nonoperating Requirement	Optimum
Temperature	5 to 32 degrees Celsius (40 to 90 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, noncondensing	Up to 93 percent relative humidity	For optimal data center rack cooling, 45 to 50 percent, noncondensing
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 nonoperating 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. Maintaining humidity in a data center is not difficult, because a high-efficiency vapor barrier and a low rate of air changes are normally present.

Ventilation and Cooling Requirements

Always provide adequate space in front and behind the rack to allow for proper ventilation. Do not obstruct the front or rear of the rack with equipment or objects that might prevent air from flowing through the rack. Rack-mountable servers and equipment typically draw cool air in through the front of the rack and let out warm air through the rear of the rack. There is no air flow requirement for the left and right sides 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 it 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 form.

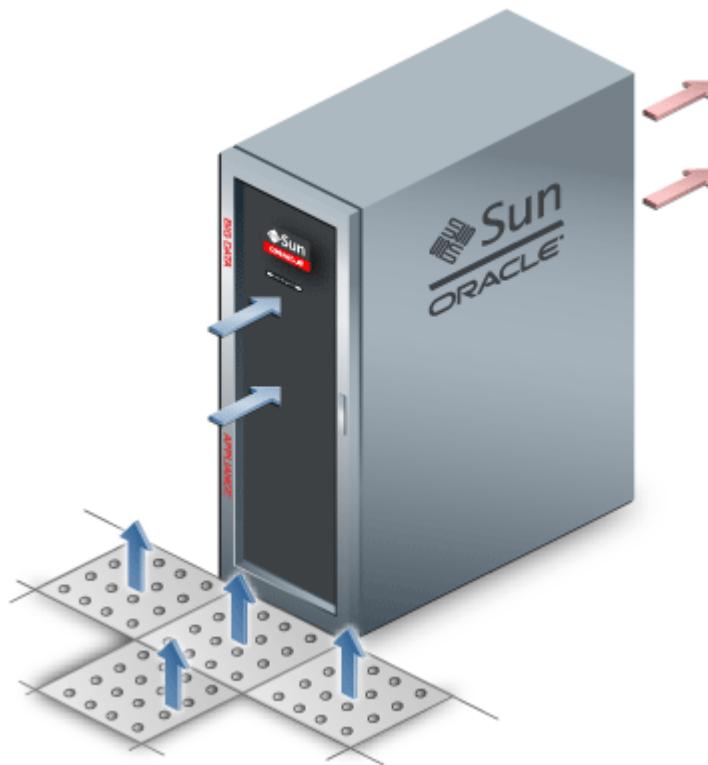
Air conditioning facilities usually do not precisely monitor or control temperature and humidity throughout an entire computer room. Generally, you should monitor individual points corresponding to multiple exhaust vents in the main unit and other units in the room, because the distribution of temperature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

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

Use perforated tiles, rated for 400 cubic feet per minute (CFM) per tile air flow, in front of the rack for cold air intake. The tiles can be arranged in any order in front of the rack, 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
 - One 1 Gbps Ethernet connection for the management switch in the rack
 - One 1 Gbps Ethernet connection for the KVM switch in the rack
- Client access network connection requirements
 - 2 (minimum) to 16 (maximum) 10 Gbps Ethernet connections split between the two Sun Datacenter InfiniBand Switch 36 modules in the rack. The exact number of connections depends on your bandwidth requirements.

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

Ensuring That the Site Is Ready

Before Oracle Big Data Appliance is delivered to the site, perform these tasks to ensure that 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 and ensure you are ready for delivery.

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 the *Oracle Big Data Appliance Configuration Worksheets* 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 performed by an Oracle representative, unless the owner performs the installation and configuration without assistance.

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 described earlier in this chapter, 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. For example, ensure that 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 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
- *Sun Rack II Power Distribution Units User's Guide* for the power distribution unit (PDU) power requirements at
<http://docs.oracle.com/cd/E19844-01/>

5. Install network cables for Oracle Big Data Appliance.

See Also: "Network Connection Requirements" on page 2-11

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.

See Also: "Temperature and Humidity Requirements" on page 2-8

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 sections:

- [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 allow clients to connect to them. Use the information in this chapter to configure the environment for Oracle Big Data Appliance.

Each server has the following network components and interfaces:

- 1 Dual-port 4X QDR (40 Gbps) 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 10 GbE connection. If they all 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 10 GbE connections in round-robin fashion. For example, if each switch has three available 10 GbE connections, then the VNIC on server 1 is assigned to 10 GbE 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 virtual local area network (VLAN). To assign different VLANs to different VNICs, you must delete the initial VNICs and manually create your own.

Minimum Requirements for the Networks

Additional configuration, such as defining multiple 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 your network meets the minimum 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 (GbE) network connects to your existing administrative network and is used to administer all components of Oracle Big Data Appliance. It connects the servers, Oracle ILOM, and switches connected to the Ethernet switch in the rack.

There are two uplinks to the administrative network:

- From the Ethernet switch in the rack
- From the KVM switch in the rack

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 Oracle 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 10 GbE 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 by 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 by using the InfiniBand switches on the rack and the bondib0 interface. This nonroutable 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.

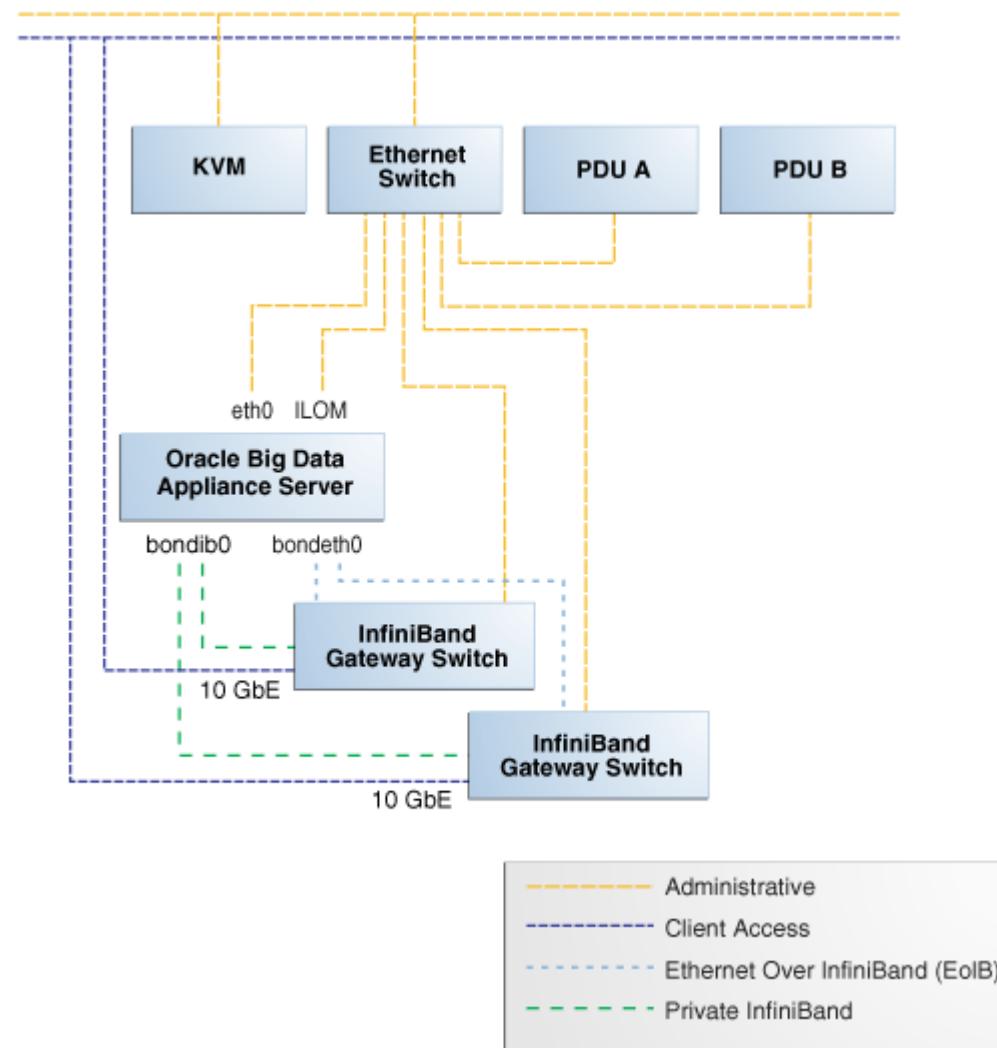
Network Diagram

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 Oracle 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 10 GbE ports. The two switches enable you to create up to 16 10 GbE 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 10 GbE network.

How the Servers Connect to the Gateway Switches

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 Oracle Big Data Appliance server has two virtual NICs (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 10 GbE 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 10 GbE 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 to 3	GW Switch 1 Link 1	GW Switch 2 Link 1
4 to 6	GW Switch 1 Link 2	GW Switch 2 Link 2
...
16 to 18	GW Switch 2 Link 3	GW Switch 1 Link 3

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

Using Splitter Cables for Connecting to the 40 Gbps Physical Ports

Although you can create up to eight 10 GbE connections for each gateway switch, its physical ports are 40 Gbps quad small form-factor pluggable (QSFP) ports. Each gateway switch has two of these physical ports reserved for 10 GbE connectivity. The ports are equipped with QSFP transceivers that take an optical cable with an MTP/MPO connector. On the site network side are typically 10 Gbps 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 10 Gbps SFP+ transceiver takes a duplex LC connector for a total of four pairs. Thus, you can use a single splitter cable to create up to four 10 GbE 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 10 GbE network from the same manufacturer as the 10 GbE switch.

What About Data Centers Without a 10 GbE Infrastructure?

If your data center does not have a 10 GbE infrastructure, you can still connect Oracle Big Data Appliance to it by using an external switch that supports both 10 GbE (or 40 GbE) and 1 GbE. The Sun Network 10 GbE Switch 72p and numerous third-party switches provide this capability. You can connect the Sun Network QDR InfiniBand

Gateway Switches to a 10 GbE or 40 GbE port and connect the data center to a 1 GbE port in the external switch.

Configuring Multiple Racks into a Single Hadoop Cluster

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

Table 3–1 (Cont.) Default IP Addresses for Oracle Big Data Appliance

Host	Administrative IP Addresses	Oracle ILOM IP Addresses	InfiniBand Bonded IP Addresses
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 you configure the network.

Table 3–2 Oracle Big Data Appliance Port Numbers

Port	Used by
22	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 to 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

Table 3–3 (Cont.) CDH Port Numbers

Port	Component	Service	Access
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
4181	ZooKeeper	Server	Peer
8001	Hue	Job Submission Server	Internal
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

Table 3–3 (Cont.) CDH Port Numbers

Port	Component	Service	Access
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 the Oracle Big Data Appliance Configuration Utility. It contains the following sections:

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

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

- [Using OpenOffice Spreadsheet](#)
- [Generating the Configuration Files](#)
- [About the Configuration Files](#)
- [Checking for Network Errors](#)

Using OpenOffice Spreadsheet

You must use Apache OpenOffice Spreadsheet 3.4 or Oracle Open Office Calc 3.3 to configure the spreadsheet. Do not use other spreadsheet programs like Microsoft Office 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 Information Center: Oracle Big Data Appliance (ID 1445762.2).

To open the Oracle Big Data Appliance Configuration Utility:

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

Generating the Configuration Files

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

To generate the configuration files:

1. Open the spreadsheet with OpenOffice 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 Installation Template to verify that the configuration is correct.

About the Configuration Files

The Oracle Big Data Appliance Configuration Utility generates the following files to use when you configure 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. Transfer this file to a USB drive for copying to Oracle Big Data Appliance. See "["Configuring the Network"](#) on page 8-20.

mammoth-rack_name.params

Contains all information provided in the spreadsheet, including the network configuration, port numbers, default user names, and passwords. Transfer this file to a USB drive for copying to Oracle Big Data Appliance. See "["Using the Mammoth Utility"](#) on page 11-1.

bda-install-template.pdf

Reproduces the Installation Preview page of the spreadsheet in a printable format. The network administrator and others can use this Installation Template 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

Cluster Name	bda	Primary Appliance	Yes
Rack Name	bda1		
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	

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

The network administrator must add the host names and IP addresses for the components of Oracle Big Data Appliance to a name server. When this task is complete, you can use the `preinstall-checkip.sh` script to check the network for errors.

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 that 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. This step ensures that the network has not changed and 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 that 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 8-20. The procedures take longer to complete because of time-outs from attempts to connect to configured services such as Domain Name System (DNS) and Network Time Protocol (NTP).

Example 4-1 shows success and error messages from the preinstall-checkip.sh script.

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 to **Table 4-5** briefly describe the fields in the Network Configuration sheet of the Oracle Big Data Appliance Configuration Utility. See the *Oracle Big Data Appliance Configuration Worksheets* for complete details.

Table 4-1 describes the general rack properties.

Table 4-1 General Properties for the Rack

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 bda has racks named bda1, bda2, and so forth.
Server Base Name	Base name for all servers. A two-digit suffix uniquely identifies each server. The rack name and server base name are used to generate the host names for all network interfaces: eth0, bondib0, bondeth0, and Oracle ILOM. For example, a cluster name of bda and a server base name of node results in host names of bda1node01, bda1node02, and so forth.
Admin Access Suffix	Suffix to the basic host name to form the bondeth0 host names
Private Name Suffix	Suffix to the basic host name to form the bondib0 host name
ILOM Name Suffix	Suffix to the basic host name to form the Oracle ILOM name

Table 4–1 (Cont.) General Properties for the Rack

Spreadsheet Field	Description
Switch Base Name	Suffix to the rack name to form the base name for all switches. For example, a rack name of bda and a switch base name of sw results in switch names of bda02sw-ip, bda02sw-ib1, and so forth.
Customer Name	Name of the enterprise. Required.
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, and 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, and 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 for the Rack

Spreadsheet Field	Description
DNS Servers	Up to four IP addresses for the Domain Name System 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 example.com and us.example.com

Table 4–3 describes the network properties for individual servers in the rack. You specify the IP addresses assigned to node01. 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 Network Properties for Servers

Spreadsheet Field	Description
Administrative - eth0	First IP address of 18 from the administrative 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 administrative network

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

Table 4–4 Network Properties for Switches

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](#).

Table 4–5 Network Properties for Power Distribution Units

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](#) to [Table 4–9](#) briefly describe the fields in the Software Configuration sheet of the Oracle Big Data Appliance Configuration Utility.

[Table 4–6](#) describes the optional software available for Oracle Big Data Appliance. If you have a license for Oracle Big Data Connectors, then you can install the software. 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 terabytes (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 <i>nfs-host:/remote/shared/path</i> , where <i>nfs-host</i> is the fully qualified host name and <i>/remote/shared/path</i> is the directory path. The <i>hdfs</i> user must have write permission on this directory. See the <i>Oracle Big Data Appliance Configuration Worksheets</i> 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 10](#) 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 the network file system (NFS) protocol.

Table 4–8 Users, Groups, and Passwords

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

Table 4–9 (Cont.) Cloudera Manager Email Alerts

Spreadsheet Field	Description
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 -- from unpacking Oracle Big Data Appliance to powering on the system. This chapter contains these sections:

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

Reviewing Safety Guidelines

Before Oracle Big Data Appliance arrives, review the following safety precautions to ensure that 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 provides proper grounding. For 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. 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 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 that 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.

WARNING: 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 that the phases are balanced on both sides, A and B, in case of a failover.

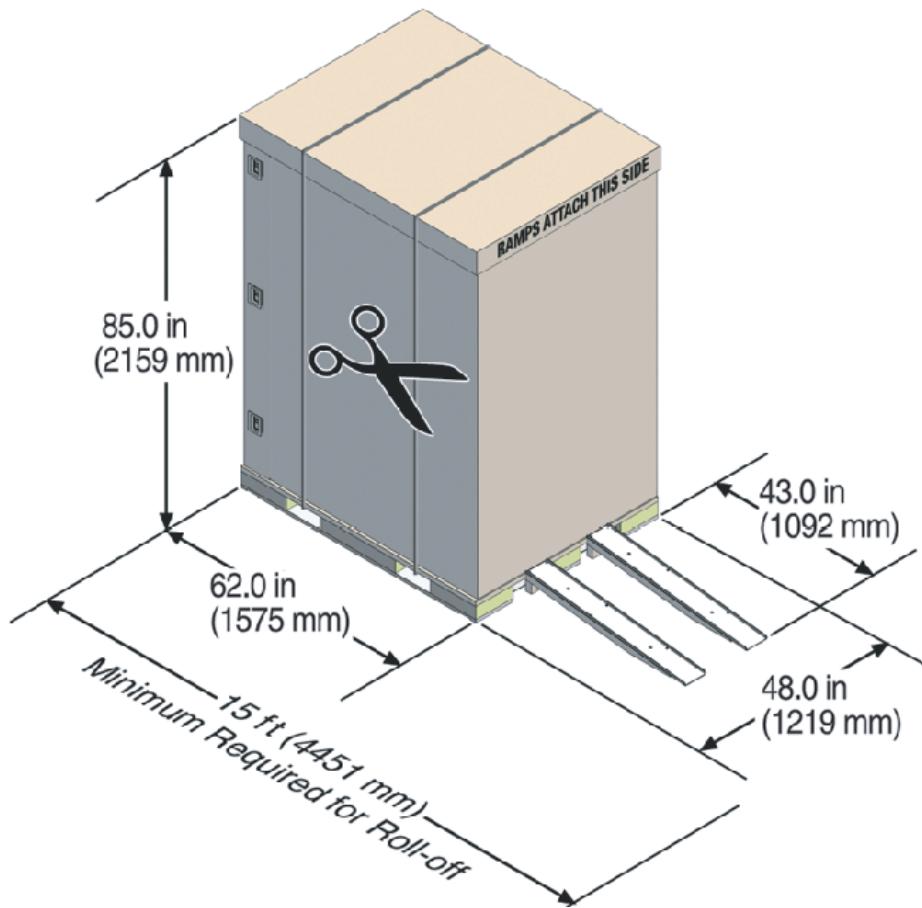
See Also:

- *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](#).

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

Figure 5–1 Oracle Big Data Appliance in the Shipping Crate

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

<http://docs.oracle.com/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 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

Removing Oracle Big Data Appliance from the Shipping Crate

Note: Oracle strongly recommends that you use professional movers to unpack and install Oracle Big Data Appliance.

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

To unpack Oracle Big Data Appliance:

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: one person on each side 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.

Note: 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

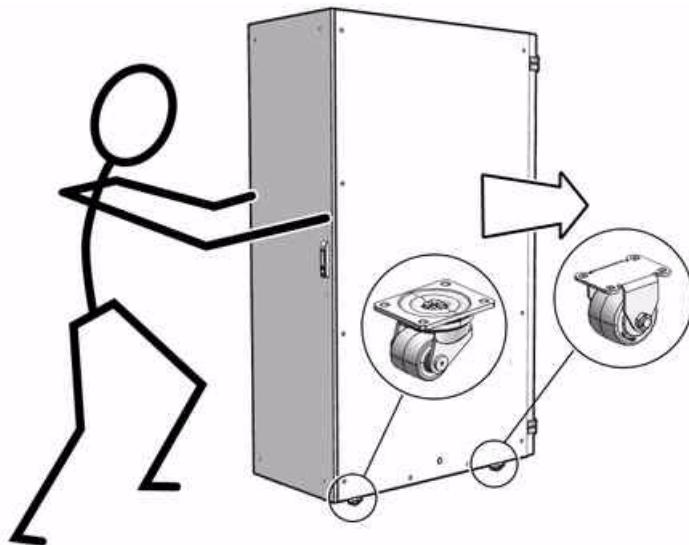
To move Oracle Big Data Appliance:

1. Ensure that the doors are closed and secured.
2. Ensure that 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, as shown in [Figure 5–2](#). You can safely maneuver Oracle Big Data Appliance by carefully pushing it.

[Figure 5–2](#) shows the correct way to maneuver Oracle Big Data Appliance.

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, move them slowly -- 0.65 meters (2 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.

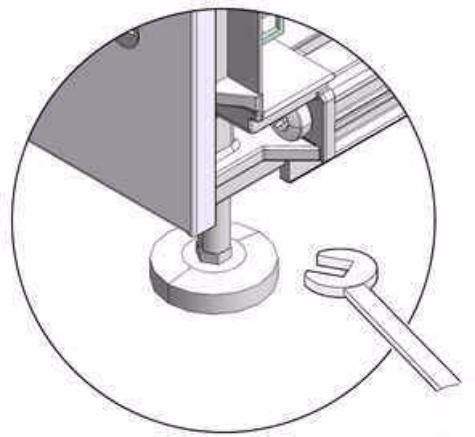
Stabilize the Oracle Big Data Appliance Rack with Leveling Feet

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

To adjust the leveling feet:

1. Locate the four leveling feet located at the bottom corners of the 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 the rack.

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



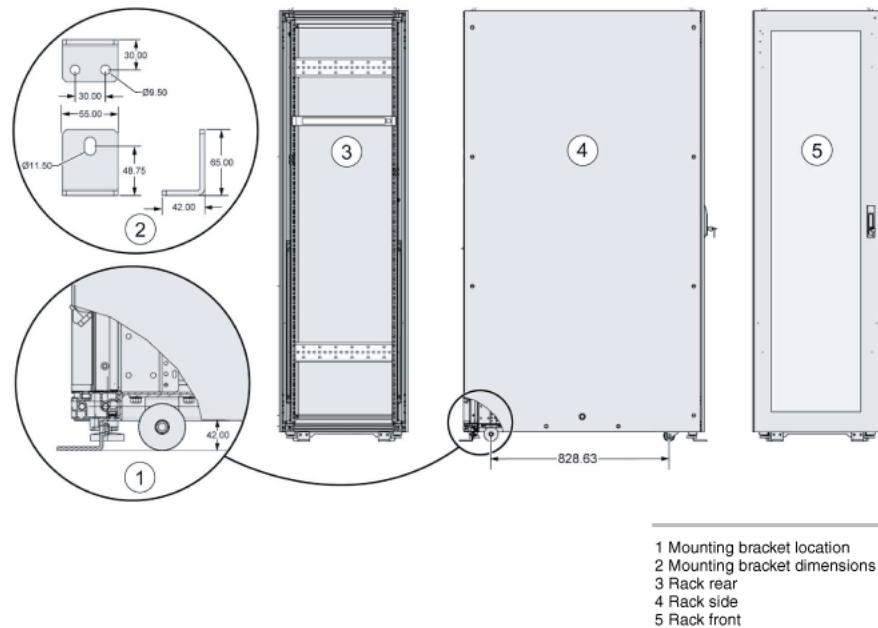
Stabilize the Oracle Big Data Appliance Rack with Mounting Brackets

You can permanently mount the rack to the installation site floor using the same four mounting brackets that secured the rack to the shipping pallet.

To install the mounting brackets:

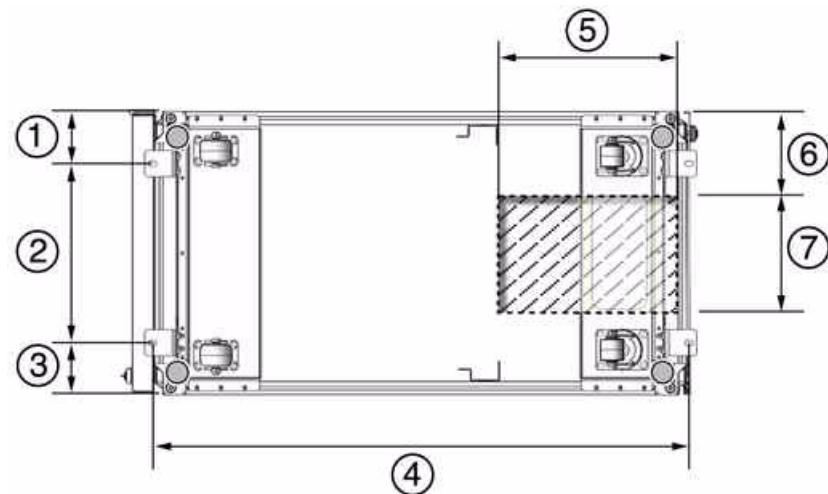
1. Drill the mounting holes in the installation floor. [Figure 5–4](#) shows the location and dimensions of the mounting brackets.

Figure 5–4 Location of Mounting Brackets on the Rack



2. Obtain four bolts and washers to mount the 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 the rack over the predrilled holes. [Figure 5–5](#) shows the bottom view of the Oracle Big Data Appliance rack and the location for the mounting hole and floor cut dimensions.

Figure 5–5 Bottom View of the Oracle Big Data Appliance Rack



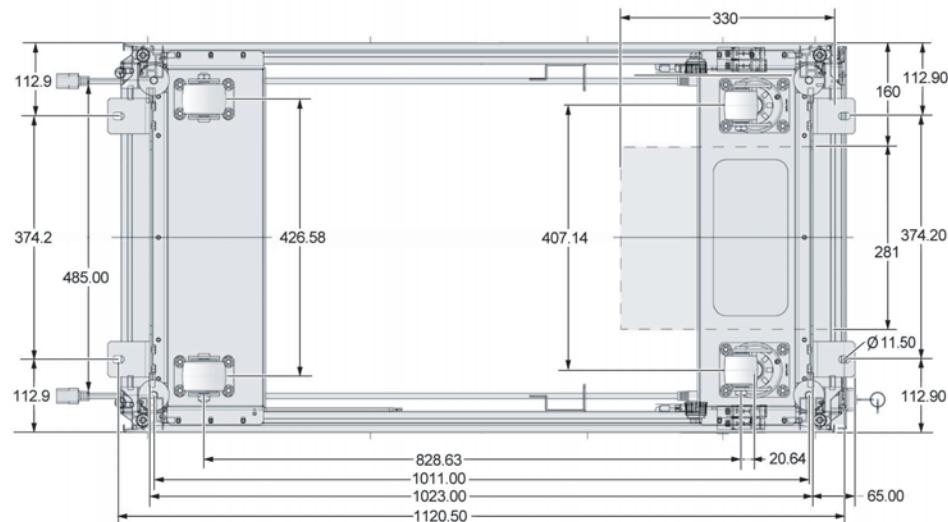
- 1: Distance from the center of the mounting bracket to the edge of the rack is 113 mm (4.45 inches).
- 2: Width between the centers of the mounting hole slots is 374 mm (14.72 inches).
- 3: Distance from the center of the mounting bracket to the edge of the rack is 113 mm (4.45 inches).
- 4: Distance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

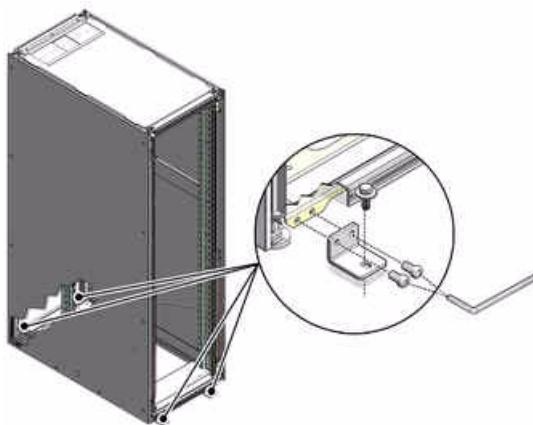
Caution: Do not create a hole where the rack casters or leveling feet brackets will be placed.

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

Figure 5–6 Base Position of the Rack to Measure (in mm)



4. Open the front and rear doors.
5. Attach the mounting brackets to the rack as shown in [Figure 5-7](#) using a 6 mm hexagon Allen wrench key.

Figure 5-7 Securing the 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 shown in [Figure 5-7](#).
7. Firmly tighten all of the bolts that secure the mounting brackets to the rack and to the floor.

Attaching a Ground Cable (Optional)

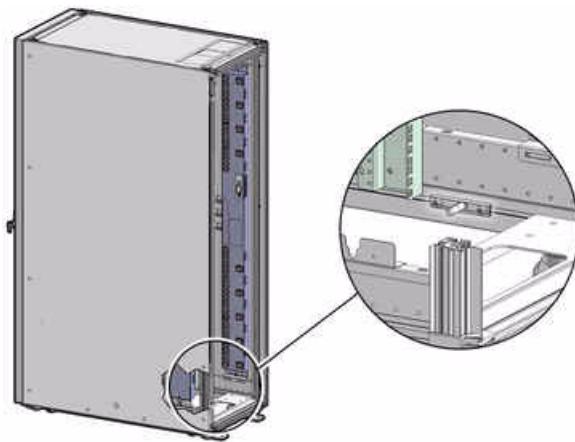
The Oracle Big Data Appliance rack power distribution units (PDUs) achieve earth ground through their power cords. Final chassis ground is achieved by way of the ground prong when you connect the power cord to a socket. For additional grounding, attach a chassis earth ground cable to the rack. The additional 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 that the installation site has properly grounded the power source in the data center. The facility PDU must have earth ground.
3. Ensure that 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.

Figure 5–8 Earth Ground Attachment Bolt Location



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

Inspecting the Oracle Big Data Appliance 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 the Oracle Big Data Appliance rack:

1. Check the rack for damage.
2. Check the rack for loose or missing screws.
3. Check the rack for the ordered configuration. Refer to the Customer Information Sheet (CIS) on the side of the packaging.
4. Check that all cable connections are secure and firmly in place 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 the 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 the Oracle Big Data Appliance rack.

To connect power cords to the 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 through the top or the bottom of the rack.
5. Secure the power cords in bundles.
6. Ensure that the breaker switches are OFF before connecting the power cables.
7. Plug the PDU power cord connectors into the facility receptacles.

Figure 5-9 Location of Power Cord Cable Ties

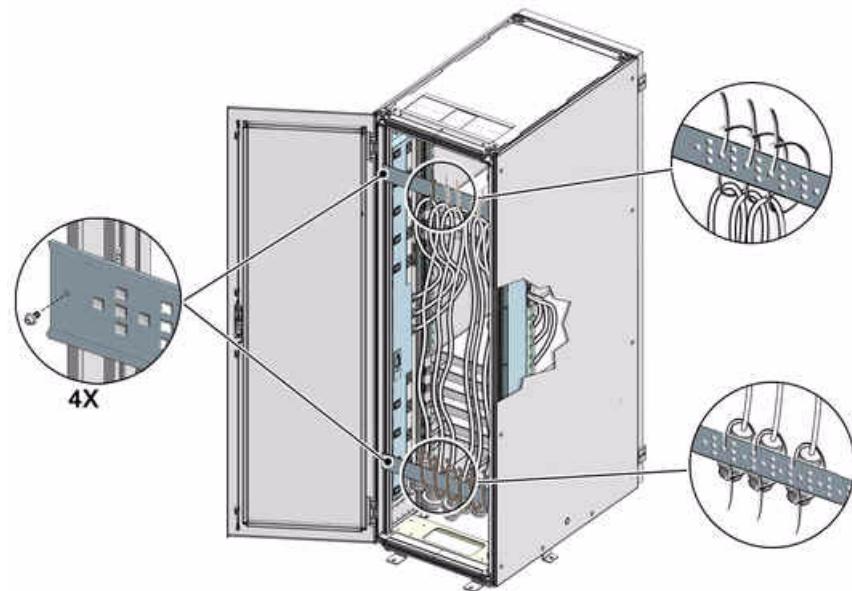


Figure 5–10 Power Cord Routing Out the Bottom

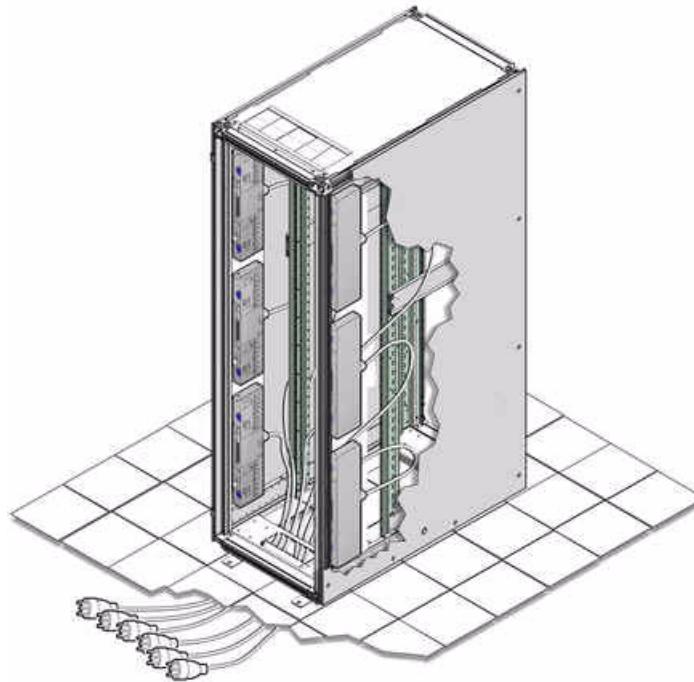
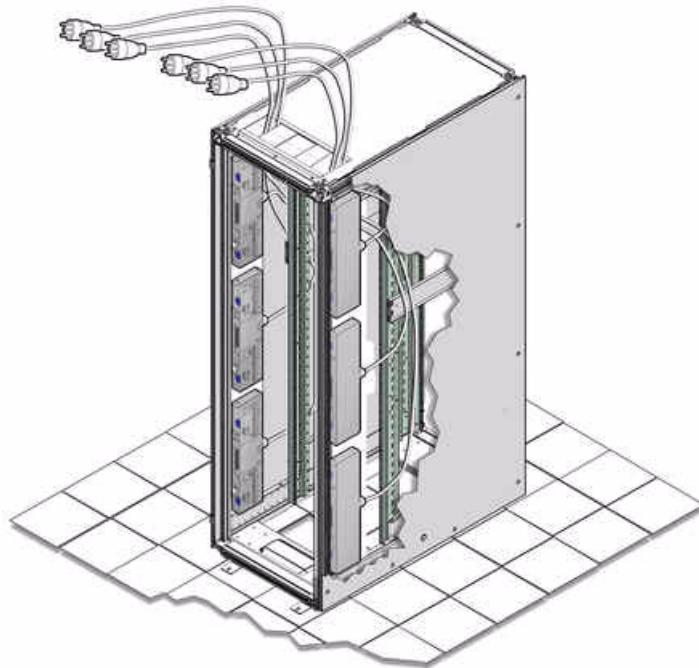


Figure 5–11 Power Cord Routing Out the Top



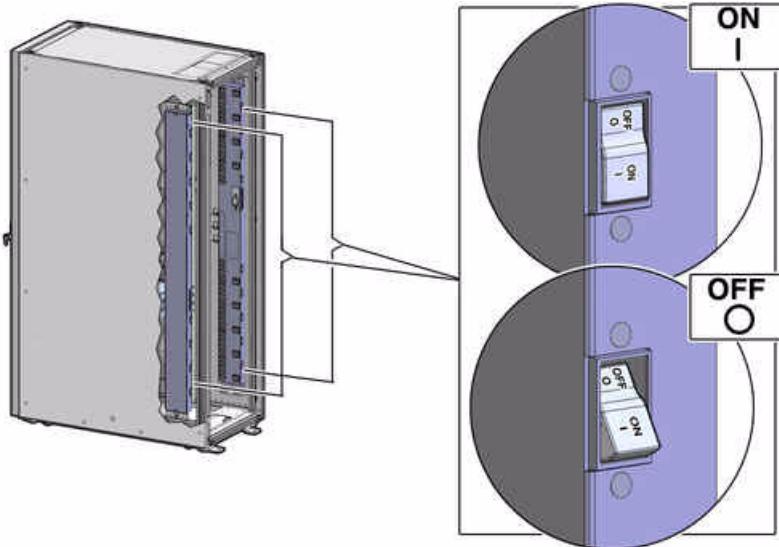
Powering On Oracle Big Data Appliance

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 rack as shown in [Figure 5-12](#). Press the ON (|) 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 that the expected power light-emitting diodes (LEDs) are on. The LEDs are located as follows:
 - 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 that 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 that 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 that 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 that the OK LED is blinking standby for all systems. The OK LED blinks every 3 seconds when in standby mode. It is dark for 2 to 3 minutes while Oracle ILOM is booting before going 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.

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

Using Oracle Integrated Lights Out Manager

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

This chapter contains the following sections:

- [Oracle ILOM Overview](#)
- [Administrative Network Diagram](#)
- [Oracle ILOM IP Addresses for Oracle Big Data Appliance Components](#)
- [Connecting to Oracle ILOM Using the Network](#)
- [Connecting to Oracle 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 an Oracle Big Data Appliance rack. Oracle ILOM dedicated hardware and software is preinstalled on these components. It 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 Oracle 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 nongraphical 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 Oracle 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 Oracle ILOM from the server operating system

(Oracle Linux). Using Oracle ILOM, you can remotely manage Oracle Big Data Appliance as if you were using a 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 to use a browser to log in to the SP, and then to perform system management and monitoring.

Command-Line Interface

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

Oracle ILOM Users

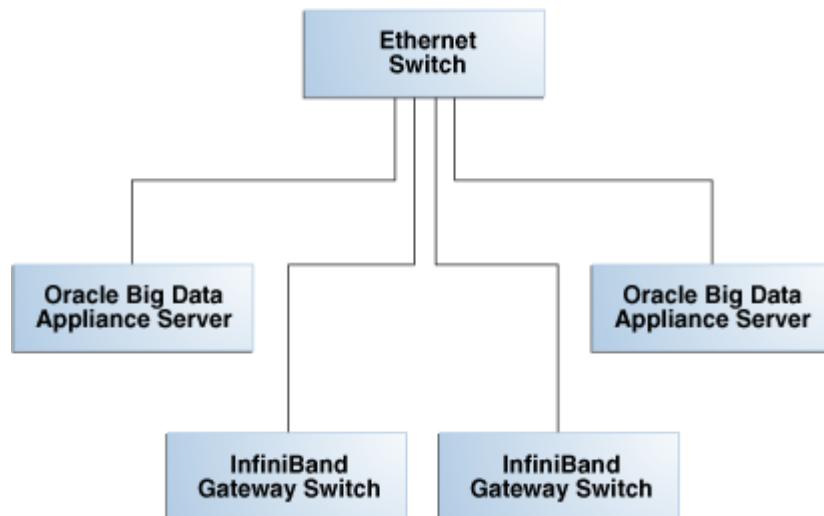
Oracle ILOM on Oracle Big Data Appliance is configured with these users:

- `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 administrative Oracle 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



Oracle ILOM IP Addresses for Oracle Big Data Appliance Components

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

After you reconfigure these IP addresses using the Oracle Big Data Appliance Configuration Utility, you must use the reconfigured IP addresses to access Oracle ILOM.

Connecting to Oracle ILOM Using the Network

You can typically access Oracle ILOM using the network over an Ethernet connection. You must know the Oracle 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 the IP address listed in the Installation Template. You can use either the CLI or the browser interface to access Oracle ILOM. Alternatively, you can launch a remote KVM session.

Note: You can use this connection method when Oracle ILOM IP addresses can be accessed over the network. Oracle recommends that you use this connection method. See [Chapter 8](#)

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

To connect to the Oracle ILOM command-line interface:

1. Open a secure shell (SSH) client, such as PuTTY.
2. In the **Host Name (or IP address)** field, enter the Oracle 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 bdalsw-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.

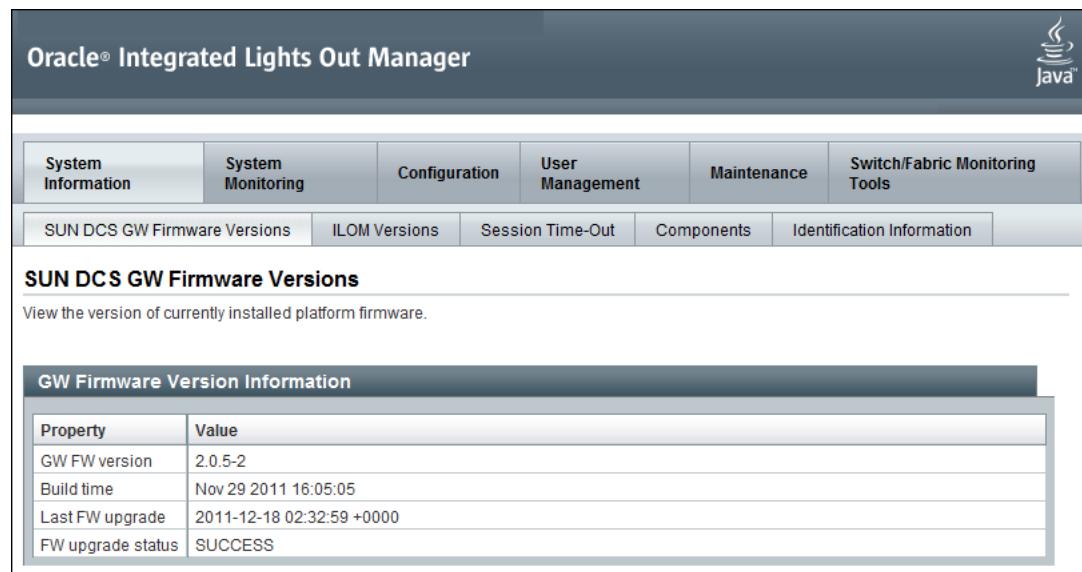
See Also: *Oracle ILOM Daily Management -- CLI Procedures Guide* at

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

Connecting to the Web GUI

To connect to the Oracle ILOM web interface:

1. Open a web browser and navigate to the Oracle ILOM name or the 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 ILOM Web Interface. The top navigation bar includes links for System Information, System Monitoring, Configuration, User Management, Maintenance, and Switch/Fabric Monitoring Tools. Below this is a sub-navigation bar with links 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 sub-section titled "GW Firmware Version Information". This section displays 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

See Also: *Oracle ILOM 3.0 Daily Management -- Web Procedures Guide*
at
<http://docs.oracle.com/cd/E19860-01/>

Opening a Remote KVM Session

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

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

Connecting to Oracle ILOM Using a Serial Connection

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

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

Connecting to the Oracle ILOM of a Server

To connect to Oracle ILOM of a 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 power is supplied to either PSU, then Oracle ILOM is functional regardless of the power state of the servers.

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

The Oracle ILOM login prompt is displayed.

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

The Oracle ILOM default command prompt is displayed.

Connecting to the Oracle ILOM of a Gateway Switch

To connect to Oracle 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 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](#).

This chapter contains the following sections:

- [Overview of the dcli Utility](#)
- [dcli Syntax](#)
- [dcli Return Values](#)
- [dcli 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 scripts 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](#) for details about these options.

To set up passwordless SSH for root:

1. Connect to an 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 basic options to the dcli command. For a complete list of options, see [Table 7-1](#).

Getting Help

To see the dcli help page, enter the dcli command with the -h or --help options. You can see a description of the commands by entering 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](#).

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 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 a command file for a series of commands that you often use together or for commands with complex syntax. See the -x option in [Table 7-1](#).

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. The syntax in these examples executes the date command on all 18 servers.

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

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

The minimal output returns OK for completed execution:

```
# dcli -n date
OK: ['bdalnode07.example.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: dcservers
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.vmstatOps None
options.execfile: None
argv: ['/opt/oracle/bda/bin/dcli', '-g', 'dcservers', '-v', 'date']
Success connecting to nodes: ['bdalnode07.example.com']
...entering thread for bdalnode07.example.com:
execute: /usr/bin/ssh -l root bdalnode07.example.com ' date'
...exiting thread for bdalnode07.example.com status: 0
bdalnode07.example.com: Tue Feb 14 10:24:43 PST 2012]
```

dcli Syntax

`dcli [options] [command]`

Parameters

options

The 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 quotation marks.

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 for information about command files.

- § (dollar sign)
- ' (quotation mark)
- < (less than)
- > (greater than)
- () (parentheses)

Table 7-1 dcli Options

Option	Description
<code>-c nodes</code>	Specifies a comma-separated list of Oracle Big Data Appliance servers where the command is executed
<code>-C</code>	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 .
<code>-d destfile</code>	Specifies a target directory or file name for the <code>-f</code> option
<code>-f file</code>	Specifies files to be copied to the user's home directory on the target servers. The files are not executed. See the <code>-l</code> option.
<code>-g groupfile</code>	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.
<code>-h, --help</code>	Displays a description of the commands
<code>-k</code>	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.
<code>-l userid</code>	Identifies the user ID for logging in to another server. The default ID is <code>root</code> .
<code>--maxlines=maxlines</code>	Identifies the maximum lines of output displayed from a command executed on multiple servers. The default is 10,000 lines.
<code>-n</code>	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 <code>-n</code> and <code>-r</code> options together.
<code>-r regexp</code>	Replaces the output with the server name for lines that match the specified regular expression
<code>-s sshoptions</code>	Specifies a string of options that are passed to SSH
<code>--scp=scpoptions</code>	Specifies a string of options that are passed to Secure Copy (SCP), when these options are different from <code>sshoptions</code>
<code>--serial</code>	Serializes execution over the servers. The default is parallel execution.
<code>-t</code>	Lists the target servers
<code>--unkey</code>	Drops the keys from the <code>authorized_key</code> files of the target servers
<code>-v</code>	Displays the verbose version of all messages
<code>--version</code>	Displays the dcli version number

Table 7-1 (Cont.) dcli Options

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

dcli Return Values

- 0: The command ran successfully on all servers.
- 1: One or more servers were inaccessible or remote execution returned a nonzero value. A message lists 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.

dcli Examples

Following are examples of the dcli utility.

This example returns the default list of target servers:

```
# dcli -t
Target nodes: ['bda1node01-adm.example.com', 'bda1node02-adm.example.com',
'bda1node03-adm.example.com', 'bda1node04-adm.example.com',
'bda1node05-adm.example.com', 'bda1node06-adm.example.com',
'bda1node07-adm.example.com', 'bda1node08-adm.example.com',
'bda1node09-adm.example.com']
```

The next example checks the temperature of all servers:

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

bda1node01-adm.example.com: value = 22.000 degree C
bda1node02-adm.example.com: value = 22.000 degree C
bda1node03-adm.example.com: value = 22.000 degree C
bda1node04-adm.example.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 (Oracle ILOM) and the `dcli` utility.

This chapter contains the following sections:

- [Configuring the KVM Switch](#)
- [Configuring the Cisco Ethernet Switch](#)
- [Configuring the InfiniBand Leaf and Spine Switches](#)
- [Configuring the Power Distribution Units](#)
- [Configuring the 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

The KVM configuration consists of these procedures:

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

- [Accessing a Server by Using the KVM Switch](#)

Starting the KVM Switch

To start 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. You see the main user interface page.
5. In the navigator on the left, select **Target Devices** under Unit View. In the main display area, verify that 18 target devices are listed with Action set to KVM Session.

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

6. If 18 sessions are not shown:
 - a. In the navigator under Appliances, expand Ports, and then select **IQ Adaptors**.
 - b. In the main display area, choose the **Port** table heading to sort the sessions by port number.
 - c. Note any missing sessions, so that you can fix them later.
 - d. In the navigator, choose **Target Devices** to return to the Target Devices page.

Connecting the KVM Switch to the Management Network

To connect the KVM switch to the management network:

1. In the navigator under User Accounts, select **Local**.
2. Under Users, choose **Admin**.
3. Set the password for the Admin account to `welcome1`, and then choose **Save**. Do not modify any other parameters.
4. Under Appliance Settings, expand Network, and then choose **IPv4**. The Network Information page appears.
5. Enter values for Address, Subnet, and Gateway, and then choose **Save**.
6. Under Appliance Settings, choose **DNS** to display the DNS Information page.
7. Enter the IP addresses of the DNS servers, and then choose **Save**.
8. Under Network, choose **General** to display the Appliance General Network Settings page.
9. Connect the KVM LAN1 Ethernet port to the management network.
10. To verify that the port has been configured correctly, ensure that the Media Access Control (MAC) address on the Network Settings page matches the label next to the LAN1/LAN2 ports at the rear of the KVM switch.
11. Under Users, select **Overview** to display the Unit Maintenance page.
12. Enter a name for the KVM switch, and then choose **Save**.

13. To reboot the KVM switch, choose **Reboot** under Overview and **Yes** to confirm.

Checking the KVM Firmware Version

You may need to upgrade the KVM firmware to the recommended version.

To check the KVM firmware version:

1. In the navigator under Appliance Settings, select **Versions**. There are two version numbers, Application and Boot. Compare the displayed versions with these recommended versions:

- Application: 1.10.2.17762
- Boot: 1.9.16473

If the application firmware version is earlier than 1.10.2, then you should upgrade it. 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. Under Appliance, select **Overview**.
6. From the Tools list, select **Upgrade Firmware**.
7. Select the connection method, such as **FTP** or **HTTP**.
8. Enter the file name of the downloaded firmware.
9. Click **Upgrade**.

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

10. Confirm the firmware version by selecting **Versions** under Appliance Settings.

Configuring the KVM Switch to Access the Servers

To configure the KVM switch to access the servers:

1. Under Unit View, select **Target Devices** to display the Target Devices page.
2. Power on the server. The power button is on the front panel.
3. Click the server name in the Name column to display the Unit Overview page.
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. Click **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 by Using the KVM Switch

To access a server by using the KVM switch:

1. Under Unit View, select **Target Devices** to display the Target Devices page.

2. Click the system name in the Name column.
3. Click **KVM Session** to open a session with the server.

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.

The Cisco Ethernet switch configuration consists of these topics and procedures:

- [Scope of the Configuration](#)
- [Prerequisites for Configuring the Ethernet Switch](#)
- [Configuring the Ethernet Switch on the Customer Network](#)
- [Setting Up Telnet Access](#)
- [Identifying the DNS Servers](#)
- [Setting the Clock and Time Zone](#)
- [Configuring the NTP Servers](#)
- [Disabling the Spanning Tree](#)
- [Verifying the Ethernet Configuration](#)

Scope of the Configuration

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 for Configuring the Ethernet Switch

To avoid disrupting the customer network, observe these prerequisites:

- Do not connect the Cisco Ethernet switch until the network administrator has verified the running configuration and 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 Ethernet Switch on the Customer Network

To configure the Ethernet switch on the customer network:

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 that the terminal session is recorded on the laptop by logging the output. The output can be used as a record 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, and 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...
Compressed configuration from 3600 bytes to 1305 bytes[OK]]
```

6. If the network requires IP routing on the switch, and 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 *bdal1sw-ip*.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# hostname bdal1sw-ip
bdal1sw-ip(config)# end
*Jan 23 15:57:50.886: %SYS-5-CONFIG_I: Configured from console by console
bdal1sw-ip# write memory
Building configuration...
Compressed configuration from 3604 bytes to 1308 bytes[OK]
bdal1sw-ip#
```

The system host name appears in the prompt.

Setting Up Telnet Access

Telnet access is optional. The following procedure describes how to enable and disable remote telnet access.

Note: Oracle Big Data Appliance ships with a version of the Cisco Ethernet switch software that supports telnet but not SSH. To obtain support for SSH, you must install an update. See My Oracle Support Information Center: Oracle Big Data Appliance (ID 1445762.2).

To set up telnet access to the Ethernet switch:

1. Set the password for telnet access if necessary; it should already be set when you receive Oracle Big Data Appliance.

```
bdal1sw-ip # configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdal1sw-ip(config)# enable password welcome1
bdal1sw-ip(config)# enable secret welcome1
The enable secret you have chosen is the same as your enable password.
This is not recommended. Re-enter the enable secret.
bdal1sw-ip(config)# end
bdal1sw-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 that 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.
bdal1sw-ip(config)# line vty 0 15
bdal1sw-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
...
```

```

bdal1sw-ip(config-line)# password welcome1
bdal1sw-ip(config-line)# login
bdal1sw-ip(config-line)# end
bdal1sw-ip# write memory
*Jan 23 15:58:53.630: %SYS-5-CONFIG_I: Configured from console by console
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.
bdal1sw-ip(config)# line vty 0 15
bdal1sw-ip(config-line)# no password
bdal1sw-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
...
bdal1sw-ip(config-line)# end
bdal1sw-ip# write memory
*Jan 23 15:58:53.630: %SYS-5-CONFIG_I: Configured from console by console
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:

```

bdal1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdal1sw-ip(config)# ip domain-name us.example.com
bdal1sw-ip(config)# ip name-server 10.7.7.3
bdal1sw-ip(config)# ip name-server 172.28.5.5
bdal1sw-ip(config)# ip name-server 10.8.160.1
bdal1sw-ip(config)# end
*Jan 23 16:01:35.010: %SYS-5-CONFIG_I:Configured from console by console
bdal1sw-ip# write memory
Building configuration...
Compressed configuration from 3662 bytes to 1348 bytes[OK]

```

Setting the Clock and Time Zone

The Cisco Ethernet switch keeps internal time in coordinated universal time (UTC) format.

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

```

bdal1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdal1sw-ip(config)# clock timezone EST -5
bdal1sw-ip(config)# clock summer-time EDT recurring
bdal1sw-ip(config)# end
bdal1sw-ip# clock set 15:00:00 January 23 2012
bdal1sw-ip# write memory
Building configuration...
Compressed configuration from 3778 bytes to 1433 bytes[OK]
bdal1sw-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, enter this command:

```
no clock timezone global configuration
```

- To use a time zone:

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

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

- To set summer time hours:

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

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

- To manually set the clock to any time:

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

In this command, *hh:mm:ss* is the hour, month, and second 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.

```
bdal1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdal1sw-ip(config)# ntp server 10.196.23.254 prefer
bdal1sw-ip(config)# ntp server 192.168.9.19
bdal1sw-ip(config)# end
Jan 23 20:00:41.235: %SYS-5-CONFIG_I:Configured from console by console
bdal1sw-ip# write memory
Building configuration...
Compressed configuration from 3870 bytes to 1487 bytes [OK ]
bdal1sw-ip# show ntp status
output varies by network
bdal1sw-ip# show clock
15:00:57.919 EST Mon Jan 23 2012
```

Disabling the Spanning Tree

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

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

To disable the spanning tree:

1. If the port must be disabled, then enter these commands:

```
bdal1sw-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
bdal1sw-ip# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes[OK]
```

2. To verify the disabling of the spanning tree:

```
bdal1sw-ip# show spanning-tree vlan 1
Spanning tree instance(s) for vlan 1 does not exist.
```

Verifying the Ethernet Configuration

To verify the Cisco Ethernet switch configuration:

1. Verify the configuration by entering the following command:

```
bdal1sw-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, enter these commands:

```
bdal1sw-ip# no ip default-gateway 10.7.7.1
bdal1sw-ip# end
bdal1sw-ip# write memory
```

2. Save the current configuration by entering this command:

```
bdal1sw-ip# copy running-config startup-config
```

3. Exit from the session with this command:

```
bdal1sw-ip#exit
```

```
bd1sw-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.

Caution: Do not connect the Cisco Ethernet switch to the management network until after the system is configured with the customer's IP addresses and the switch configuration is complete.

Configuring the InfiniBand Leaf and Spine 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 the switches, follow these procedures for each one:

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

Configuring an InfiniBand Switch

To configure an InfiniBand switch:

1. Connect to the switch using a serial or an Ethernet connection.
 - For an Ethernet connection to 192.168.1.201, 192.168.1.202, or 192.168.1.203, see ["Connecting to Oracle ILOM Using the Network" on page 6-3](#).
 - For a serial connection, see ["Connecting to Oracle ILOM Using a Serial Connection" on page 6-5](#).
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. Enter these commands to configure the switch:

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

In these commands, *ip_address*, *ip_netmask*, and *ip_gateway* represent the appropriate settings on your network.

5. Enter a show command to view the changes. If any values are wrong, reenter 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 bda1sw-ib2. Do not include the domain name.

```

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

```

```

/SP
Properties:
  hostname = bda1sw-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

```

In these commands, *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 and Clock on an InfiniBand Switch

To set 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 in 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 Server 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 usenntpserver=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 Oracle Big Data Appliance server 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

To check the health of an InfiniBand leaf or spine 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 Information Center: Oracle Big Data Appliance (ID 1445762.2) for the current version.

```
FabMan@bdal1sw-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@bdal1sw-02->
```

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

```
FabMan@bdal1sw-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@bdal1sw-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@bdal1sw-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@bdal1sw-ib2-> disablesm
```

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@bdal1sw-ib2-> setsmpriority 5
```

c. Restart the InfiniBand Subnet Manager:

```
FabMan@bdal1sw-ib2-> enablesm
```

7. If you are connecting this Oracle Big Data Appliance rack to an Oracle Exadata Database Machine or an Oracle Exalogic rack:

a. Verify that the Exadata InfiniBand switches and the Exalogic spine switch are running firmware version 1.3.3_2 or later.

b. Ensure that the subnet manager runs only on the switches with the highest firmware version.

c. 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.

For example, if Oracle Big Data Appliance has the highest firmware version, then make its spine switch the master and its gateway switches the failover. Then, on the other engineered system, disable the subnet manager on any InfiniBand switch that has a lower firmware version than the version on Oracle Big Data Appliance.

8. Exit the Fabric Management shell:

```
FabMan@bdal1sw-ib2-> exit
->
```

9. Exit the Oracle ILOM shell:

```
-> exit
```

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

```
reboot
```

11. Repeat these steps for the other InfiniBand switches.

Configuring the Power Distribution Units

The power distribution unit (PDU) configuration consists of these procedures:

- [Connecting the PDUs to the Network](#)
- [Verifying the PDU Firmware Version](#)
- [Configuring the Threshold Settings for the PDUs](#)

Connecting the PDUs to the Network

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

- Static IP address
- Subnet mask
- Default gateway

To connect the PDUs 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 that 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.
7. Repeat Steps 5 and 6 for the second PDU.

Verifying the PDU Firmware Version

To verify 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**, and then click **Logout**.

Configuring the Threshold Settings for the PDUs

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 the 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 the 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 the 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 the 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

Configuring the Oracle Big Data Appliance Servers

Before configuring the network, ensure that the Oracle Big Data Appliance servers are set up correctly.

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 Oracle ILOM serial management port, then the baud rate setting on the servers changes from the default Oracle ILOM setting of 9600 to 115200 baud, 8 bits, no parity, and 1 stop bit.

To check the Oracle Big Data Appliance servers:

1. Power on all servers by pressing the power button on the front panel of each server. The servers take 5 to 10 minutes to run through the normal startup tests.
2. To configure the KVM switch 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 C-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 switch or a laptop:
 - KVM switch: Under Unit View, select **Target Devices**. Left-click the system name, and 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 entering 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 that 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
```

 Enter 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 Gbps (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:


```
bda1node02-adm.example.com: Connected. Use ^D to exit.
bda1node02-adm.example.com: -> show faulty
bda1node02-adm.example.com: Target      | Property      | Value
bda1node02-adm.example.com:-----+-----+-----+
bda1node02-adm.example.com:
bda1node02-adm.example.com: -> Session closed
```

```
bda1node02-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 that 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 that the status is OK for both power supplies:

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

- To verify that 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 that 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 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 bdachecksw > 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 ["Generating the Configuration Files" on page 4-2](#) if you do not have this file.

The network configuration consists of these procedures:

- [Verifying the Factory Software Image](#)
- [Copying the Configuration Files to Oracle Big Data Appliance](#)
- [Starting the Network Configuration](#)
- [Connecting to the Network](#)
- [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` files are in the `/root` directory of each server. If you see a `BDA_IMAGING_FAILED` or `BDA_REBOOT_FAILED` file in the output, then check the `/root/bda_imaging_status` file on that server for more information. Do not proceed with network configuration until all problems are resolved.

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

```
# dccli ls -1 /root | grep BDA
IP address BDA_IMAGING_SUCCEEDED
IP address BDA_REBOOT_SUCCEEDED
.
.
.
```

You can also confirm the image version:

```
# dccli 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 to Oracle Big Data Appliance:

1. Copy the configuration files to a USB flash drive.
2. Use the KVM switch to open a console session to the first server. The first server is the lowest server in the rack. See [Figure C-1](#).
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. Enter 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
```

Note: If `mammoth-rack_name.params` is also on the drive, you can copy it to `/opt/oracle/BDAMammoth` for use in [Chapter 11](#).

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 Oracle ILOM names for all servers and configures the administrative network and the private InfiniBand network.

To start the network configuration:

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

.
.

bdalnode02.example.com BdaUserConfigNetwork: reset network
bdalnode03.example.com BdaUserConfigNetwork: reset network
bdalnode04.example.com BdaUserConfigNetwork: reset network

.
.

networksetup-one: deploying this server
networksetup-one: network will restart momentarily, pardon our dust
bdalnode01.example.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.example.com
hello from bda1node03.example.com
hello from bda1node04.example.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 1 GbE administrative network by connecting the Cisco Ethernet switch to the data center.
- Connect the 10 GbE 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 10 GbE 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 10 GbE 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 8-23](#).

2. Run the following 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.example.com
hello from bda1node03.example.com
hello from bda1node04.example.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 bdalnode02.example.com
hello from bdalnode03.example.com

.

.

networksetup-two: passed
networksetup-two: check existence of default vlan for port 0A-ETH-1 on bdalsw-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 bdalsw-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 bdalnode02 eth9 on switch bdalsw-ib2
check and delete vNIC for bdalnode02 eth9 on switch bdalsw-ib3
create vNIC eth9 bdalnode02 using switch bdalsw-ib3
vNIC created
check and delete vNIC for bdalnode02 eth8 on switch bdalsw-ib2

.

.

networksetup-two: ping server ips on client network
networksetup-two: passed
networksetup-two: test ssh server ips on client network
hello from bdalnode02.example.com
hello from bdalnode03.example.com
```

```

.
.
networksetup-two: passed
networksetup-two: end setup client network

```

Reinstalling the Base Image

The operating system and various utilities are factory installed on Oracle Big Data Appliance, 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 zip file with the correct version of the base image and copy it to `node01` (bottom server). See [My Oracle Support Information Center: Oracle Big Data Appliance \(ID 1445762.2\)](#) for the download location. You can copy the file to any directory, such as `/tmp`.

The name of the file is in the format `BDABaseImage-version.zip` (for example, `BDABaseImage-1.1.0.zip`).

Note: You can also download `BDABaseImage-version.zip` 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. Locate the appropriate configuration file and verify that it reflects the intended network configuration. Edit the file and copy it to `/opt/oracle/bda/` as needed.
 - To reimagine to the custom network settings, copy `BdaDeploy.json`.
 - To reimagine to the factory default network settings, copy `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 4 GB 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 zip file, for example:

```
# unzip BDABaseImage-1.1.0.zip
$ unzip BDABaseImage-1.1.0.zip
Archive: BDABaseImage-1.1.0.zip
  inflating: README.txt
  creating: BDABaseImage-1.1.0/
  inflating: BDABaseImage-1.1.0/BDABaseImage-1.1.0.iso
  inflating: BDABaseImage-1.1.0/makebdimage
  extracting: BDABaseImage-1.1.0/BDABaseImage-1.1.0.md5sum
  inflating: BDABaseImage-1.1.0/reimagerack
```

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

```
cd BDABaseImage-1.1.0
```

9. Complete one of the following procedures:

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

```
./reimagerack
```

- To reimagine an appliance that still has the factory settings:

- Ensure that /opt/oracle/bda/BdaDeploy.json does not exist.
- Enter the ./reimagerack command.

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

- Copy /opt/oracle/bda/BdaDeploy.json to a safe location outside Oracle Big Data Appliance.

- Disconnect the rack from the network.

- Reimage the rack:

```
./reimagerack deploy ship
```

The reimagerack utility creates an ISO image, 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 11](#).

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.

- [bdacheckcluster](#)
- [bdacheckhw](#)
- [bdacheckib](#)
- [bdachecknet](#)
- [bdachecksw](#)

- [bdadiag](#)
- [bdaid](#)
- [bdaimagevalidate](#)
- [bdaredoclientnet](#)
- [bdaserials](#)
- [iblinkinfo](#)
- [imagehistory](#)
- [imageinfo](#)
- [listlinkup](#)
- [showvlan](#)
- [showvnics](#)

bdacheckcluster

Checks the health of a CDH cluster, including the software, hardware, and network.

This example shows the output from the utility:

```
# bdacheckcluster
SUCCESS: Mammoth configuration file is valid.
SUCCESS: All cluster host names are pingable
SUCCESS: All cluster hosts passed checks on last reboot
INFO: Starting cluster host hardware checks
SUCCESS: All cluster hosts pass hardware checks
INFO: Starting cluster host software checks
SUCCESS: All cluster hosts pass software checks
SUCCESS: All ILOM hosts are pingable
SUCCESS: All client interface IPs are pingable
SUCCESS: All admin eth0 interface IPs are pingable
SUCCESS: All private Infiniband interface IPs are pingable
Warning: Permanently added 'bda1node01-master' (RSA) to the list of known hosts.
SUCCESS: Puppet master is running on bda1node01-master
SUCCESS: Puppet running on all cluster hosts
SUCCESS: Cloudera SCM server is running on bda1node02
SUCCESS: Cloudera SCM agent running on all cluster hosts
SUCCESS: Name Node is running on bda1node01
SUCCESS: Secondary Name Node is running on bda1node02
SUCCESS: Job Tracker is running on bda1node01
SUCCESS: Data Nodes running on all cluster hosts
SUCCESS: Task Trackers running on all cluster slave hosts
SUCCESS: Hadoop filesystem is healthy.
SUCCESS: MySQL server is running on MySQL master node bda1node03
SUCCESS: MySQL server is running on MySQL backup node bda1node02
SUCCESS: Big Data Appliance cluster health checks succeeded
```

bdacheckhw

Checks the hardware profile of the server. See ["Configuring the Oracle Big Data Appliance Servers"](#) on page 8-17 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
.
.
.
.
```

bdacheckib

Checks the InfiniBand cabling between the servers and switches of a single rack, when entered 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 bdacheckib command has these options:

-s

The 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]# bdacheckib
LINK bda1sw-ib3.15A ... bdainode02.HCA-1.2 UP
LINK bda1sw-ib3.15B ... bdainode01.HCA-1.2 UP
LINK bda1sw-ib3.14A ... bdainode04.HCA-1.2 UP
LINK bda1sw-ib3.14B ... bdainode03.HCA-1.2 UP
.
.
.
```

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

```
[root@bdainode01 bda]# bdacheckib -g
[root@bdainode01 bda]# cat sample-multi-rack.json
# This json multirack spec is generated. The array elements are sorted
# alphabetically. A properly 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": "bdainode01", "LEAF1_NAME": "bdainode02", "LEAF2_NAME": "bdainode03"},
  {"SPINE_NAME": "bdainode04", "LEAF1_NAME": "bdainode05", "LEAF2_NAME": "bdainode06"}]
```

The final 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: bdaisw-ib2
    LINK ... to rack1 UP
    LINK ... to rack1 UP
    .
    .
    .
```

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 ["Configuring the Oracle Big Data Appliance Servers" on page 8-17](#) 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](#), and then [bdachecksw](#).

bdaredoclientnet

Re-creates the virtual NICs (VNICs) 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 the following subset of tasks done by the `networksetup-two` script 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

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:

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

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

3. Remove passwordless SSH:

```
/opt/oracle/bda/bin/remove-root-ssh
```

See "[Setting Up Passwordless SSH](#)" on page 7-1 for more information about this command.

4. Change directories:

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

5. Run the utility:

```
bdaredoclientnet
```

The output is similar to that shown in [Example 8-2](#).

6. Restore passwordless SSH (optional):

```
/opt/oracle/bda/bin/setup-root-ssh
```

bdaserials

Returns the serial numbers and media access control (MAC) addresses for most components of the Oracle Big Data Appliance server that you are connected to.

This example shows the output from the utility:

```
# bdaserials
Rack serial number :
System serial number : 1137FMM0BY
System UUID : 080020FF-FFFF-FFFF-7E97D6282100
Motherboard serial number : 0338MSL-1131BA2194
Chassis serial number : 1137FMM0BY
Memory serial numbers : 87948175 87949173 87948163 8794816B 87948130 87948176
Infiniband HCA serial number : 1388FMH-1122501437
Disk controller serial number : SV11713731
Hard disk serial numbers :
SEAGATE ST32000SSSUN2.0T061A1125L6M89X
SEAGATE ST32000SSSUN2.0T061A1125L6LFH0
SEAGATE ST32000SSSUN2.0T061A1125L6M94J
SEAGATE ST32000SSSUN2.0T061A1125L6LLEZ
SEAGATE ST32000SSSUN2.0T061A1125L6M5S2
SEAGATE ST32000SSSUN2.0T061A1125L6LSD4
SEAGATE ST32000SSSUN2.0T061A1127L6M58L
SEAGATE ST32000SSSUN2.0T061A1127L6R40S
SEAGATE ST32000SSSUN2.0T061A1125L6M3WX
SEAGATE ST32000SSSUN2.0T061A1125L6M65D
SEAGATE ST32000SSSUN2.0T061A1127L6NW3K
SEAGATE ST32000SSSUN2.0T061A1127L6N4G1

MAC addresses :
bondeth0 Ethernet : CE:1B:4B:85:2A:63
```

```

bondib0 InfiniBand : 80:00:00:4A:FE:80:00:00:00:00:00:00:00:00:00:00:00:00:00
bond0 Ethernet : 00:00:00:00:00:00
eth0 Ethernet : 00:21:28:E7:97:7E
eth1 Ethernet : 00:21:28:E7:97:7F
eth2 Ethernet : 00:21:28:E7:97:80
eth3 Ethernet : 00:21:28:E7:97:81
eth8 Ethernet : CE:1B:4B:85:2A:63
eth9 Ethernet : CE:1B:4C:85:2A:63
ib0 InfiniBand : 80:00:00:4A:FE:80:00:00:00:00:00:00:00:00:00:00:00:00:00
ib1 InfiniBand : 80:00:00:4B:FE:80:00:00:00:00:00:00:00:00:00:00:00:00

```

iblinkinfo

Lists all InfiniBand connections in the InfiniBand network. Run this command as `root` from any server.

This example shows two Oracle Big Data Appliances and one Oracle Exadata Database Machine on the InfiniBand network:

```

[root@bda1node01 network]# iblinkinfo
Switch 0x002128df348ac0a0 SUN IB QDR GW switch bdalsw-ib2 10.133.43.36:
 149 1[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 130 2[ ] "SUN IB QDR GW switch bdalsw-ib2 10.133...
 149 2[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 127 1[ ] "SUN IB QDR GW switch bdalsw-ib2 10.133...
 149 3[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 111 2[ ] "SUN IB QDR GW switch bdalsw-ib2 10.133...
 149 4[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 109 1[ ] "SUN IB QDR GW switch bdalsw-ib2 10.133...
 149 5[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 143 1[ ] "bda1node02 BDA 192.168.41.20 HCA-1" ( )
 149 6[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 137 1[ ] "bda1node01 BDA 192.168.41.19 HCA-1" ( )
 149 7[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 141 1[ ] "bda1node04 BDA 192.168.41.22 HCA-1" ( )
 149 8[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 123 1[ ] "bda1node03 BDA 192.168.41.21 HCA-1" ( )
 149 9[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 151 1[ ] "bda1node06 BDA 192.168.41.24 HCA-1" ( )
 149 10[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 112 1[ ] "bda1node05 BDA 192.168.41.23 HCA-1" ( )
 149 11[ ] ==( 4X 10.0 Gbps Active/ LinkUp)==> 139 1[ ] "bda1node07 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)==> [ ] "" ( )

```

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

```

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.el5uek
BDA_RPM_VERSION           : bda-1.0.3-1
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

```

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 to combine multiple Oracle Big Data Appliance racks into one super cluster. It contains the following sections:

- [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 you are 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 must not have problems related to InfiniBand restarts.
- The new racks must be configured with the appropriate IP addresses to be migrated into the expanded system prior to any cabling, and duplicate IP addresses are not allowed.
- The existing spine switch is set to priority 10 during the cabling procedure. This setting gives the spine switch a higher priority than any other switch in the network 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 C](#) 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.

To cable two racks together:

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 the Subnet Manager.
 - e. Use the `setsmpriority 10` command to set the priority to 10.
 - f. Use the `enablesm` command to restart the Subnet Manager.
 - g. Repeat Step b to ensure that the Subnet Manager Master is running on the spine switch.
2. Ensure that the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
3. Completely shut down the new rack (R2).
4. Cable the leaf switch in the new rack according to [Table D-2](#).
5. Power off leaf switch R1 IB2. This causes all servers to fail over their InfiniBand traffic to R1 IB3.
6. Disconnect all interswitch links between R1 IB2 and R1 IB3.
7. Cable leaf switch R1 IB2 according to [Table D-1](#).
8. Power on leaf switch R1 IB2.
9. Wait for 3 minutes for R1 IB2 to become completely operational.
To check the switch, log in to it and run the `ibswitches` command. The output should show three switches: R1 IB1, R1 IB2, and R1 IB3.
10. Power off leaf switch R1 IB3. This causes all servers to fail over their InfiniBand traffic to R1 IB2.
11. Cable leaf switch R1 IB3 according to [Table D-1](#).
12. Power on leaf switch R1 IB3.
13. Wait for 3 minutes for R1 IB3 to become completely operational.

To check the switch, log in to it and run the `ibswitches` command. The output should show three switches: R1 IB1, R1 IB2, and R1 IB3.

14. Power on all the InfiniBand switches in R2.
15. Wait for 3 minutes for the switches to become completely operational.

To check the switch, log in to it and run the `ibswitches` command. The output should show six switches: R1 IB1, R1 IB2, R1 IB3, R2 IB1, R2 IB2, and R2 IB3.

16. Ensure that 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 `disablesm` command to stop the Subnet Manager.
 - b. Use the `setsmpriority 8` command to set the priority to 8.
 - c. Use the `enablesm` command to restart the Subnet Manager.
19. Ensure that the 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,... R n , the new rack is R $n+1$, and the Subnet Manager Master is running on R1 IB1.

To cable several racks together:

1. Set the priority of the current, active Subnet Manager Master to 10 on the spine switch, 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 the Subnet Manager.
 - e. Use the `setsmpriority 10` command to set the priority to 10.
 - f. Use the `enablesm` command to restart the Subnet Manager.
 - g. Repeat Step b to ensure that the Subnet Manager Master is running on the spine switch.
2. Ensure that the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
3. Completely shut down the new rack (R $n+1$).
4. Cable the leaf switch in the new rack according to the appropriate table in [Appendix D](#). For example, if rack R $n+1$ was R4, then use [Table D-9](#).
5. Complete the following procedure for each of the original racks. In these steps, Rx represents a rack number from R1 to R n .

- a. Power off leaf switch Rx IB2. This causes all servers to fail over their InfiniBand traffic to Rx IB3.
- b. Cable leaf switch Rx IB2 according to [Appendix D](#).
- c. Power on leaf switch Rx IB2.
- d. Wait for 3 minutes for Rx IB2 to become completely operational.

To check the switch, log in to it and run the `ibswitches` command. The output should show $n*3$ switches for IB1, IB2, and IB3 in racks R1, R2, ..., R n .

- e. Power off leaf switch Rx IB3. This causes all servers to fail over their InfiniBand traffic to Rx IB2.
- f. Cable leaf switch Rx IB3 according to [Appendix D](#).
- g. Power on leaf switch Rx IB3.
- h. Wait for 3 minutes for Rx IB3 to become completely operational.

To check the switch, log in to the switch and enter the `ibswitches` command. The output should show $n*3$ switches for IB1, IB2, and IB3 in racks R1, R2, ..., R n .

All racks should now be rewired according to [Appendix D](#).

6. Power on all the InfiniBand switches in R $n+1$.
7. Wait for 3 minutes for the switches to become completely operational.

To check the switch, log in to the switch and run the `ibswitches` command. The output should show $(n+1)*3$ switches for IB1, IB2, and IB3 in racks R1, R2, ..., R $n+1$.

8. Ensure that the Subnet Manager Master is running on R1 IB1 by entering the `getmaster` command from any switch.
9. Power on all servers in R $n+1$.
10. Log in to spine switch R1 IB1, and lower its priority to 8 as follows:
 - a. Enter the `disablesm` command to stop the Subnet Manager.
 - b. Enter the `setsmpriority 8` command to set the priority to 8.
 - c. Enter the `enablesm` command to restart the Subnet Manager.
11. Ensure that the Subnet Manager Master is running on one of the spine switches by entering the `getmaster` command from any switch.
12. Ensure that the Subnet Manager is running on every spine switch by entering the following command from any switch:
`ibdiagnet -r`
Each spine switch should show as running in the Summary Fabric SM-state-priority section of the output. If a spine switch is not running, then log in to the switch and enable the Subnet Manager by entering the `enablesm` command.
13. If there are now four or more racks, then log in to the leaf switches in each rack and disable the Subnet Manager by entering the `disablesm` command.

Setting Up Auto Service Request

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

- [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. 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 you configure 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 detects faults in the most common server components, such as disks, fans, and power supplies, and automatically opens a service request when a fault occurs. ASR monitors only server components and does not detect all possible faults.

ASR is not a replacement for other monitoring mechanisms, such as SMTP and SNMP alerts, within the customer data center. It is a complementary mechanism that

expedites and simplifies the delivery of replacement hardware. ASR should not be used for downtime events in high-priority systems. For high-priority events, contact Oracle Support Services directly.

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/systems/asr/overview/index.html>

Oracle Auto Service Request user documentation at

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

Getting Ready for ASR Installation

Before installing ASR, complete these prerequisites:

1. Create a My Oracle Support account at <http://support.oracle.com>.
2. Ensure that the following are set up correctly:
 - Oracle Premier Support for Systems, Oracle Premier Support for Operating Systems, or Oracle Limited Warranty
 - The technical contact person at the customer site who is responsible for Oracle Big Data Appliance
 - A valid shipping address at the customer site for Oracle Big Data Appliance 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/systems/asr/overview/hardware-recommendations-330108.html>
5. Confirm that Java Development Kit 6 (JDK 1.6.0_04 or later) is running on the designated ASR Manager system:

```
java -version
```

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

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

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

See Also: *Oracle Auto Service Request Installation and Operations Guide*
at

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

Installing ASR Manager

This section contains these topics:

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

Downloading the Software

ASR Manager requires installation of these components:

- ASR Manager Version 3.6 or later
- Oracle Service Tags

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

<http://www.oracle.com/technetwork/systems/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 Oracle Automated Service Manager Package

To install the Oracle Automated Service Manager (OASM) package:

1. Connect to the ASR Manager system.
2. Check whether OASM 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 10-3.

3. As the `root` user, install the OASM 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/SUNWswasr/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. Verify that ASR is running before entering any other ASR commands:

```
/opt/SUNWswasr/bin/sasm status
```

Registering ASR Manager

To register ASR Manager:

1. As the root user on the ASR Manager server, enter 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 user name and password when prompted. ASR validates the login. After validation, the registration is complete. Passwords are not stored.

Your My Oracle Support email address receives output from ASR reports, notification of ASR problems, and notice when a service request (SR) has been generated.

Verifying ASR Manager

Perform these checks on ASR Manager to ensure that 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 to trap component fault indicators. See ["Mammoth Utility Steps"](#) on page 11-6.

Activating ASR Assets

To activate ASR assets:

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

```
asr list_asset -i asset_ip
```

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

```
asr list_asset
```

The following is an example of the output. You should see a host name and an Oracle ILOM name for each of the 18 servers, for a total of 36 activations. If you see fewer than 36, then identify the missing ones and activate them. An example of a server name is bda1node15 and an Oracle 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, then verify that all steps of the Mammoth Utility have run successfully. See [Chapter 11](#).

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, then 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
```

In this command, *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 Oracle ILOM:

```
asr activate_asset -i ilom_ip
```

In this command, *ilom_ip* is the Oracle ILOM IP address.

4. Activate the servers:

```
asr activate_asset -i host_eth0_ip
```

In this command, *host_eth0_ip* is the IP address of the first server on the administrative network.

5. Repeat these `activate_asset` commands on each 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 Oracle ILOM:

1. Log in to the Oracle ILOM of any server in Oracle Big Data Appliance.
2. Change to the following directory:

```
cd /SP/alertmgmt/rules/3
```

3. Enter 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 that 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 support identifier (CSI) for Oracle Big Data Appliance.

To set up a test in the operating system environment:

1. Enter 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 support identifier (CSI) for Oracle Big Data Appliance.

This test event does not open a case, but it makes a verifiable entry in the ASR back end.

You received the emails from both of the previous tests (Oracle ILOM and operating system environment), then your validation is complete. If not, then proceed to ["Troubleshooting ASR"](#) on page 10-7.

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 Chapter 5 of the *Oracle ASR Installation and Operations Guide* at:

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

If you continue to have issues, contact ASR support. 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 sections:

- [Using the Mammoth Utility](#)
- [Installing the Software on a Single or Primary Rack](#)
- [Upgrading the Software on Oracle Big Data Appliance](#)
- [Mammoth Utility Syntax](#)
- [What If an Error Occurs During the Installation?](#)
- [Adding a Rack to an Existing Cluster](#)
- [Mammoth Utility Steps](#)
- [Using the Mammoth Reconfiguration Utility](#)
- [Mammoth Reconfiguration Utility Syntax](#)

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 11-2](#).
- 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 11-2](#) 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 11-2](#).
- For the other racks of the cluster, follow the procedure in ["Adding a Rack to an Existing Cluster" on page 11-4](#).

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 8-20](#).
2. Verify that the software is not installed on the rack already. If the software is installed and you want to reinstall it, then use the `mammoth -p` option in Step 9.
3. Download the BDAMammoth zip file to any directory on node01 (such as `/tmp`). See My Oracle Support Information Center: Oracle Big Data Appliance (ID 1445762.2) for the download location.
4. Log in to node01 as root and extract all files from the downloaded zip file:

```
$ unzip p14479858_110_Linux-x86-64.zip
Archive: p14479858_110_Linux-x86-64.zip
  inflating: README.txt
  creating: BDAMammoth-1.1.0/
  inflating: BDAMammoth-1.1.0/bda-configuration-1.1.0.ods
  inflating: BDAMammoth-1.1.0/BDAMammoth-1.1.0.run
```

5. Change to the BDAMammoth-*version* directory:

```
$ cd BDAMammoth-1.1.0
```

6. Extract all files from BDAMammoth-*version*.run:

```
$ ./BDAMammoth-1.1.0.run
```

7. Change to the BDAMammoth directory.

```
$ cd /opt/oracle/BDAMammoth
```

8. Copy `mammoth-rack_name.params` to the current directory. See ["About the Configuration Files" on page 4-3](#).

9. Run the `mammoth` command with the appropriate options. See [Table 11-1](#). 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.

Upgrading the Software on Oracle Big Data Appliance

The procedure for upgrading the software is the same whether you are upgrading from one major release to another or just applying a patch set. The procedure is also the same whether your Hadoop cluster consists of one Oracle Big Data Appliance rack or multiple racks.

The process upgrades all components of the software stack including the firmware, operating system, CDH, JDK, and Oracle Big Data Connectors (if previously installed).

Software downgrades are not supported.

Note: Because the upgrade process automatically stops and starts services as needed, the cluster is unavailable while the `mammoth` command is executing.

To upgrade the software:

1. Download the BDAMammoth zip file to any directory on node01 (such as `/tmp`). See My Oracle Support Information Center: Oracle Big Data Appliance (ID 1445762.2) for the download location.
2. Log in to node01 as `root` and extract all files from the downloaded zip file:

```
$ unzip p14479858_110_Linux-x86-64.zip
Archive: p14479858_110_Linux-x86-64.zip
  inflating: README.txt
  creating: BDAMammoth-1.1.0/
  inflating: BDAMammoth-1.1.0/bda-configuration-1.1.0.ods
  inflating: BDAMammoth-1.1.0/BDAMammoth-1.1.0.run
```

3. Change to the BDAMammoth-*version* directory:

```
$ cd BDAMammoth-1.1.0
```

4. Extract all files from BDAMammoth-*version*.run:

```
$ ./BDAMammoth-1.1.0.run
```

The new version of the Mammoth software is installed in `/opt/oracle/BDAMammoth`, and the previous version is saved in `/opt/oracle/BDAMammoth/previous-BDAMammoth`.

5. Change to the BDAMammoth directory.

```
$ cd /opt/oracle/BDAMammoth
```

6. Run the `mammoth` command with the `-p` option:

```
./mammoth -p rack_name
```

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]
```

In this command, `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.

You must finish installing one rack before starting the installation of another rack.

[Table 11-1](#) lists the Mammoth Utility options.

Table 11-1 Mammoth Utility Options

Option	Description
-h	Displays command Help including command usage and a list of steps.
-i	Runs all mandatory steps, equivalent to -r 1-18.
-l	List the steps of the Mammoth Utility.
-p	Upgrades the software on the cluster to the current version.
-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> .
-v	Displays the version number of the Mammoth Utility.

Example 11-1 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:

```
STEP-i-yyyymmddhhmmss.log
```

In this format, *i* is the step number and *yyyymmddhhmmss* 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.

To install the software on additional racks in 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 11-2.
2. Ensure that all racks are running the same software version. See "[About Software Version Differences](#)" on page 11-5.
3. Ensure that all racks that form a single Hadoop cluster are cabled together. See [Chapter 9](#).
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, enter the `mammoth` command with the appropriate option. See "[Mammoth Utility Syntax](#)" on page 11-3. 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, if you have a license for them, 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 upgrade the image version:

- Run the Mammoth Utility with the `-p` option to upgrade the software on the cluster to the latest version. See "[Upgrading the Software on Oracle Big Data Appliance](#)" on page 11-3.

To downgrade the image version:

- Reimage the new rack to the older version installed on the cluster. See My Oracle Support Information Center: Oracle Big Data Appliance (ID 1445762.2) 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 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 3 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 4 PreInstallChecks

This step performs a variety of hardware and software checks. A failure in any of these checks causes the Mammoth Utility to fail:

- The ARP cache querying time is 2 seconds or less.
- All server clocks are synchronized within 10 seconds of the current server.
- All servers succeeded on the last restart and generated a /root/BDA_REBOOT_SUCCEEDED file.
- The bdacheckhw utility succeeds.
- The bdachecksw utility succeeds.

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 11-4.

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 and corresponds to device /dev/lvg1/lv1.

After this step finishes, the Linux file systems table in /etc/fstab shows the logical disk instead of /u12, or /u11 and /u12.

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 EXTFILE and be owned by root.
- The /scratch/bda/bda1/nn and /scratch/bda/bda1/snn directories must exist on EXTFILE and be owned by hdfs in group hadoop.

See Also: *Oracle Big Data Appliance Configuration Worksheets* for detailed information about external NameNode backups.

Step 12 SetupMySQL

Installs and configures MySQL Database. This step creates the primary database and several databases on node02 for use by Cloudera Manager. It also sets up replication of the primary database to a backup database on node03.

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`

In this example, 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](#).

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 10](#).

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 10-5](#).

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 2](#).

Using the Mammoth Reconfiguration Utility

The following is the general procedure for running the Mammoth Reconfiguration Utility.

To run the Mammoth Reconfiguration Utility:

1. Log into the HDFS node (node01) of the primary rack and change to the BDAMammoth directory:

```
cd /opt/oracle/BDAMammoth
```

Note: If the HDFS node is in failure, then log in to the noncritical node that you want to reconfigure as the new HDFS node.

2. Enter the `mammoth-reconfig` command with the appropriate subcommand option. See "[Mammoth Reconfiguration Utility Syntax](#)" on page 11-10.

Specific procedures for performing common software installations are provided in this topic:

- [Changing the Software Configuration](#)

Changing the Software Configuration

During the initial configuration of Oracle Big Data Appliance, the optional software components may or may not be installed. Using the Mammoth Reconfiguration Utility, you can reverse those decisions. In this release, you can turn Auto Service Request on.

To turn on Auto Service Request:

1. Set up your My Oracle Support account and install ASR Manager. See [Chapter 10](#).
2. Log into the HDFS node (node01) of the primary rack and change to the BDAMammoth directory:

```
cd /opt/oracle/BDAMammoth
```

3. Turn on Auto Service Request monitoring and activate the assets:

```
mammoth-reconfig add asr
```

Mammoth Reconfiguration Utility Syntax

The `mammoth-reconfig` command has this basic syntax:

```
mammoth-reconfig option parameter
```

This utility uses the configuration settings stored in /opt/oracle/bda/install/state/mammoth-saved.params. It prompts for any missing

information, such as passwords. When the utility makes a change, it modifies this file to reflect the new configuration.

Options

add

Adds a service to the cluster. The *parameter* is a keyword that identifies the service:

- asr: Turns on Auto Service Request monitoring on Oracle Big Data Appliance and activates assets on ASR Manager. The installation process prompts you for the ASR Manager host name, port number, and root password. See [Chapter 10](#) for more information about Auto Service Request.

This example adds Auto Service Request support to all servers in the cluster:

```
mammoth-reconfig add asr
```

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 sections:

- [Monitoring the Ambient Temperature of Servers](#)
- [Powering On and Off Oracle Big Data Appliance](#)
- [Adding Memory to a Server](#)
- [Maintaining the Physical Disks of Servers](#)
- [Replacing a Server Disk](#)
- [Changing InfiniBand IP Addresses](#)
- [Maintaining the InfiniBand Network](#)
- [Changing the NTP Servers](#)

See Also: [Chapter 7, "Using the dcli Utility"](#)

Monitoring the 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. 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 entering 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](#).

The following is an example of the command output:

```
bda1node01-adm.example.com: value = 22.000 degree C
bda1node02-adm.example.com: value = 22.000 degree C
bda1node03-adm.example.com: value = 22.000 degree C
```

```
bdalnode04-adm.example.com: value = 23.000 degree C
.
.
```

Powering On and Off Oracle Big Data Appliance

This section includes the following topics:

- [Nonemergency Power Procedures](#)
- [Emergency Power-Off Considerations](#)
- [Cautions and Warnings](#)

Nonemergency Power Procedures

This section contains the procedures for powering on and off the components of Oracle Big Data Appliance in an orderly fashion.

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 Oracle ILOM interface and applying power to the system.

To power on Oracle Big Data Appliance:

1. Turn on all 12 breakers on both PDUs.
Allow 1 to 2 minutes for Oracle ILOM to start.
2. Power up the servers.

Powering On Servers Remotely Using Oracle ILOM

You can power on the servers remotely using the Oracle ILOM interface. You can access Oracle ILOM using the web console, the command-line interface (CLI), the intelligent platform management interface (IPMI), or the simple network management protocol interface (SNMP). For example, to apply power to server bdalnode01 using IPMI, run the following command as `root` from a server that has `ipmitool` installed:

```
ipmitool -H bdalnode01-c -U root chassis power on
```

In this example, `bdalnode01-c` is the host name of Oracle ILOM for the server to be powered on. You are prompted for the password.

See Also: [Chapter 6, "Using Oracle Integrated Lights Out Manager"](#)

Powering Off Oracle Big Data Appliance

To power off Oracle Big Data Appliance:

1. Power off 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. Enter 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:

```
dcli -l root -g group_name shutdown -hP now
```

In this command, `group_name` is a file that contains a list of servers.

The following example shuts down all Oracle Big Data Appliance servers listed in the `server_group` file:

```
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 5 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:

WARNING: Do not touch the parts of this product that use high-voltage power. Touching them might result in serious injury.

Caution:

- Do not power off Oracle Big Data Appliance unless there is an emergency. In that case, follow the "Emergency Power-Off Procedure" on page 12-3.
- 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 populate the 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 to specific nodes that need more memory, such as the NameNode (node01).

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. However, individual servers can be taken outside of the cluster temporarily and require downtime.

See Also: "Parts for Sun Fire Servers" on page B-4 for the repair procedures

Verifying the Server Configuration

The 12 disk drives in each Oracle Big Data Appliance 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 degradation or an outage. The effect on the server of validating the RAID devices is minimal. The corrective actions may affect operation of the server and can range from simple reconfiguration to an outage, depending on the specific issue uncovered.

Verifying Disk Controller Configuration

Enter 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 controllers 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

Enter this command to verify the virtual drive configuration:

```
# MegaCli64 -LDInfo -lAll -a0
```

The 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. The 12 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
.
.
.
```

Replacing a Server Disk

The failure of a disk is never catastrophic on Oracle Big Data Appliance. No user data should be lost. Data stored in HDFS or Oracle NoSQL Database is automatically replicated.

This section contains the following topics:

- [Overview of the Disk Replacement Process](#)
- [About Disk Drive Identifiers](#)
- [Dismounting Partitions Before Replacing a Working Disk](#)
- [What If a Server Fails to Restart?](#)
- [Replacing a Disk Drive](#)
- [Identifying the Function of a Disk Drive](#)
- [Configuring an Operating System Disk](#)
- [Configuring an Oracle NoSQL Database Disk](#)
- [Configuring an HDFS Disk](#)
- [Verifying the Disk Configuration](#)

Overview of the Disk Replacement Process

The following are the basic steps for replacing a server disk drive:

1. Replace the failed disk drive.
2. Perform the basic configuration steps for the new disk.
3. Identify the dedicated function of the failed disk, either as an HDFS disk, an operating system disk, or an Oracle NoSQL Database disk.
4. Configure the disk for its dedicated function.
5. Verify that the configuration is correct.

See Also:

"Servicing Customer-Replaceable Devices" in the *Sun Fire X4170 M2 and X4270M2 Servers Installation Guide* at

http://docs.oracle.com/cd/E19245-01/E21671/swap.html#50503714_61628

About Disk Drive Identifiers

The disk drives are configured as raid0 logical drives, with logical drive 0 recognized as disk /dev/sda by the operating system. The physical disk position (or slot number) does not necessarily correspond with the logical disk number.

Standard Disk Drive Mappings

Table 12-1 shows the mappings between the RAID logical drives and the operating system identifiers, and the dedicated function of each drive in an Oracle Big Data Appliance server. Nonetheless, you must confirm that these mappings are correct on your system.

Table 12–1 Disk Drive Identifiers

RAID Logical Drive	Operating System Location	Dedicated Function
0	/dev/sda	Operating system
1	/dev/sdb	Operating system
2	/dev/sdc	HDFS
3	/dev/sdd	HDFS
4	/dev/sde	HDFS
5	/dev/sdf	HDFS
6	/dev/sdg	HDFS
7	/dev/sdh	HDFS
8	/dev/sdi	HDFS
9	/dev/sdj	HDFS
10	/dev/sdk	HDFS or Oracle NoSQL Database
11	/dev/sdl	HDFS or Oracle NoSQL Database

Standard Mount Points

Table 12–2 show the mappings between HDFS partitions and mount points.

Table 12–2 Mount Points

HDFS Partition	Mount Point
/dev/sda4	/u01
/dev/sdb4	/u02
/dev/sdc1	/u03
/dev/sdd1	/u04
/dev/sde1	/u05
/dev/sdf1	/u06
/dev/sdg1	/u07
/dev/sdh1	/u08
/dev/sdi1	/u09
/dev/sdj1	/u10
/dev/sdk1	/u11
/dev/sdl1	/u12

Obtaining the Physical Slot Number of a Disk Drive

Use the following `MegaCli64` command to verify the mapping of virtual drive numbers to physical slot numbers. See "Replacing a Disk Drive" on page 12-10.

```
# MegaCli64 LdPdInfo a0 | more
```

Dismounting Partitions Before Replacing a Working Disk

If you plan to replace an HDFS disk or an operating system disk before it fails, then you should first dismount the HDFS partitions.

For operating system disks (disks 0 and 1), you must also inactivate the swap partition. Replacing an operating system disk before inactivating the swap partition triggers a restart of the server.

Note: Only dismount HDFS partitions. For an operating system disk, ensure that you do not dismount operating system partitions. Only partition 4 (sda4 or sdb4) of an operating system disk is used for HDFS.

To dismount the HDFS partitions of HDFS and operating system disks:

1. Log in to the server with the failed drive.
2. For operating system disks, inactivate the swap partition:

```
# swapoff
```

3. List the mounted HDFS partitions:

```
# mount -l
```

```
/dev/md2 on / type ext3 (rw,noatime)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
/dev/md0 on /boot type ext3 (rw)
tmpfs on /dev/shm type tmpfs (rw)
/dev/sda4 on /u01 type ext4 (rw,nodev,noatime) [/u01]
/dev/sdb4 on /u02 type ext4 (rw,nodev,noatime) [/u02]
/dev/sdc1 on /u03 type ext4 (rw,nodev,noatime) [/u03]
/dev/sdd1 on /u04 type ext4 (rw,nodev,noatime) [/u04]
```

```
.
```

```
.
```

4. Check the list of mounted partitions for the failed disk. If the disk has no partitions listed, then proceed to ["Replacing a Disk Drive"](#) on page 12-10. Otherwise, continue to the next step.

Caution: For operating system disks, look for partition 4 (sda4 or sdb4). Do not dismount an operating system partition.

5. Dismount the HDFS mount points for the failed disk:

```
# umount mountpoint
```

For example, `umount /u12` removes the mount point for disk `/dev/sd1`.

If the `umount` commands succeed, then proceed to ["Replacing a Disk Drive"](#) on page 12-10. If a `umount` command fails with a device busy message, then the partition is still in use. Continue to the next step.

6. Open a browser window to Cloudera Manager. For example:

```
http://bda1node02.example.com:7180
```

7. Complete these steps in Cloudera Manager:

- a. Log in as admin.

- b. On the Services page, click **hdfs1**
- c. Click the **Instances** subtab.
- d. In the Host column, locate the server with the failed disk. Then click the service in the Name column, such as **hdfs1-DATANODE-12**, to open its page.
- e. Click the **Configuration** subtab.
- f. Remove the mount point from the Directory field.
- g. Click **Save Changes**.
- h. From the Actions list, choose **Restart**.

Note: If you removed the mount point in Cloudera Manager, then you must restore the mount point in Cloudera Manager after finishing all other configuration procedures.

8. Return to your session on the server with the failed drive.
9. Reissue the `umount` command:

```
# umount mountpoint
```

10. Complete the steps in "Replacing a Disk Drive" on page 12-10.

What If a Server Fails to Restart?

The server may restart during the disk replacement procedures, either because you issued a reboot command or made an error in a MegaCli64 command. In most cases, the server restarts successfully, and you can continue working. However, in other cases, an error occurs so that you cannot reconnect using ssh. In this case, you must complete the reboot using Oracle ILOM.

To restart a server using Oracle ILOM:

1. Use your browser to open a connection to the server using Oracle ILOM. For example:

```
http://bda1node12-c.example.com
```

Note: Your browser must have a JDK plug-in installed. If you do not see the Java coffee cup on the log-in page, then you must install the plug-in before continuing.

2. Log in using your Oracle ILOM credentials.
3. Select the **Remote Control** tab.
4. Click the **Launch Remote Console** button.
5. Enter `Ctrl+d` to continue rebooting.
6. If the reboot fails, then enter the server `root` password at the prompt and attempt to fix the problem.
7. After the server restarts successfully, open the Redirection menu and choose **Quit** to close the console window.

See Also: Oracle Integrated Lights Out Manager (ILOM) 3.0 documentation at

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

Replacing a Disk Drive

Complete this procedure for all failed disk drives.

1. If you are replacing a working disk, then see "["Dismounting Partitions Before Replacing a Working Disk" on page 12-7](#)".
2. Replace the failed disk drive.
See "[Parts for Sun Fire Servers](#)" on page B-4.
3. Power on the server if you powered it off to replace the failed disk.
4. Connect to the server as root using either the KVM or an SSL connection to a laptop.
5. Store the physical drive information in a file:

```
# MegaCli64 pdlist a0 > pdinfo.tmp
```

Note: This command redirects the output to a file so that you can perform several searches using a text editor. If you prefer, you can pipe the output through the more or grep commands.

The utility returns the following information for each slot. This example shows a Firmware State of Unconfigured(good) , Spun Up.

```
Enclosure Device ID: 20
Slot Number: 8
Drive's position: DiskGroup: 8, Span: 0, Arm: 0
Enclosure position: 0
Device Id: 11
WWN: 5000C5003487075C
Sequence Number: 2
Media Error Count: 0
Other Error Count: 0
Predictive Failure Count: 0
Last Predictive Failure Event Seq Number: 0
PD Type: SAS
Raw Size: 1.819 TB [0xe8e088b0 Sectors]
Non Coerced Size: 1.818 TB [0xe8d088b0 Sectors]
Coerced Size: 1.817 TB [0xe8b6d000 Sectors]
Firmware state: Unconfigured(good) , Spun Up
Is Commissioned Spare : NO
Device Firmware Level: 061A
Shield Counter: 0
Successful diagnostics completion on : N/A
SAS Address(0): 0x5000c5003487075d
SAS Address(1): 0x0
Connected Port Number: 0(path0)
Inquiry Data: SEAGATE ST32000SSSUN2.0T061A1126L6M3WX
FDE Enable: Disable
Secured: Unsecured
Locked: Unlocked
Needs EKM Attention: No
Foreign State: None
Device Speed: 6.0Gb/s
Link Speed: 6.0Gb/s
```

```
Media Type: Hard Disk Device
```

- .
- .
- .
- 6. Open the file you created in Step 5 in a text editor and search for the following:
 - Disks that have a Foreign State of True
 - Disks that have a Firmware State of Unconfigured
- 7. For disks that have a Foreign State of Foreign, clear that status:

```
# MegaCli64 CfgForeign clear a0
```

A foreign disk is one that the controller saw previously, such as a reinserted disk.

- 8. For disks that have a Firmware State of Unconfigured (Bad), complete these steps:
 - a. Note the enclosure device ID number and the slot number.
 - b. Enter a command in this format:

```
# MegaCli64 pdmakegood physdrv[enclosure:slot] a0
```

For example, [20:10] repairs the disk identified by enclosure 20 in slot 10.

- 9. For disks that have a Firmware State of Unconfigured (Good), use the following command. If multiple disks are unconfigured, then configure them in order from the lowest to the highest slot number:

```
# MegaCli64 CfgLdAdd r0[enclosure:slot] a0
```

```
Adapter 0: Created VD 1
```

```
Adapter 0: Configured the Adapter!!
```

```
Exit Code: 0x00
```

For example, [20:5] repairs the disk identified by enclosure 20 in slot 5.

- 10. Verify the disk is recognized by the operating system.

```
# lsscsi
```

The disk may appear with its original device name (such as /dev/sdc) or under a new device name (such as /dev/sdn). If the operating system does not recognize the disk, then the disk is missing from the list generated by the lsscsi command.

This example output shows two disks with new device names: /dev/sdn in slot 5, and /dev/sdo in slot 10.

[0:0:20:0]	enclosu	SUN	HYDE12	0341	-
[0:2:0:0]	disk	LSI	MR9261-8i	2.12	/dev/sda
[0:2:1:0]	disk	LSI	MR9261-8i	2.12	/dev/sdb
[0:2:2:0]	disk	LSI	MR9261-8i	2.12	/dev/sdc
[0:2:3:0]	disk	LSI	MR9261-8i	2.12	/dev/sdd
[0:2:4:0]	disk	LSI	MR9261-8i	2.12	/dev/sde
[0:2:5:0]	disk	LSI	MR9261-8i	2.12	/dev/sdn
[0:2:6:0]	disk	LSI	MR9261-8i	2.12	/dev/sdg
[0:2:7:0]	disk	LSI	MR9261-8i	2.12	/dev/sdh
[0:2:8:0]	disk	LSI	MR9261-8i	2.12	/dev/sdi
[0:2:9:0]	disk	LSI	MR9261-8i	2.12	/dev/sdj
[0:2:10:0]	disk	LSI	MR9261-8i	2.12	/dev/sdo
[0:2:11:0]	disk	LSI	MR9261-8i	2.12	/dev/sdl

```
[7:0:0:0]      disk      Unigen      PSA4000      1100  /dev/sdm
[
```

- If the disk is listed with a new device name, such as /dev/sdn, then reboot the server to restore the original device mappings:

```
# reboot
```

- Reconnect to the server and reissue the lsscsi command to verify the original disk mappings. See [Table 12-1](#).

- Check the hardware profile of the server:

```
# bdachckhw
```

- Identify the function of the drive, so you configure it properly. See ["Identifying the Function of a Disk Drive"](#) on page 12-12.

Identifying the Function of a Disk Drive

Most disks are used for HDFS, as shown in [Table 12-1](#). Nonetheless, you should verify that the failed disk was not used for either the operating system or Oracle NoSQL Database before configuring it for a particular function.

Checking for Use by the Operating System

Oracle Big Data Appliance is configured with the operating system on the first two disks.

To confirm that a failed disk supported the operating system:

- Check whether the replacement disk corresponds to /dev/sda or /dev/sdb, which are the operating system disks.

```
# lsscsi
```

See the output from Step 10 of ["Replacing a Disk Drive"](#) on page 12-10..

- Verify that /dev/sda and /dev/sdb are the operating system mirrored partitioned disks:

```
# mdadm -Q --detail /dev/md2
```

```
Version : 0.90
Creation Time : Fri Oct 19 19:19:07 2012
Raid Level : raid1
.
.
.
Number  Major  Minor  RaidDevice State
      0      8      2          0  active sync  /dev/sda2
      1      8     18          1  active sync  /dev/sdb2
```

- If the previous steps indicate that the failed disk is an operating system disk, then proceed to ["Configuring an Operating System Disk"](#) on page 12-13.
- If the failed disk did not support the operating system, then proceed to ["Checking for Use by Oracle NoSQL Database"](#) on page 12-13..

Checking for Use by Oracle NoSQL Database

Oracle Big Data Appliance can be configured to allocate the last 0, 1, or 2 disks for the exclusive use of Oracle NoSQL Database. HDFS data does not reside on the same disks.

To discover whether a failed disk supported Oracle NoSQL Database:

1. Open an SSH connection to the first server in the rack and log in as the `root` user.
2. Obtain the value of `NOSQLDB_DISKS` from the `mammoth-rack_name.params` configuration file:


```
# cat /opt/oracle/BDAMammoth/mammoth-rackname.params | grep NOSQL
```
3. Use the value of `NOSQLDB_DISKS` to determine whether the replacement disk is allocated to Oracle NoSQL Database:
 - 0: No disks are allocated to Oracle NoSQL Database.
 - 1: The `/dev/sdl` disk is allocated to Oracle NoSQL Database.
 - 2: The `/dev/sdk` and `/dev/sdl` disks are allocated to Oracle NoSQL Database.
4. To verify that the disks are part of a logical volume, you can run either `pvscan` or `pvdisplay`. All disks allocated for use by Oracle NoSQL Database are presented to it as a single logical volume named `lvg1`.

These commands verify that `/dev/sdl1` is part of `lvg1`:

```
# pvscan
```

```
.
.
.
PV /dev/sdl1   VG lvg1   lvm2 [1.82 TB / 93.09 GB free]
Total: 1 [1.82 TB] / in use: 1 [1.82 TB] / in no VG: 0 [0 ]
```

```
# pvdisplay
```

```
--- Physical volume ---
PV Name          /dev/sdl1
VG Name          lvg1
.
.
```

5. If the previous steps indicate that the failed disk supported Oracle NoSQL Database, then proceed to "[Configuring an Oracle NoSQL Database Disk](#)" on page 12-19. Otherwise, proceed to "[Configuring an HDFS Disk](#)" on page 12-22.

Configuring an Operating System Disk

The first two disks (`/dev/sda`, `/dev/sdb`) support the Linux operating system. These disks store a copy of the mirrored operating system, a swap partition, a mirrored boot partition, and an HDFS data partition.

To configure an operating system disk, you must copy the partition table from the surviving disk, create an HDFS partition (ext4 file system), and add the software raid partitions and boot partitions for the operating system.

Complete these procedures after replacing logical drive 0 or 1 (`/dev/sda` or `/dev/sdb`).

- Partitioning the Operating System Disk
- Repairing the RAID Arrays
- Formatting the HDFS Partition of an Operating System Disk
- Restoring the Swap Partition
- Restoring the GRUB Master Boot Records

Partitioning the Operating System Disk

Note: Replace /dev/sdx in the following commands with the appropriate device name, either /dev/sda or /dev/sdb.

To partition a logical drive:

1. Complete the steps in "Replacing a Disk Drive" on page 12-10.
2. Confirm that the new disk does not have a partition table:

```
# parted /dev/sdx -s print
```

You should see a message about a missing partition table.

3. If the parted command displays a partition table, then clear it:

```
# dd if=/dev/zero of=/dev/sdx bs=1M count=100
```

Tip: You can use this command to restart an operating system disk configuration, if you make a mistake.

4. Create the partition table:

```
# parted /dev/sdx -s mklabel gpt print
```

5. List the Cylinder, Head, Sector (CHS) partition information of the *surviving* disk. Thus, if you are partitioning /dev/sda, then enter /dev/sdb for /dev/sdy in the following command:

```
# parted /dev/sdy -s unit chs print
```

```
Model: LSI MR9261-8i (scsi)
Disk /dev/sda: 243031,30,6
Sector size (logical/physical): 512B/512B
BIOS cylinder,head,sector geometry: 243031,255,63.  Each cylinder is 8225kB.
Partition Table: gpt
```

Number	Start	End	File system	Name	Flags
1	0,0,34	25,127,7	ext3		raid
2	25,127,8	21697,116,20	ext3		raid
3	21697,116,21	23227,61,35	linux-swap		
4	23227,61,36	243031,29,36	ext3		primary

6. Create partitions 1 to 3 on the new drive by duplicating the partitions of the surviving disk. Issue three commands in this format:

```
# parted /dev/sdx -s mkpart file_system start end
```

Use the start and end addresses that you obtained in Step 3 instead of the addresses shown in the following example:

```
# parted /dev/sdx -s mkpart ext3 0,0,34 25,127,7
# parted /dev/sdx -s mkpart ext3 25,127,8 21697,116,20
# parted /dev/sdx -s mkpart linux-swap 21697,116,21 23227,61,35
```

7. Create primary partition 4 using the start address obtained in Step 5 and an end address of 100%:

```
# parted /dev/sdx -s mkpart primary ext3 23227,61,36 100%
```

Partition 4 stores HDFS data, and this syntax makes the partition as large as possible.

8. Set the RAID flags:

```
# parted -s /dev/sdx set 1 raid
# parted -s /dev/sdx set 2 raid
```

9. Clear the names, using parted in interactive mode:

```
# parted /dev/sdx
```

```
GNU Parted 1.8.1
Using /dev/sdx
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) name 1 " "
(parted) name 2 " "
(parted) name 3 " "
(parted) quit
```

10. Complete the steps in "[Repairing the RAID Arrays](#)" on page 12-15..

Repairing the RAID Arrays

After partitioning the disks, you can repair the two logical RAID arrays:

- /dev/md0 contains /dev/sda1 and /dev/sdb1. It is mounted as /boot.
- /dev/md2 contains /dev/sda2 and /dev/sdb2. It is mounted as / (root).

Caution: Do not dismount the /dev/md devices, because that action shuts down the system.

To repair the RAID arrays:

1. Mark the partitions as failed:

```
# mdadm --fail /dev/md0 /dev/sdx1
# mdadm --fail /dev/md2 /dev/sdx2
```

You can ignore "No such device" messages about /dev/sdx in the mdadm commands.

2. Remove the partitions from the RAID arrays:

```
# mdadm --remove /dev/md0 /dev/sdx1
# mdadm --remove /dev/md2 /dev/sdx2
```

3. Verify that the RAID arrays are degraded:

```
# mdadm -Q --detail /dev/md0
```

```
# mdadm -Q --detail /dev/md2
```

4. Verify that the degraded file for each array is set to 1:

```
# cat /sys/block/md0/md/degraded
1
# cat /sys/block/md2/md/degraded
1
```

5. Restore the partitions to the RAID arrays:

```
# mdadm --add /dev/md0 /dev/sdx1
# mdadm --add /dev/md2 /dev/sdx2
```

6. Check that resynchronization is started, so that /dev/md2 is in a state of repair and not idle:

```
# cat /sys/block/md2/md/sync_action
repair
```

7. To verify that resynchronization is proceeding, you can monitor the mdstat file. A counter identifies the percentage complete.

```
# cat /proc/mdstat

Personalities : [raid1]
md0 : active raid1 sdb1[1] sda1[0]
      204736 blocks [2/2] [UU]

md2 : active raid1 sdb2[2] sda2[0]
      174079936 blocks [2/1] [U_]
      [=====>.....]  recovery = 61.6% (107273216/174079936)
finish=18.4min speed=60200K/sec
```

The following output shows that synchronization is complete:

```
Personalities : [raid1]
md0 : active raid1 sdb1[1] sda1[0]
      204736 blocks [2/2] [UU]

md2 : active raid1 sdb2[1] sda2[0]
      174079936 blocks [2/2] [UU]

unused devices: <none>
```

8. Compare the output of the following command with /etc/mdadm.conf:

```
# mdadm --examine --brief --scan --config=partitions
```

9. If the lines for /dev/md0 and /dev/md2 are different from the output, then use a text editor to replace them with the output of the mdadm command.

```
# cat /etc/mdadm.conf

# mdadm.conf written out by anaconda
DEVICE partitions
MAILADDR root
ARRAY /dev/md0 level=raid1 num-devices=2 UUID=df1bd885:c1f0f9c2:25d6...
ARRAY /dev/md2 level=raid1 num-devices=2 UUID=6c949a1a:1d45b778:a6da...
```

10. Restart the server:

```
# reboot
```

11. Complete the steps in ["Formatting the HDFS Partition of an Operating System Disk" on page 12-17..](#)

Formatting the HDFS Partition of an Operating System Disk

Partition 4 (sda4) on an operating system disk is used for HDFS. After you format the partition and set the correct label, HDFS rebalances the job load to use the partition if the disk space is needed.

To format the HDFS partition:

1. Format the HDFS partition as an ext4 file system:

```
# mke4fs -t ext4 /dev/sdx4

mke4fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
110354432 inodes, 441393655 blocks
22069682 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
13471 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
102400000, 214990848
```

Writing inode tables:

Note: If this command fails because the device is mounted, then skip to Step 3. After dismounting the device, return to Step 1. You do not need to repeat Step 3.

2. Verify that the partition label (such as /u01 for /dev/sda) is missing:

```
# ls -l /dev/disk/by-label
```

3. Dismount the appropriate HDFS partition, either /u01 for /dev/sda, or /u02 for /dev/sdb:

```
# umount /u0n
```

4. Reset the partition label:

```
# tune4fs -c -1 -i 0 -m 0.2 -L /u0n /dev/sdx4
```

5. Mount the HDFS partition:

```
# mount /u0n
```

6. Complete the steps in ["Restoring the Swap Partition" on page 12-18.](#)

Restoring the Swap Partition

After formatting the HDFS partition, you can restore the swap partition.

To restore the swap partition:

1. Set the swap label:

```
# mkswap -L SWAP-sdx3 /dev/sdx3
```

2. Verify that the swap partition is restored:

```
# swapon -s
```

Filename	Type	Size	Used	Priority
/dev/sdb3	partition	12287992	0	-1
/dev/sda3	partition	12287992	0	-2

3. Complete the steps in "Restoring the GRUB Master Boot Records" on page 12-18.

Restoring the GRUB Master Boot Records

After restoring the swap partition, you can restore the Grand Unified Bootloader (GRUB) master boot record.

To restore the GRUB boot record:

1. Open GRUB:

```
# grub --device-map=/boot/grub/device.map

GNU GRUB version 0.97 (640K lower / 3072K upper memory)

[ Minimal BASH-like line editing is supported. For the first word, TAB
lists possible command completions. Anywhere else TAB lists the possible
completions of a device/filename.

]

The device.map file maps the BIOS drives to operating system devices. The
following is an example of a device map file:
```

```
# this device map was generated by anaconda
(hd0)      /dev/sda
(hd1)      /dev/sdb
```

2. Set the root device, entering hd0 for /dev/sda, or hd1 for /dev/sdb:

```
grub> root (hdn,0)

root (hdn,0)
Filesystem type is ext2fs, partition type 0x83
```

3. Install grub, entering hd0 for /dev/sda, or hd1 for /dev/sdb:

```
grub> setup (hdn)

setup (hdn)
Checking if "/boot/grub/stage1" exists... no
Checking if "/grub/stage1" exists... yes
Checking if "/grub/stage2" exists... yes
Checking if "/grub/e2fs_stage1_5" exists... yes
Running "embed /grub/e2fs_stage1_5 (hdn)"... failed (this is not fatal)
Running "embed /grub/e2fs_stage1_5 (hdn,0)"... failed (this is not fatal)
Running "install /grub/stage1 (hdn) /grub/stage2 p /grub/grub.conf"...
succeeded
```

Done.

4. Close the GRUB command-line interface:

```
grub> quit
```

5. Ensure that logical drive L0 (L + zero) is set as the boot drive in the HBA:

```
# MegaCli64 -AdpBootDrive -get a0
Adapter 0: Boot Virtual Drive - #0 (target id - 0).
```

6. If the previous command does not report L0 or virtual drive 0 target 0, then enter:

```
# MegaCli64 AdpBootDrive set L0 a0
```

7. Ensure that the auto-select boot drive feature is enabled:

```
# MegaCli64 adpBIOS EnblAutoSelectBootLd a0
Auto select Boot is already Enabled on Adapter 0.
```

8. Check the configuration. See "["Verifying the Disk Configuration"](#) on page 12-23.

Configuring an Oracle NoSQL Database Disk

Depending on the owner's configuration choice, the last zero, one, or two disks (none, /dev/sd1, or /dev/sdk and /dev/sd1) store the data files for Oracle NoSQL Database.

Removing the Logical Volume for Oracle NoSQL Database

All disks allocated for use by Oracle NoSQL Database are configured as a single logical volume. To configure an Oracle NoSQL Database disk, you must destroy and re-create the volume.

To remove the logical volume:

1. Complete the steps in "["Replacing a Disk Drive"](#) on page 12-10.
2. Open /etc/fstab in a text editor and make these changes:
 - a. Delete the mount entry for /dev/lvg1/lv1.
 - b. Add mount points for either u12, or both u11 and 12.
3. Verify that the logical volume exists:

```
# lvdisplay
/dev/sdf: read failed after 0 of 4096 at 0: Input/output error
--- Logical volume ---
  LV Name                /dev/lvg1/lv1
  VG Name                lvg1
  LV UUID                BaWmSp-Ha8J-Qm18-9qEr-i1qH-mKJm-eEj1wD
  LV Write Access         read/write
  LV Status               available
  # open                  1
  LV Size                 3.45 TB
  Current LE              905536
  Segments                1
  Allocation              inherit
  Read ahead sectors      auto
  - currently set to      256
  Block device            253:0
```

4. Verify that the physical volume exists:

```
# pvdisplay
/dev/sdf: read failed after 0 of 4096 at 0: Input/output error
--- Physical volume ---
PV Name          /dev/sd11
VG Name          lvg1
PV Size          1.82 TB / not usable 1.97 MB
Allocatable      yes
PE Size (KByte) 4096
Total PE         476598
Free PE          23830
Allocated PE     452768
PV UUID          kHtDSR-gmvv-uLIm-69Y1-fp4Q-7bgf-tMhzZU

--- Physical volume ---
PV Name          /dev/sdk1
VG Name          lvg1
PV Size          1.82 TB / not usable 1.97 MB
Allocatable      yes
PE Size (KByte) 4096
Total PE         476598
Free PE          23830
Allocated PE     452768
PV UUID          FyYegA-xQd9-mff1-YhxX-3g9P-2ArQ-V2DYG0
```

5. If neither the logical nor the physical volumes exist, then skip to Step 13. Otherwise, continue with the next step.
6. Dismount the logical volume.

```
# umount /lv1
```

Note: If you are unable to dismount the volume, then complete the procedure in ["Stopping the Oracle NoSQL Database Service"](#) on page 12-21..

7. Deactivate the logical volume:

```
# lvchange -a n /dev/lvg1/lv1
```

8. Remove the logical volume:

```
# lvremove -f /dev/lvg1/lv1
```

Logical volume "lv1" successfully removed

9. Deactivate the volume group:

```
# vgchange -a n /dev/lvg1
```

0 logical volume(s) in volume group "lvg1" now active

10. Remove the volume group.

```
# vgremove -f lvg1
```

Volume group "lvg1" successfully removed

11. Identify the physical partitions allocated to logical volume management:

```
# pvdisplay
```

```

"/dev/sdk1" is a new physical volume of "1.82 TB"
--- NEW Physical volume ---
PV Name           /dev/sdk1
VG Name
PV Size          1.82 TB
Allocatable      NO
PE Size (KByte) 0
Total PE         0
Free PE          0
Allocated PE     0
PV UUID          FyYegA-xQd9-mff1-YhxX-3g9P-2ArQ-V2DYG0

"/dev/sd11" is a new physical volume of "1.82 TB"
--- NEW Physical volume ---
PV Name           /dev/sd11
VG Name
PV Size          1.82 TB
Allocatable      NO
PE Size (KByte) 0
Total PE         0
Free PE          0
Allocated PE     0
PV UUID          kHtDSR-gmvv-uLIm-69Y1-fp4Q-7bgf-tMhzZU

```

12. Remove the physical volumes:

```
# pvremove -ff disk1 [disk2]
```

This example removes two allocated partitions:

```
# pvremove -ff /dev/sdk1 /dev/sd11
Labels on physical volume "/dev/sdk1" successfully wiped
Labels on physical volume "/dev/sd11" successfully wiped
```

13. Configure the drive or drives for HDFS. See "["Configuring an HDFS Disk" on page 12-22.](#)

Stopping the Oracle NoSQL Database Service

You may need to stop the service if two disks are allocated to Oracle NoSQL Database, or if you are replacing a working disk.

To stop the service:

1. Obtain the process ID (PID) of the Oracle NoSQL Database service:

```
# fuser -m /dev/mapper/lvg1-lv1
```

```
/dev/mapper/lvg1-lv1: 11154 11430 11443
```

2. Use the PID to obtain more information about the process:

```
# ps -fp process_id
```

This example uses PID 11154 from the example output Step 1.

```
# ps -fp 11154
UID PID PPID C STIME TTY TIME CMD
oracle 11154 1 0 09:29 ? 00:00:00 java -Xms64m -Xmx64m
oracle.kv.impl.sna.StorageNodeAgentImpl -root /lv1/kvroot
```

3. If your output is similar to the example shown in Step 2, then stop the Oracle NoSQL Database service:

```
# service nsdbservice stop
```

4. Restart the configuration beginning with Step 6 in "Removing the Logical Volume for Oracle NoSQL Database" on page 12-19, which dismounts the device.

Configuring an HDFS Disk

An HDFS disk is not used by the operating system nor by Oracle NoSQL Database. See "Identifying the Function of a Disk Drive" on page 12-12.

To configure an HDFS disk, you must partition and format it.

Note: Replace /dev/sdx in the following commands with the appropriate device name, such as /dev/sdc.

To configure a disk for use by HDFS:

1. Complete the steps in "Replacing a Disk Drive" on page 12-10.
2. Partition the drive:

```
# parted /dev/sdx -s mklabel gpt mkpart primary ext3 0% 100%
```

3. Format the HDFS partition as an ext4 file system.

```
# mke4fs -t ext4 /dev/sdx1
```

```
mke4fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
122011648 inodes, 488036855 blocks
24401842 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
14894 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
102400000, 214990848

Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

This filesystem will be automatically checked every 23 mounts or 180 days, whichever comes first. Use tune4fs -c or -i to override.

4. Verify that the device is missing:

```
# ls -l /dev/disk/by-label
```

In this example output, sdf1 is missing:

```
total 0
lrwxrwxrwx 1 root root 10 Nov  8 12:08 BDAUSB -> ../../sdm1
```

```
lrwxrwxrwx 1 root root 10 Nov  8 12:08 u01 -> ../../sda4
lrwxrwxrwx 1 root root 10 Nov  8 12:08 u02 -> ../../sdb4
lrwxrwxrwx 1 root root 10 Nov  8 12:08 u03 -> ../../sdc1
lrwxrwxrwx 1 root root 10 Nov  8 12:08 u04 -> ../../sdd1
lrwxrwxrwx 1 root root 10 Nov  8 12:08 u05 -> ../../sde1
lrwxrwxrwx 1 root root 10 Nov  8 12:08 u07 -> ../../sdg1
lrwxrwxrwx 1 root root 10 Nov  8 12:08 u08 -> ../../sdh1
lrwxrwxrwx 1 root root 10 Nov  8 12:08 u09 -> ../../sdi1
lrwxrwxrwx 1 root root 10 Nov  8 12:08 u10 -> ../../sdj1
```

5. Reset the appropriate partition label to the missing device. See [Table 12-1](#).

```
# tune4fs -c -1 -i 0 -m 0.2 -L /unn /dev/sdx1
```

For example, this command resets the label for /dev/sdf to /u06:

```
# tune4fs -c -1 -i 0 -m 0.2 -L /u06 /dev/sdf1
Setting maximal mount count to -1
Setting interval between checks to 0 seconds
Setting reserved blocks percentage to 0.2% (976073 blocks)
```

6. Mount the HDFS partition, entering the appropriate mount point:

```
# mount /unn
```

For example, mount /u03.

7. If you removed a mount point in Cloudera Manager earlier, then restore it:

- Open a browser window to Cloudera Manager. For example:

```
http://bdalnode02.example.com:7180
```

- Log in as admin.

- On the Services page, click **hdfs1**.

- Click the **Instances** subtab.

- In the Host column, locate the server with the replaced disk. Then click the service in the Name column, such as **hdfs1-DATANODE-12**, to open its page.

- Click the **Configuration** subtab.

- If the mount point is missing from the Directory field, then add it to the list.

- Click **Save Changes**.

- From the Actions list, choose **Restart**.

The disk is now configured for use by HDFS.

8. Check the configuration. See ["Verifying the Disk Configuration"](#) on page 12-23.

Verifying the Disk Configuration

Be sure to verify that the configuration is correct on the new disk drive.

To verify the disk configuration:

1. Check the software configuration:

```
# bdachcksw
```

2. If there are errors, then redo the configuration steps as necessary to correct the problem.

3. When bdachecksw returns no errors, restart the server.
4. Check the /root directory for a file named BDA_REBOOT_SUCCEEDED. This file must exist for the Mammoth Reconfiguration Utility to run.

Note: Allow several minutes for the file to refresh after you restart the server. Check the time stamp to ensure that you are seeing the current file.

5. If you find a file named BDA_REBOOT_FAILED, then read the file to identify the problem.
6. Return to Step 2.
7. Repeat these steps until restarting the server results in a BDA_REBOOT_SUCCEEDED file.

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 quality assurance (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.

The following is an example of an 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 8-26 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 Settings](#)
- [Replacing a Failed InfiniBand Switch](#)
- [Verifying InfiniBand Network Operation](#)
- [Understanding the Network Subnet Manager Master](#)

Backing Up and Restoring Oracle ILOM Settings

Oracle 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 Mammoth Utility.

See Also: Oracle Integrated Lights Out Manager 3.0 documentation at
<http://docs.oracle.com/cd/E19860-01/>

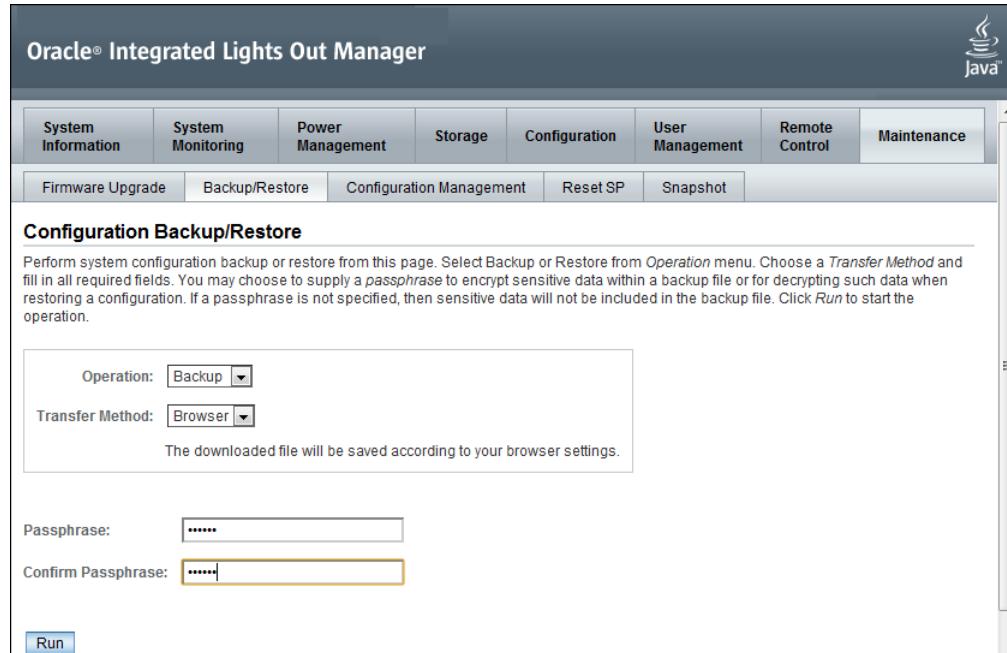
Backing Up Oracle ILOM Configuration Settings

To back up the Oracle ILOM configuration settings:

1. Open your browser on any system that is on the same network as Oracle Big Data Appliance, and enter the Oracle ILOM address of a server. This example uses the Oracle ILOM address of node08:
`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 12-1](#).
6. Enter a pass phrase. The phrase is used to encrypt sensitive information, such as passwords, in the backup file.

7. Click **Run** to start the backup operation. 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 12–1 Oracle ILOM Configuration Backup



Restoring Oracle ILOM Configuration Settings

To restore the Oracle ILOM configuration settings:

1. Open your browser on any system on the same network as Oracle Big Data Appliance and navigate to an Oracle ILOM service processor. This example uses Oracle ILOM on node08:

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 operation.
8. Click **Run** to restore the configuration.

Replacing a Failed InfiniBand Switch

Complete these steps to replace a Sun Network QDR InfiniBand Gateway Switch or a Sun Datacenter InfiniBand Switch 36.

See Also:

- ["InfiniBand Switch-to-Server Cable Connections" on page C-7](#) for information about 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 file, as described in ["Backing Up and Restoring Oracle ILOM Settings" on page 12-25](#).
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@bda1sw-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@bda1sw-02-> ibdiagnet -c 1000 -r
```

10. Enable the Subnet Manager:

```
FabMan@bda1sw-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 reenable 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 that the InfiniBand network is operating properly. The following procedure describes how to verify network operation:

Note: Use this procedure any time the InfiniBand network is performing below expectations.

To verify InfiniBand network operation:

1. Enter 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.

See Also: *Sun Network QDR InfiniBand Gateway Switch Command Reference* at

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

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
```

See Also: Linux man page for `ibqueryerrors.S`

3. Check the status of the hardware:

```
# bdachckhw
```

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
[1m[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:

```
# 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
```

Understanding the Network Subnet Manager Master

The Subnet Manager manages all operational characteristics of the InfiniBand network, such as the ability to:

- 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 the leaf switches are configured as priority 5, and the Subnet Managers on the spine switches are configured as priority 8.

The following guidelines determine where the Subnet Managers run on Oracle Big Data Appliance:

- Run the Subnet Managers only on the switches in Oracle Big Data Appliance. Running a 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 a 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 the spine switches run a Subnet Manager. The leaf switches must disable the 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://docs.oracle.com/cd/E19197-01/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 and Clock on an InfiniBand Switch"](#) on page 8-12.

To update the Cisco Ethernet switch:

1. Use telnet to connect to the Cisco Ethernet switch.
2. Delete 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 8-4](#)

Restart Oracle Big Data Appliance after changing the servers and switches.

A

Error Messages for Oracle Big Data Appliance

This appendix describes the hardware messages for Oracle Big Data Appliance.

Auto Service Request Alerts: DISKALRT-02001 to DISKALRT-02006

The following messages are Auto Service Request (ASR) alerts. ASR automatically submits a service request to Oracle Support Services in response to these alerts.

Information is also written to the system logs.

DISKALRT-02001 System OS hard disk failure

Cause: A physical disk that is hosting the operating system failed or an operating system RAID volume has lost redundancy. HDFS might also raise an alert.

Action: Replace the physical disk as soon as possible. See ["Replacing a Server Disk" on page 12-6](#) and follow the instructions for configuring an operating system disk.

DISKALRT-02002 System OS hard disk predictive failure

Cause: A physical disk that is hosting the operating system is in a state of predictive failure.

Action: Replace the physical disk as soon as possible. See ["Replacing a Server Disk" on page 12-6](#) and follow the instructions for configuring an operating system disk.

DISKALRT-02003 System OS hard disk poor performance

Cause: A physical disk that is hosting the operating system is performing poorly, which indicates that it is failing.

Action: Plan to replace the disk. If multiple disk disks in a server are identified as poor performers, then check for problems with the cable connections, the controllers, and the disk backplane before replacing the disks. To replace a disk, see ["Replacing a Server Disk" on page 12-6](#) and follow the instructions for configuring an operating system disk.

DISKALRT-02004 Data hard disk failure

Cause: A physical disk that is *not* hosting the operating system failed.

Action: Replace the physical disk as soon as possible. See ["Replacing a Server Disk" on page 12-6](#) and follow the instructions for configuring a disk for either HDFS or Oracle NoSQL Database.

DISKALRT-02005 Data hard disk predictive failure

Cause: A physical disk that is *not* hosting the operating system is in a state of predictive failure.

Action: Replace the physical disk as soon as possible. See ["Replacing a Server Disk"](#) on page 12-6 and follow the instructions for configuring a disk for either HDFS or Oracle NoSQL Database.

DISKALRT-02006 Data hard disk poor performance

Cause: A physical disk that is *not* hosting the operating system is performing poorly, which indicates that it is failing.

Action: Replace the physical disk as soon as possible. See ["Replacing a Server Disk"](#) on page 12-6 and follow the instructions for configuring a disk for either HDFS or Oracle NoSQL Database.

B

Replacement Units

This appendix lists the replacement units for Oracle Big Data Appliance. All replacement units are field replaceable units (FRUs), which are installed by Oracle field engineers.

This appendix contains the following sections:

- [Understanding Repair Categories](#)
- [Disk Controller Batteries](#)
- [LED Status Descriptions](#)
- [Parts for Sun Fire Servers](#)
- [Parts for the Sun Datacenter InfiniBand Switch 36](#)
- [Parts for the Sun Network QDR InfiniBand Gateway Switches](#)
- [Parts for the Cisco Ethernet Switch](#)
- [Parts for the KVM Switch and KMM Tray Components](#)
- [Parts for Power Distribution Units](#)
- [Cables and Cable Accessories for an Oracle Big Data Appliance Rack](#)

Understanding Repair Categories

Repair categories describe the potential effects of a repair action. [Table B-1](#) describes the categories and related actions.

Table B-1 Oracle Big Data Appliance Repair Categories

Category	Description	Scenario Examples
Hot Swap (HS)	The 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 an 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 or KMM tray

Table B-1 (Cont.) Oracle Big Data Appliance 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 workload and data duplication may temporarily affect 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 12-2 for additional information.</p>	<ul style="list-style-type: none"> ▪ System boards ▪ PCIe cards ▪ Memory ▪ Processors

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 the battery charge is insufficient, when the temperature is high, and when the battery should be replaced.

Battery charge capacity degrades over time, and the life expectancy of a battery is inversely proportional to the operating temperature. The worst-case life expectancy of a battery in an Oracle Big Data Appliance rack is as follows:

Inlet Ambient Temperature	Battery Lifetime
< 25 degrees Celsius (77 degrees Fahrenheit)	3 years
< 32 degrees Celsius (90 degrees Fahrenheit)	2 years

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 Oracle Premier Support for Systems, Oracle attempts to proactively replace the batteries in an Oracle Big Data Appliance rack before the end of the estimated lifetime, on a best effort basis.

LED Status Descriptions

This section describes the light-emitting diodes (LEDs) in an Oracle Big Data Appliance rack. It contains these topics:

- [Sun Fire Server LEDs](#)
- [Sun Datacenter InfiniBand Switch 36 LEDs](#)

Sun Fire Server LEDs

[Table B-2](#) describes the LED color codes on Sun Fire X4270 M2 servers.

Table B-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 B-3 describes the color codes of the LEDs on Sun Datacenter InfiniBand Switch 36 switches.

Table B-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 the link is established, it is off when the 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.

Table B-3 (Cont.) Sun Datacenter InfiniBand Switch 36 LED Status Descriptions

Component	LED Status
Power supply	<ul style="list-style-type: none"> ■ OK LED is green: 12 VDC is supplied. Flashing indicates no function. ■ Attention LED is amber: There is a fault, and 12 VDC is shut down. Flashing indicates no function. ■ AC LED is green: AC power is supplied. Flashing indicates no function.

Parts for Sun Fire Servers

Table B-4 lists the replaceable parts for Sun Fire servers, which are covered under the Oracle 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 an Oracle service representative arrives if circumstances warrant immediate action.

Table B-4 Replacement Parts for a Sun Fire X4270 M2 Server

Part Number	Description	Repair Category	Repair Procedure
150-3993	Battery, 3V, 390 MAH, 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, 1200W, 12V, 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 12-2 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 Guide</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
371-4966	LV DIMM, 8 GB, DDR3, 1RX4, 1333 MHz	BDA-O	See Section 4.2, "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

Table B-4 (Cont.) Replacement Parts for a Sun Fire X4270 M2 Server

Part Number	Description	Repair Category	Repair Procedure
371-4982	6 Gigabit SAS RAID PCI battery module, RoHS:Y (LION), BBU-08	BDA-O	<ol style="list-style-type: none"> If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 12-2 for additional information. See <i>Sun Storage 6 Gb SAS PCIe HBA, Internal Installation Guide For HBA Models SGX-SAS6-INT-Z and SG-SAS6-INT-Z</i> (internal only) at http://docs.oracle.com/cd/E19337-01/index.html
375-3696	InfiniBand (CX2) dual port CX2 4x QDR PCI-e	BDA-O	<ol style="list-style-type: none"> If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 12-2 for additional information. 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
375-3701	6 Gbps SAS RAID disk PCIe HBA, RoHS:Y	BDA-O	<ol style="list-style-type: none"> If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 12-2 for additional information. 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	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
530-4228	Cable, mini-SAS, 36 POS LONG, LY, RoHS:Y	BDA-O	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
541-2075	Bus_bar, power, 1U+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_85936
541-2884	Assembly, riser, X8, X8, 2U, PCI-E, RoHS:Y	BDA-O	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 B-4 (Cont.) Replacement Parts for a Sun Fire X4270 M2 Server

Part Number	Description	Repair Category	Repair Procedure
541-3310	Xeon Heatsink with grease, RoHS:Y, 2U	BDA-O	See Section 4.7, "Servicing CPUs" 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
541-3356	x8/x8 switched PCIe riser assembly, RoHS:Y with the following: <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 12-2 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
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 12-2 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 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

Table B-4 (Cont.) Replacement Parts for a Sun Fire X4270 M2 Server

Part Number	Description	Repair Category	Repair Procedure
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
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
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.
7010036	FRU DR 3 TB 3.5-inch 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.
7012433	3.06 GHz processor	BDA-O	See Section 4.7, "Servicing CPUs" 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

Parts for the Sun Datacenter InfiniBand Switch 36

Table B-5 lists the replaceable parts for the Sun Datacenter InfiniBand Switch 36, which are covered under the Oracle Premier Support for Systems warranty. Unless otherwise noted, the repair procedures reference the *Sun Datacenter InfiniBand Switch 36 Service Manual* at

http://docs.oracle.com/cd/E26698_01/index.html

See "Cables and Cable Accessories for an Oracle Big Data Appliance Rack" on page B-10 for information about InfiniBand cables.

Table B-5 Replacement Parts for the Sun Datacenter InfiniBand Switch 36

Part Number	Description	Repair Category	Repair Procedure
300-2299	A247A 760 watt power supply, RoHS:Y	HS	See "Servicing Power Supplies" at http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f49112.html
350-1312	Fan module for Sun Datacenter InfiniBand Switch 36, RoHS:Y	HS	See "Servicing Fans" at http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f59112.html
371-2210	CR2032 3V battery, RoHS:Y	IR	See "Servicing the Battery" at http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f79112.html
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 12-26. 2. See "Installing the Switch" in the <i>Sun Datacenter InfiniBand Switch 36 Installation Guide</i> at http://docs.oracle.com/cd/E26698_01/html/E26434/cggfbhhf.html
594-6603	10 Gbps QSFP short wavelength transceiver	IR	See "Servicing Data Cables" at http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f69112.html

Parts for the Sun Network QDR InfiniBand Gateway Switches

Table B-6 lists the replaceable parts for the Sun Network QDR InfiniBand Gateway Switch, which are covered under the Oracle Premier Support for Systems warranty. Unless otherwise noted, the repair procedures refer to the *Sun Network QDR InfiniBand Gateway Switch Service Manual* at

http://docs.oracle.com/cd/E26699_01/index.html

See "[Cables and Cable Accessories for an Oracle Big Data Appliance Rack](#)" on page B-10 for information about InfiniBand cables.

Table B-6 Replacement Parts for the Sun Network QDR InfiniBand Gateway Switch

Part Number	Description	Repair Category	Repair Procedure
300-2299	760 watt power supply, RoHS:Y	HS	See "Servicing Power Supplies" at http://docs.oracle.com/cd/E26699_01/html/E26706/z40001f49112.html
350-1312	Fan module, RoHS:Y	HS	See "Servicing Fans" at http://docs.oracle.com/cd/E26699_01/html/E26706/z40001f59112.html

Table B-6 (Cont.) Replacement Parts for the Sun Network QDR InfiniBand Gateway Switch

Part Number	Description	Repair Category	Repair Procedure
371-2210	CR2032 3V battery, RoHS:Y	IR	See "Servicing the Battery" at http://docs.oracle.com/cd/E26699_01/html/E26706/z40001f79112.html
594-6603	10 Gbps QSFP short wavelength transceiver	IR	See "Servicing Data Cables" at http://docs.oracle.com/cd/E26699_01/html/E26706/z40001f69112.html
7014378	Assembly, system, NM2-Gateway	IR	<ol style="list-style-type: none"> 1. See "Replacing a Failed InfiniBand Switch" on page 12-26. 2. See "Installing the Gateway" in the <i>Sun Network QDR InfiniBand Gateway Switch Installation Guide</i> at http://docs.oracle.com/cd/E26699_01/html/E26706/cggfbhhf.html

Parts for the Cisco Ethernet Switch

Table B-7 lists replaceable parts for the Cisco Ethernet switch, which are covered under the Oracle Premier Support for Systems warranty.

See Also: *Cisco Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at

http://www.cisco.com/en/US/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html

Table B-7 Replacement Parts for Cisco Catalyst Ethernet Switch

Part Number	Description	Repair Category	Repair
7055121	Cooling fan	HS	See Chapter 4 of <i>Cisco Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> .
7055122	Cisco Catalyst 4948 switch, RoHS:Y	IR	See Chapter 3 of <i>Cisco Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> .
7055123	Power supply	HS	See Chapter 4 of <i>Cisco Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> .

Parts for the KVM Switch and KMM Tray Components

Table B-8 lists the replacement parts for the KVM switch and KMM tray components, which are covered under the Oracle 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 B-8 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 B-9 lists the replacement parts for the power distribution units (PDUs), which are covered under the Oracle Premier Support for Systems warranty.

Because of the number of cables and connections, Oracle recommends shutting down the Oracle Big Data Appliance rack during the replacement of 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 B-9 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: 15 kVA, IEC309 60A, 4-pin, 250 VAC, three-phase 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: 15 kVA, 5-pin, IEC309 32A, 5-pin 230/400V, three-phase 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: 15 kVA, 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: 15 kVA, 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 an Oracle Big Data Appliance Rack

Table B-10 lists the replacement cables and cable accessories for an Oracle Big Data Appliance rack, which are covered under the Oracle Premier Support for Systems warranty.

See Also:

- *Sun Network QDR InfiniBand Gateway Switch Service Manual* at
http://docs.oracle.com/cd/E26699_01/html/E26706/index.html
- *Sun Datacenter InfiniBand Switch 36 Service Manual* at
http://docs.oracle.com/cd/E26698_01/html/E26434/index.html
- See *Cisco Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide* at
http://www.cisco.com/en/US/docs/switches/lan/catalyst4900/4948E/installation/guide/4948E_ins.html

Table B-10 Replacement Cables for an Oracle Big Data Appliance Rack

Part Number	Description	Repair Category	Repair
350-1287	Slim rail, CMA, 1U-2U for Sun Fire servers	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 12-2 for additional information. 2. See "Installing the Cable Management Arm" in <i>Sun Fire X4170, X4270, and X4275 Servers Installation Manual</i> at http://docs.oracle.com/cd/E19477-01/820-5827-13/rack_install.html#50614309_99582
350-1519	Serial cable kit, RoHS:Y with the following: <ul style="list-style-type: none"> ▪ USB to DB 9-pin M serial cable ▪ DB 9-pin F to DB 9-pin F null modem cable 	HS	See "Understanding Cabling" in the <i>Sun Datacenter InfiniBand Switch 36 Installation Guide</i> at http://docs.oracle.com/cd/E26698_01/html/E26434/cggfhjcc.html
530-4403	4x QSFP copper QDR InfiniBand cable, 2 m, RoHS:Y	HS	<ol style="list-style-type: none"> 1. See "Maintaining the InfiniBand Network" on page 12-25. 2. See "Understanding Cabling" in the <i>Sun Datacenter InfiniBand Switch 36 Installation Guide</i> at http://docs.oracle.com/cd/E26698_01/html/E26434/cggfhjcc.html <p>Note: Redundancy is lost during cable replacement.</p>
530-4404	4x QSFP copper QDR InfiniBand cable, 3 m, RoHS:Y	HS	<ol style="list-style-type: none"> 1. See "Maintaining the InfiniBand Network" on page 12-25. 2. See "Understanding Cabling" in the <i>Sun Datacenter InfiniBand Switch 36 Installation Guide</i> at http://docs.oracle.com/cd/E26698_01/html/E26434/cggfhjcc.html

Table B-10 (Cont.) Replacement Cables for an Oracle Big Data Appliance Rack

Part Number	Description	Repair Category	Repair
530-4415	4x QSFP copper QDR InfiniBand cable, 5 m, RoHS:Y	HS	<p>1. See ""Maintaining the InfiniBand Network" on page 12-25.</p> <p>2. See "Understanding Cabling" in the <i>Sun Datacenter InfiniBand Switch 36 Installation Guide</i> at http://docs.oracle.com/cd/E26698_01/html/E26434/cggfhjcc.html</p>
530-4432	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, blue, RoHS:Y	HS	See <i>Cisco Catalyst 4948E and Catalyst 4948E-F 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 and Catalyst 4948E-F Switch Installation Guide</i> .
530-4434	Ethernet cable, Cat 5/5E, RJ45 to RJ45, 7 foot, red, RoHS:Y	HS	See <i>Cisco Catalyst 4948E and Catalyst 4948E-F 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 and Catalyst 4948E-F 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 and Catalyst 4948E-F 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 and Catalyst 4948E-F 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 and Catalyst 4948E-F Switch Installation Guide</i> .

C

In-Rack Cabling Tables

The tables in this appendix show the in-rack cable connections for Oracle Big Data Appliance. This appendix contains the following sections:

- [Cable Color Coding](#)
- [Oracle Big Data Appliance Rack Layout](#)
- [KVM Network Cable Connections](#)
- [Administrative Gigabit Ethernet Cable Connections](#)
- [Oracle Integrated Lights Out Manager Cable Connections](#)
- [Single-Phase Power Distribution Unit Cable Connections](#)
- [Three-Phase Power Distribution Unit Cable Connections](#)
- [InfiniBand Switch-to-Switch Cable Connections](#)
- [InfiniBand Switch-to-Server Cable Connections](#)

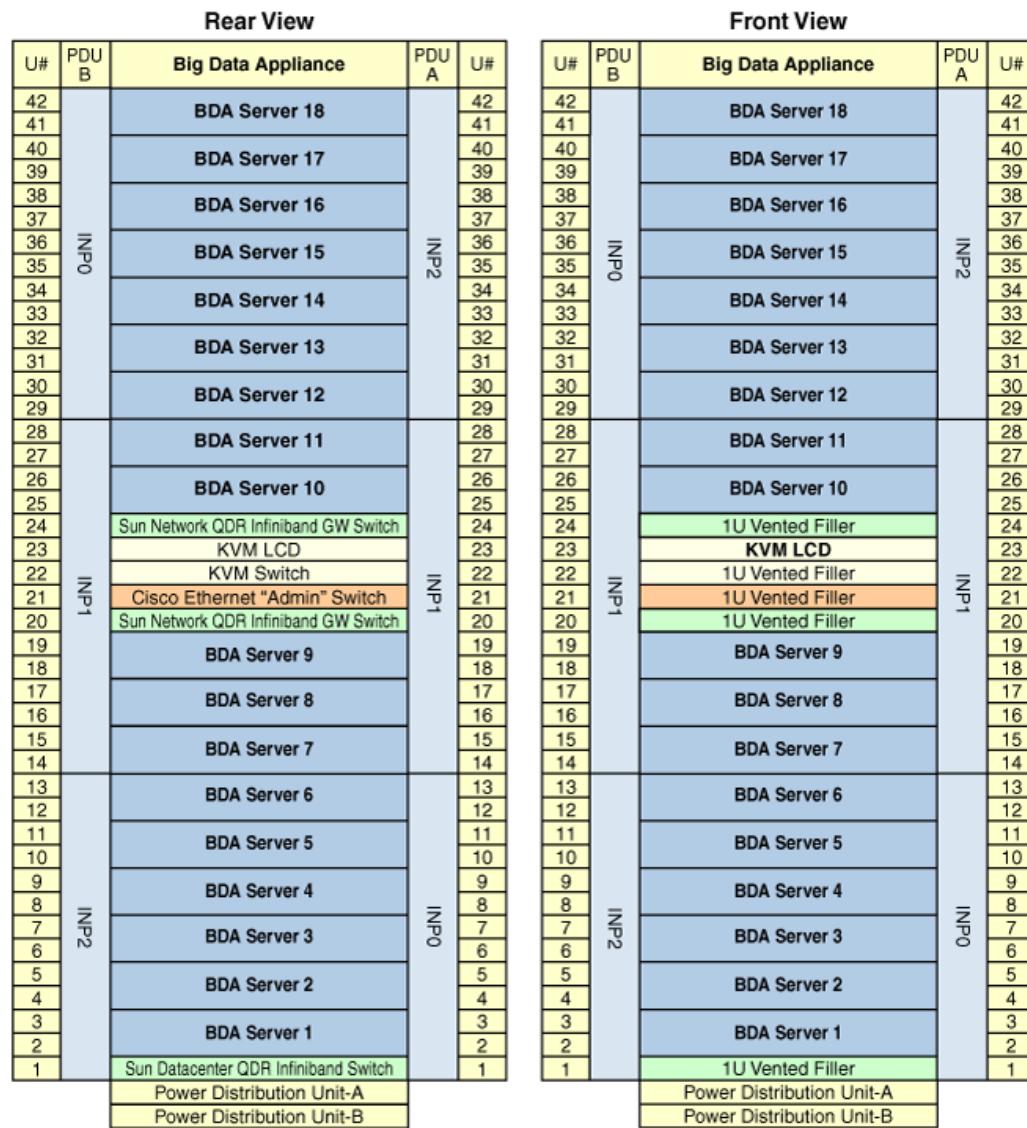
Cable Color Coding

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 (Oracle ILOM) Ethernet management cables
- Blue: Gigabit Ethernet management (eth0) cables
- Orange: KVM switch to dongle cables

Oracle Big Data Appliance Rack Layout

[Figure C-1](#) shows the front and rear layout of an Oracle Big Data Appliance rack.

Figure C-1 Rack Layout of Oracle Big Data Appliance

KVM Network Cable Connections

A KVM switch is provided only in Sun Fire X4270 M2-based rack.

Table C-1 shows the keyboard, video, and mouse (KVM) network cabling to the servers. The cables are orange.

Table C-1 KVM Cabling

From Rack Unit	Component	To KVM Port
U41	Server 18	1
U39	Server 17	2
U37	Server 16	3
U35	Server 15	4
U33	Server 14	5

Table C-1 (Cont.) KVM Cabling

From Rack Unit	Component	To KVM Port
U31	Server 13	6
U29	Server 12	7
U28	Server 11	8
U27	Server 10	9
U26	NM2-IB leaf switch	10
U25	KVM switch	11
U19	KVM tray	12
U18	Cisco switch	13
U17	NM2-IB leaf switch	14
U16	Server 9	15
U14	Server 8	16
U12	Server 7	17
U10	Server 6	18
U8	Server 5	19
U6	Server 4	20
U4	Server 3	21
U2	Server 2	22

Administrative Gigabit Ethernet Cable Connections

Table C-2 shows the cable connections for the administrative Gigabit Ethernet network (Net 0). The Ethernet switch is located in rack unit 21. The cables are 10 feet long and black.

Table C-2 Administrative Gigabit Ethernet Cabling

From Rack Unit ¹	Component	To Ethernet Port
U41	Server 18	1
U39	Server 17	3
U37	Server 16	5
U35	Server 15	7
U33	Server 14	9
U31	Server 13	11
U29	Server 12	13
U27	Server 11	17
U25	Server 11	21
U24	NM2-IB Switch	45
U20	NM2-IB Switch	46
U18	Server 10	25
U16	Server 9	29

Table C-2 (Cont.) Administrative Gigabit Ethernet Cabling

From Rack Unit ¹	Component	To Ethernet Port
U14	Server 8	31
U12	Server 7	33
U10	Server 6	35
U8	Server 5	37
U6	Server 4	39
U4	Server 3	41
U2	NM2-IB Switch	43
U1	Server 2	47

¹ *Un* is the unit location in the rack, where *n* is the number.

Oracle Integrated Lights Out Manager Cable Connections

Table C-3 shows the cable connections from the servers to the Oracle ILOM switch. The Oracle ILOM ports on the servers are labeled Net Mgt and connect to the Cisco Ethernet switch located in rack unit 21.

Table C-3 Oracle ILOM Cabling

From Rack Unit ¹	Component	Ethernet Port	Cable Length	Cable Color
U41	Server 18	2	10 feet	Red
U39	Server 17	4	10 feet	Red
U37	Server 16	6	10 feet	Red
U35	Server 15	8	10 feet	Red
U33	Server 14	10	10 feet	Red
U31	Server 13	12	10 feet	Red
U29	Server 12	14	10 feet	Red
U27	Server 11	18	10 feet	Red
U25	Server 10	22	10 feet	Red
U18	Server 9	26	10 feet	Red
U16	Server 8	30	10 feet	Red
U14	Server 7	32	10 feet	Red
U12	Server 6	34	10 feet	Red
U10	Server 5	36	10 feet	Red
U8	Server 4	38	10 feet	Red
U6	Server 3	40	10 feet	Red
U4	Server 2	42	10 feet	Red
U2	Server 1	44	10 feet	Red
PDU-A	PDU-A	15	1 meter	White
PDU-B	PDU-B	20	1 meter	White
Not applicable	Service	48	10 feet	Blue

¹ *Un* is the unit location in the rack, where *n* is the number.

Single-Phase Power Distribution Unit Cable Connections

Table C-4 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 C-4 Single-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U41	G5-6	G0-0	2 meters
U39	G5-3	G0-3	2 meters
U37	G5-0	G0-6	2 meters
U35	G4-6	G1-0	2 meters
U33	G4-4	G1-2	2 meters
U31	G4-2	G1-4	2 meters
U29	G3-6	G2-0	2 meters
U27	G3-5	G2-1	2 meters
U25	G3-3	G2-3	2 meters
U24	G3-1	G2-5	2 meters
U23	Not Applicable	G3-0	included
U22	G2-5	G3-1	1 meter
U21	G3-0	G2-6	2 meters
U20	G2-4	G3-2	2 meters
U18	G2-2	G3-4	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters
U1	G0-0	G5-6	2 meters

¹ *Un* is the unit location in the rack, where *n* is the number.

Three-Phase Power Distribution Unit Cable Connections

Table C-5 describes three-phase cabling from each power distribution unit (PDU) to the power supplies in the rack. The cables are terminated to PDU-A on the left, are routed to the right to enter the cable management arm (CMA), and are bundled in groups of four.

Table C-5 Three-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U41	G5-6	G2-0	2 meters
U39	G5-3	G2-3	2 meters
U37	G5-0	G2-6	2 meters
U35	G4-6	G1-0	2 meters
U33	G4-4	G1-2	2 meters
U31	G4-2	G1-4	2 meters
U29	G3-6	G0-0	2 meters
U27	G3-5	G0-1	2 meters
U25	G3-3	G0-3	2 meters
U24	G3-1	G0-5	2 meters
U23	Not Applicable	G5-0	Included
U22	G2-5	G5-1	1 meter
U21	G3-0	G0-6	2 meters
U20	G2-4	G5-2	2 meters
U18	G2-2	G5-4	2 meters
U16	G1-6	G4-0	2 meters
U14	G2-0	G5-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters
U1	G0-0	G3-6	2 meters

¹ Un is the unit location in the rack, where n is the number.

InfiniBand Switch-to-Switch Cable Connections

Table C-6 lists the location, ports, and cables for the InfiniBand connections from switch to switch. Figure C-2 identifies the locations of the ports on a Sun Network QDR InfiniBand Gateway Switch.

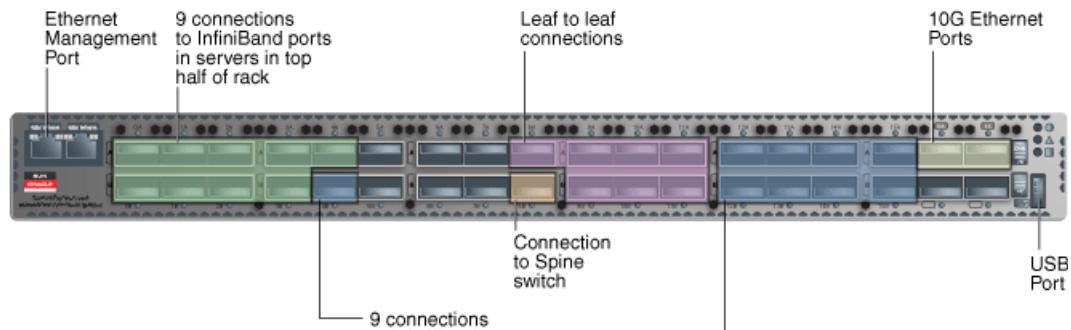
Table C-6 InfiniBand Switch-to-Switch Cabling

From InfiniBand Switch Rack Unit ¹	Port	To InfiniBand Switch Rack Unit ¹	Port	Cable Length	Description
U20	9B	U24	9A	2 meters	Leaf to leaf
U20	10B	U24	10A	2 meters	Leaf to leaf
U20	11B	U24	11A	2 meters	Leaf to leaf

Table C–6 (Cont.) InfiniBand Switch-to-Switch Cabling

From InfiniBand Switch Rack Unit ¹	Port	To InfiniBand Switch Rack Unit ¹	Port	Cable Length	Description
U20	8A	U24	8A	2 meters	Leaf to leaf
U20	9A	U24	9B	2 meters	Leaf to leaf
U20	10A	U24	10B	2 meters	Leaf to leaf
U20	11A	U24	11B	2 meters	Leaf to leaf
U1	0B	U20	8B	3 meters	Spine to leaf
U1	1B	U24	8B	3 meters	Spine to leaf

¹ *Un* is the unit location in the rack, where *n* is the number.

Figure C–2 Sun Network QDR InfiniBand Gateway Switch Ports

InfiniBand Switch-to-Server Cable Connections

Table C–7 lists the location, ports, and cables for the InfiniBand connections from the leaf switches to the servers.

Table C–7 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
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

Table C-7 (Cont.) InfiniBand Switch-to-Server Cabling

From InfiniBand Switch Rack Unit¹	Port	To Rack Unit¹	Port²	Cable Length
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 P1	2 meters
U20	12A	U16	PCIe 3 P1	2 meters
U20	12B	U14	PCIe 3 P1	3 meters
U20	13A	U12	PCIe 3 P1	3 meters
U20	13B	U10	PCIe 3 P1	3 meters
U20	14A	U8	PCIe 3 P1	3 meters
U20	14B	U6	PCIe 3 P1	3 meters
U20	15A	U4	PCIe 3 P1	3 meters
U20	15B	U2	PCIe 3 P1	3 meters
U24	4B	U18	PCIe 3 P2	2 meters
U24	12A	U16	PCIe 3 P2	2 meters
U24	12B	U14	PCIe 3 P2	3 meters
U24	13A	U12	PCIe 3 P2	3 meters
U24	13B	U10	PCIe 3 P2	3 meters
U24	14A	U8	PCIe 3 P2	3 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 *Un* is the unit location in the rack, where *n* is the number.² *Pn* is the InfiniBand port, where *n* is the port number.

D

Multirack Cabling Tables

This appendix contains the tables for multirack cabling. It contains the following sections:

- [Understanding Multirack Cabling](#)
- [Key to Cabling Table Abbreviations](#)
- [Two-Rack Cabling](#)
- [Three-Rack Cabling](#)
- [Four-Rack Cabling](#)
- [Five-Rack Cabling](#)
- [Six-Rack Cabling](#)
- [Seven-Rack Cabling](#)
- [Eight-Rack Cabling](#)

Understanding Multirack Cabling

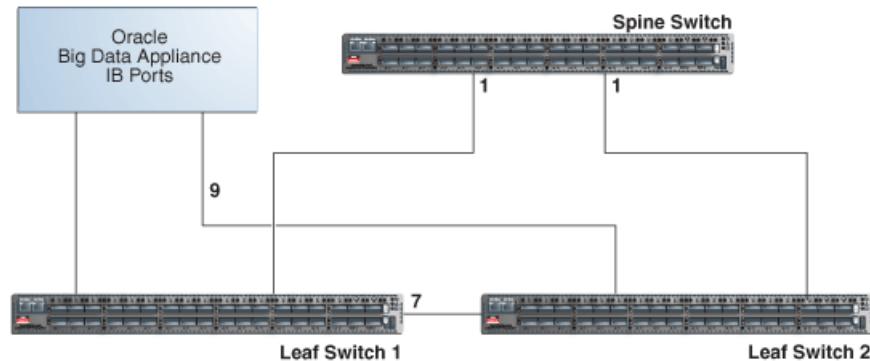
You can connect multiple Oracle Big Data Appliance racks or a combination of Oracle Big Data Appliance racks and Oracle Exadata Database Machine full racks, half racks, or quarter racks. Exadata half racks must have a spine switch.

The 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 C-1](#).

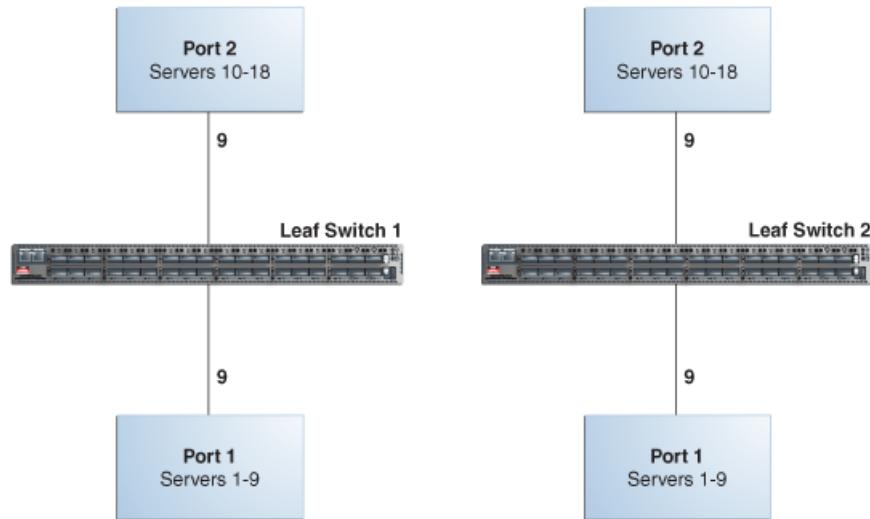
These switches attach to standard quad small form-factor pluggable (QSFP) connectors at the end of the InfiniBand cables. You can cable together up to eight racks without external InfiniBand switches. The cable lengths in this appendix are for adjacent racks. If they are apart, then longer cables may be required for the connections.

Cabling the Switches in One Rack

In a single rack, the two leaf switches are connected by seven cables. Each leaf switch has one cable to the spine switch. Nine servers connect to one leaf switch, and the other nine servers connect to the other leaf switch, as shown in [Figure D-1](#).

Figure D-1 Connections Between Spine Switch and Leaf Switches

The Sun Fire servers connect to the leaf switches as shown in [Figure D-2](#).

Figure D-2 Connections Between the Servers and the Leaf Switches

Cabling the Switches in Two Racks

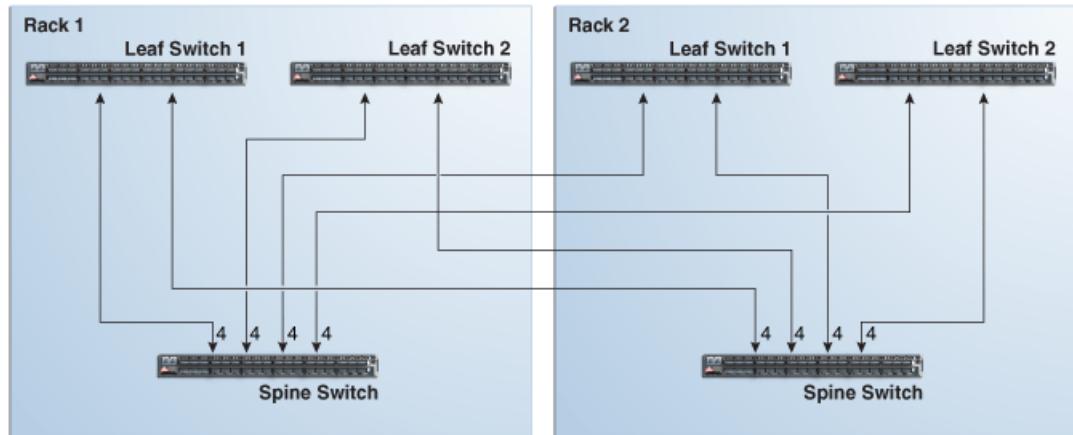
When you connect multiple racks, remove the seven interswitch connections between leaf switches and the two connections between the leaf switches and the spine switch. From each leaf switch, distribute eight connections over the spine switches in all racks. In multirack environments, the leaf switches inside a rack are not directly interconnected, as shown in [Figure D-3](#).

[Figure D-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:

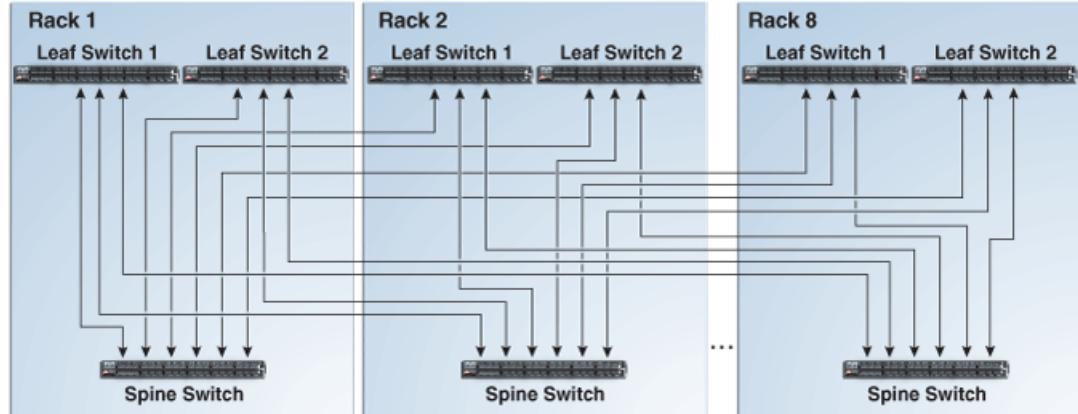
- Eight connections to the internal leaf switches (four to each rack 1 switch)
- Eight connections to the leaf switches in rack 2 (four to each rack 2 switch)

Figure D-3 Connecting Switches Across Two Racks

Cabling the Switches in Up to Eight Racks

As the number of racks increases from two to eight, the pattern continues as shown in [Figure D-4](#). Because not all of the racks are shown, not all of the cables are shown either.

Each leaf switch has eight interswitch connections distributed over all spine switches. Each spine switch has 16 interswitch 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.

Figure D-4 Connecting Switches Across Multiple Racks

Key to Cabling 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 D-1](#) shows the cable connections for the first spine switch (R1-U1) when two racks are cabled together.

Note:

- The spine switch is in U1 (IB1) for all racks.
- The leaf switches are in U20 (IB2) and U24 (IB3).
- In this section, the leaf switches are called IB2 and IB3.

Table D-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 D-2](#) shows the cable connections for the second spine switch (R2-U1) when two racks are cabled together.

Table D-2 Leaf Switch Connections for the Second Rack in a Two-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A R2-IB3-P9A to R2-U1-P5A R2-IB3-P9B to R2-U1-P6A	3 meters
R2 IB3 to Rack 1	R2-IB3-P10A to R1-U1-P7A R2-IB3-P10B to R1-U1-P8A R2-IB3-P11A to R1-U1-P9A R2-IB3-P11B to R1-U1-P10A	5 meters

Table D-2 (Cont.) Leaf Switch Connections for the Second Rack in a Two-Rack System

Leaf Switch	Connection	Cable Length
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 D-3](#) shows the cable connections for the first spine switch (R1-U1) when three racks are cabled together.

Note:

- The spine switch is in U1 (IB1) for all racks.
- The leaf switches are in U20 (IB2) and U24 (IB3).
- In this section, the leaf switches are called IB2 and IB3.

Table D-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
R1 IB2 to Rack 3	R1-IB2-P11A to R3-U1-P9B R1-IB2-P11B to R3-U1-P10B	5 meters

[Table D-4](#) shows the cable connections for the second spine switch (R2-U1) when three racks are cabled together.

Table D-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 D-5 shows the cable connections for the third spine switch (R3-U1) when three racks are cabled together.

Table D-5 Leaf Switch Connections for the Third Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A R3-IB3-P8B to R3-U1-P4A R3-IB3-P9A to R3-U1-P5A	3 meters
R3 IB3 to Rack 1	R3-IB3-P9B to R1-U1-P6A R3-IB3-P10A to R1-U1-P7A R3-IB3-P10B to R1-U1-P8A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11A to R2-U1-P9A R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B R3-IB2-P8B to R3-U1-P4B R3-IB2-P9A to R3-U1-P5B	3 meters
R3 IB2 to Rack 1	R3-IB2-P9B to R1-U1-P6B R3-IB2-P10A to R1-U1-P7B R3-IB2-P10B to R1-U1-P8B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11A to R2-U1-P9B R3-IB2-P11B to R2-U1-P10B	5 meters

Four-Rack Cabling

Table D-6 shows the cable connections for the first spine switch (R1-U1) when four racks are cabled together.

Note:

- The spine switch is in U1 (IB1) for all racks.
- The leaf switches are in U20 (IB2) and U24 (IB3).
- In this section, the leaf switches are called IB2 and IB3.

Table D-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 D-7 shows the cable connections for the second spine switch (R2-U1) when four racks are cabled together.

Table D-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

Table D-7 (Cont.) Leaf Switch Connections for the Second Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
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 D-8 shows the cable connections for the third spine switch (R3-U1) when four racks are cabled together.

Table D-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 D-9 shows the cable connections for the fourth spine switch (R4-U1) when four racks are cabled together.

Table D-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

Table D-9 (Cont.) Leaf Switch Connections for the Fourth Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
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 D-10 shows the cable connections for the first spine switch (R1-U1) when five racks are cabled together.

Note:

- The spine switch is in U1 (IB1) for all racks.
- The leaf switches are in U20 (IB2) and U24 (IB3).
- In this section, the leaf switches are called IB2 and IB3.

Table D-10 Leaf Switch Connections for the First Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A R1-IB3-P8B to R1-U1-P4A	3 meters
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A R1-IB3-P9B to R2-U1-P6A	5 meters
R1 IB3 to Rack 3	R1-IB3-P10A to R3-U1-P7A R1-IB3-P10B to R3-U1-P8A	5 meters
R1 IB3 to Rack 4	R1-IB3-P11A to R4-U1-P9A	10 meters
R1 IB3 to Rack 5	R1-IB3-P11B to R5-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B	3 meters
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B R1-IB2-P9B to R2-U1-P6B	5 meters
R1 IB2 to Rack 3	R1-IB2-P10A to R3-U1-P7B R1-IB2-P10B to R3-U1-P8B	5 meters

Table D-10 (Cont.) Leaf Switch Connections for the First Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
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 D-11 shows the cable connections for the second spine switch (R2-U1) when five racks are cabled together.

Table D-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 D-12 shows the cable connections for the third spine switch (R3-U1) when five racks are cabled together.

Table D-12 Leaf Switch Connections for the Third Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A R3-IB3-P8B to R3-U1-P4A	3 meters
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P9A to R4-U1-P5A R3-IB3-P9B to R4-U1-P6A	5 meters
R3 IB3 to Rack 5	R3-IB3-P10A to R5-U1-P7A R3-IB3-P10B to R5-U1-P8A	5 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B R3-IB2-P8B to R3-U1-P4B	3 meters

Table D-12 (Cont.) Leaf Switch Connections for the Third Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
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 D-13](#) shows the cable connections for the fourth spine switch (R4-U1) when five racks are cabled together.

Table D-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 D-14](#) shows the cable connections for the fifth spine switch (R5-U1) when five racks are cabled together.

Table D-14 Leaf Switch Connections for the Fifth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A R5-IB3-P8B to R5-U1-P4A	3 meters
R5 IB3 to Rack 1	R5-IB3-P9A to R1-U1-P5A R5-IB3-P9B to R1-U1-P6A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10A to R2-U1-P7A R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters

Table D-14 (Cont.) Leaf Switch Connections for the Fifth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
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 D-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 this section, the leaf switches are called IB2 and IB3.

Table D-15 Leaf Switch Connections for the First Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A R1-IB3-P8B to R1-U1-P4A	3 meters
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A R1-IB3-P9B to R2-U1-P6A	5 meters
R1 IB3 to Rack 3	R1-IB3-P10A to R3-U1-P7A	5 meters
R1 IB3 to Rack 4	R1-IB3-P10B to R4-U1-P8A	10 meters
R1 IB3 to Rack 5	R1-IB3-P11A to R5-U1-P9A	10 meters
R1 IB3 to Rack 6	R1-IB3-P11B to R6-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B	3 meters
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B R1-IB2-P9B to R2-U1-P6B	5 meters
R1 IB2 to Rack 3	R1-IB2-P10A to R3-U1-P7B	5 meters
R1 IB2 to Rack 4	R1-IB2-P10B to R4-U1-P8B	10 meters
R1 IB2 to Rack 5	R1-IB2-P11A to R5-U1-P9B	10 meters
R1 IB2 to Rack 6	R1-IB2-P11B to R6-U1-P10B	10 meters

Table D-16 shows the cable connections for the second spine switch (R2-U1) when six racks are cabled together.

Table D-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 D-17 shows the cable connections for the third spine switch (R3-U1) when six racks are cabled together.

Table D-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
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

Table D-17 (Cont.) Leaf Switch Connections for the Third Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
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 D-18 shows the cable connections for the fourth spine switch (R4-U1) when six racks are cabled together.

Table D-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 D-19 shows the cable connections for the fifth spine switch (R5-U1) when six racks are cabled together.

Table D-19 Leaf Switch Connections for the Fifth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A R5-IB3-P8B to R5-U1-P4A	3 meters
R5 IB3 to Rack 1	R5-IB3-P10A to R1-U1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P9A to R6-U1-P5A R5-IB3-P9B to R6-U1-P6A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B R5-IB2-P8B to R5-U1-P4B	3 meters

Table D-19 (Cont.) Leaf Switch Connections for the Fifth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
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 D-20 shows the cable connections for the sixth spine switch (R6-U1) when six racks are cabled together.

Table D-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 D-21 shows the cable connections for the first spine switch (R1-U1) when seven racks are cabled together.

Note:

- The spine switch is in U1 (IB1) for all racks.
- The leaf switches are in U20 (IB2) and U24 (IB3).
- In this section, the leaf switches are called IB2 and IB3.

Table D-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 D-22 shows the cable connections for the second spine switch (R2-U1) when cabling seven racks.

Table D-22 Leaf Switch Connections for the Second Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P9A to R3-U1-P5A	5 meters
R2 IB3 to Rack 4	R2-IB3-P9B to R4-U1-P6A	5 meters
R2 IB3 to Rack 5	R2-IB3-P10A to R5-U1-P7A	10 meters
R2 IB3 to Rack 6	R2-IB3-P10B to R6-U1-P8A	10 meters
R2 IB3 to Rack 7	R2-IB3-P11A to R7-U1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P9A to R3-U1-P5B	5 meters
R2 IB2 to Rack 4	R2-IB2-P9B to R4-U1-P6B	5 meters
R2 IB2 to Rack 5	R2-IB2-P10A to R5-U1-P7B	10 meters
R2 IB2 to Rack 6	R2-IB2-P10B to R6-U1-P8B	10 meters
R2 IB2 to Rack 7	R2-IB2-P11A to R7-U1-P9B	10 meters

Table D-23 shows the cable connections for the third spine switch (R3-U1) when seven racks are cabled together.

Table D-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 D-24 shows the cable connections for the fourth spine switch (R4-U1) when seven racks are cabled together.

Table D-24 Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A R4-IB3-P8B to R4-U1-P4A	3 meters
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-U1-P5A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9B to R6-U1-P6A	5 meters
R4 IB3 to Rack 7	R4-IB3-P10A to R7-U1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B R4-IB2-P8B to R4-U1-P4B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-U1-P5B	5 meters

Table D-24 (Cont.) Leaf Switch Connections for the Fourth Rack in a Seven-Rack

Leaf Switch	Connection	Cable Length
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 D-25 shows the cable connections for the fifth spine switch (R5-U1) when seven racks are cabled together.

Table D-25 Leaf Switch Connections for the Fifth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R5 IB3 within Rack 5	R5-IB3-P8A to R5-U1-P3A R5-IB3-P8B to R5-U1-P4A	3 meters
R5 IB3 to Rack 1	R5-IB3-P10A to R1-U1-P7A	10 meters
R5 IB3 to Rack 2	R5-IB3-P10B to R2-U1-P8A	10 meters
R5 IB3 to Rack 3	R5-IB3-P11A to R3-U1-P9A	5 meters
R5 IB3 to Rack 4	R5-IB3-P11B to R4-U1-P10A	5 meters
R5 IB3 to Rack 6	R5-IB3-P9A to R6-U1-P5A	5 meters
R5 IB3 to Rack 7	R5-IB3-P9B to R7-U1-P6A	5 meters
R5 IB2 within Rack 5	R5-IB2-P8A to R5-U1-P3B R5-IB2-P8B to R5-U1-P4B	3 meters
R5 IB2 to Rack 1	R5-IB2-P10A to R1-U1-P7B	10 meters
R5 IB2 to Rack 2	R5-IB2-P10B to R2-U1-P8B	10 meters
R5 IB2 to Rack 3	R5-IB2-P11A to R3-U1-P9B	5 meters
R5 IB2 to Rack 4	R5-IB2-P11B to R4-U1-P10B	5 meters
R5 IB2 to Rack 6	R5-IB2-P9A to R6-U1-P5B	5 meters
R5 IB2 to Rack 7	R5-IB2-P9B to R7-U1-P6B	5 meters

Table D-26 shows the cable connections for the sixth spine switch (R6-U1) when seven racks are cabled together.

Table D-26 Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-U1-P3A R6-IB3-P8B to R6-U1-P4A	3 meters
R6 IB3 to Rack 1	R6-IB3-P9B to R1-U1-P6A	10 meters
R6 IB3 to Rack 2	R6-IB3-P10A to R2-U1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-U1-P8A	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

Table D-26 (Cont.) Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
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 D-27 shows the cable connections for the seventh spine switch (R7-U1) when seven racks are cabled together.

Table D-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 D-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 this section, the leaf switches are called IB2 and IB3.

Table D-28 Leaf Switch Connections for the First Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A	3 meters
R1 IB3 to Rack 2	R1-IB3-P8B to R2-U1-P4A	5 meters
R1 IB3 to Rack 3	R1-IB3-P9A to R3-U1-P5A	5 meters
R1 IB3 to Rack 4	R1-IB3-P9B to R4-U1-P6A	10 meters
R1 IB3 to Rack 5	R1-IB3-P10A to R5-U1-P7A	10 meters
R1 IB3 to Rack 6	R1-IB3-P10B to R6-U1-P8A	10 meters
R1 IB3 to Rack 7	R1-IB3-P11A to R7-U1-P9A	10 meters
R1 IB3 to Rack 8	R1-IB3-P11B to R8-U1-P10A	10 meters
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
R1 IB2 to Rack 2	R1-IB2-P8B to R2-U1-P4B	5 meters
R1 IB2 to Rack 3	R1-IB2-P9A to R3-U1-P5B	5 meters
R1 IB2 to Rack 4	R1-IB2-P9B to R4-U1-P6B	10 meters
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 D-29 shows the cable connections for the second spine switch (R2-U1) when eight racks are cabled together.

Table D-29 Leaf Switch Connections for the Second Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R2 IB3 within Rack 2	R2-IB3-P8A to R2-U1-P3A	3 meters
R2 IB3 to Rack 1	R2-IB3-P11B to R1-U1-P10A	5 meters
R2 IB3 to Rack 3	R2-IB3-P8B to R3-U1-P4A	5 meters
R2 IB3 to Rack 4	R2-IB3-P9A to R4-U1-P5A	5 meters
R2 IB3 to Rack 5	R2-IB3-P9B to R5-U1-P6A	10 meters
R2 IB3 to Rack 6	R2-IB3-P10A to R6-U1-P7A	10 meters
R2 IB3 to Rack 7	R2-IB3-P10B to R7-U1-P8A	10 meters
R2 IB3 to Rack 8	R2-IB3-P11A to R8-U1-P9A	10 meters
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B	3 meters
R2 IB2 to Rack 1	R2-IB2-P11B to R1-U1-P10B	5 meters
R2 IB2 to Rack 3	R2-IB2-P8B to R3-U1-P4B	5 meters
R2 IB2 to Rack 4	R2-IB2-P9A to R4-U1-P5B	5 meters
R2 IB2 to Rack 5	R2-IB2-P9B to R5-U1-P6B	10 meters
R2 IB2 to Rack 6	R2-IB2-P10A to R6-U1-P7B	10 meters
R2 IB2 to Rack 7	R2-IB2-P10B to R7-U1-P8B	10 meters
R2 IB2 to Rack 8	R2-IB2-P11A to R8-U1-P9B	10 meters

[Table D-30](#) shows the cable connections for the third spine switch (R3-U1) when eight racks are cabled together.

Table D-30 Leaf Switch Connections for the Third Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R3 IB3 within Rack 3	R3-IB3-P8A to R3-U1-P3A	3 meters
R3 IB3 to Rack 1	R3-IB3-P11A to R1-U1-P9A	5 meters
R3 IB3 to Rack 2	R3-IB3-P11B to R2-U1-P10A	5 meters
R3 IB3 to Rack 4	R3-IB3-P8B to R4-U1-P4A	5 meters
R3 IB3 to Rack 5	R3-IB3-P9A to R5-U1-P5A	5 meters
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 D-31](#) shows the cable connections for the fourth spine switch (R4-U1) when eight racks are cabled together.

Table D-31 Leaf Switch Connections for the Fourth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A	3 meters
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P8B to R5-U1-P4A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9A to R6-U1-P5A	5 meters
R4 IB3 to Rack 7	R4-IB3-P9B to R7-U1-P6A	10 meters
R4 IB3 to Rack 8	R4-IB3-P10A to R8-U1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P8B to R5-U1-P4B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9A to R6-U1-P5B	5 meters

Table D-31 (Cont.) Leaf Switch Connections for the Fourth Rack in an Eight-Rack

Leaf Switch	Connection	Cable Length
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 D-32 shows the cable connections for the fifth spine switch (R5-U1) when eight racks are cabled together.

Table D-32 Leaf Switch Connections for the Fifth Rack in an 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 D-33 shows the cable connections for the sixth spine switch (R6-U1) when eight racks are cabled together.

Table D-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

Table D-33 (Cont.) Leaf Switch Connections for the Sixth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
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 D-34 shows the cable connections for the seventh spine switch (R7-U1) when eight racks are cabled together.

Table D-34 Leaf Switch Connections for the Seventh Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R7 IB3 within Rack 7	R7-IB3-P8A to R7-U1-P3A	3 meters
R7 IB3 to Rack 1	R7-IB3-P9A to R1-U1-P5A	10 meters
R7 IB3 to Rack 2	R7-IB3-P9B to R2-U1-P6A	10 meters
R7 IB3 to Rack 3	R7-IB3-P10A to R3-U1-P7A	10 meters
R7 IB3 to Rack 4	R7-IB3-P10B to R4-U1-P8A	10 meters
R7 IB3 to Rack 5	R7-IB3-P11A to R5-U1-P9A	5 meters
R7 IB3 to Rack 6	R7-IB3-P11B to R6-U1-P10A	5 meters
R7 IB3 to Rack 8	R7-IB3-P8B to R8-U1-P4A	5 meters
R7 IB2 within Rack 7	R7-IB2-P8A to R7-U1-P3B	3 meters
R7 IB2 to Rack 1	R7-IB2-P9A to R1-U1-P5B	10 meters
R7 IB2 to Rack 2	R7-IB2-P9B to R2-U1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-U1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-U1-P8B	10 meters
R7 IB2 to Rack 5	R7-IB2-P11A to R5-U1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB2-P11B to R6-U1-P10B	5 meters
R7 IB2 to Rack 8	R7-IB3-P8B to R8-U1-P4B	5 meters

Table D-35 shows the cable connections for the eighth spine switch (R8-U1) when eight racks are cabled together.

Table D-35 Leaf Switch Connections for the Eighth Rack in an 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

Table D-35 (Cont.) Leaf Switch Connections for the Eighth Rack in an Eight-Rack

Leaf Switch	Connection	Cable Length
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

AC breaker panels, 2-8
AC power, location, 2-8
access requirements, 2-3
access route requirements, 2-3
accessing servers from KVM, 8-3
acoustic levels, 2-1
add subcommand, 11-11
adding racks, 11-4
administrative network, 3-2
 cabling requirements, 2-11
 configuration, 8-20
 connections, 2-11
 ports, 3-1
administrative network diagram, 6-2
administrative utilities, 1-3
Africa
 PDU specifications, 2-6, 2-7
air conditioning, 2-9, 2-13, 5-2
airflow, 2-2, 2-8
airflow requirements, 2-9
alerts, A-1
 generating, 6-1
 sent by Cloudera Manager, 4-8
Allen wrench keys, 5-8
altitude requirements, 2-8
ambient relative humidity range, 2-9
ambient temperature range, 2-9, 12-1
Americas
 PDU specifications, 2-4, 2-5
Apache OpenOffice Spreadsheet, 4-2
Asia Pacific
 PDU specifications, 2-6, 2-7
ASR
 about, 10-1
 activating assets, 10-5
 ASR Manager registration, 10-4
 configuration, 10-2
 troubleshooting, 10-7
 See also Auto Service Request
ASR Manager
 port number, 4-8
ASR validation
 generating test events, 10-6
Auto Service Request

alerts, A-1
installing, 4-7, 11-9
turning on, 11-10
See also ASR

B

backups
 external NFS directory, 4-7
base image reinstallation, 8-25
batteries
 life expectancy, B-2
 replacement policy, B-2
BDA server offline repair category, B-2
bda_imaging_status file, 8-20
BDA_REBOOT files, 5-14
bdacheckcluster utility, 8-27
bdacheckhw command, 8-19, 12-28
bdacheckhw utility, 8-27
bdacheckib utility, 8-28
bda-checkip.out file, 4-4
bdachecknet command, 8-29
bdachecksw command, 8-19, 12-29
bdachecksw utility, 8-30
BdaDeploy.json file, 4-3, 8-20, 8-26
bdadiag utility, 8-30
bdaid utility, 8-31
bdaimagevalidate utility, 8-31
bda-install-template.pdf file, 4-3
BDAMammoth directory, 11-3
BDA-O repair category, B-2
bdaredoclientnet utility, 8-31
bdaserials utility, 8-32
Big Data Appliance
 See Oracle Big Data Appliance
Big Data Connectors
 See Oracle Big Data Connectors
biosconfig command, 8-19
black cables, C-1
blue cables, C-1
bolts, 5-7
bondeth0 interface
 host names, 4-5
 network connections, 3-1
bondib0 interface
 description, 3-2

host names, 4-5
maintenance, 12-25
bondib0 network
 verifying operations, 12-27
boot order check, 8-19
breaker panels, 2-8

C

cable tie,
 removing, 5-11
cable-routing floor cutout, 5-8
cables
 color coding, C-1
 spares, 1-3
 splitter, 3-4
cabling
 client network, 3-3
 eight racks, D-19
 existing network requirements, 2-11
 five racks, D-9
 four racks, D-7
 multirack, 9-2, D-1
 seven racks, D-15
 six racks, D-12
 spare parts, B-10
 three racks, D-5
 two racks, D-4
cage nuts, 5-3
casters, 5-5
cautions, 12-3
CDH
 cluster configuration, 1-5, 11-5
 installing, 11-9
CDH ports, 3-6
CDH services, starting, 11-9
ceiling height requirements, 2-2
checking firmware, 8-19
circuit breakers, 2-8
 general requirements, 2-8
circuit breakers for PDUs, 5-13
Cisco switch
 replacement parts, B-9
 See Ethernet switch
CleanupInstall step, 11-9
CleanupSSHroot step, 11-10
CLI, Oracle ILOM, 6-2
client access network
 cabling requirements, 2-11
 description, 3-2
 requirements, 2-11
client network
 bandwidth, 3-4
 cabling, 3-3
clock
 configuring Ethernet switch, 8-7
 configuring InfiniBand switches, 8-12
Cloudera Manager, 4-8
Cloudera's Distribution including Apache Hadoop
 See CDH

cluster configurations, 1-5
cluster names, 4-5
color coding, cables, C-1
commands
 remote execution across servers, 1-3, 7-1
component failure
 environmental causes, 2-8
components, location in the rack, C-2
configuration
 disk controller, 12-4
 physical drives, 12-5
 virtual drive, 12-5
configuration files, 8-20
 generating, 4-2
 location, 4-3
 ntp.conf file, 12-30
configuration utility, 4-1
configuration worksheets, 2-12, 4-2
configure terminal command, 8-5
configuring a new disk, 12-10
configuring ASR, 10-2
configuring Ethernet switch, 8-4, 8-5
configuring KVM, 8-1, 8-3
configuring servers, 8-17
connecting KVM to the network, 8-2
connectors
 See Oracle Big Data Connectors
cooling requirements, 2-1
copying files, 7-2
CopyLicenseFiles step, 11-7
CopySoftwareSource step, 11-7
corrosion problems, 2-9
CreateLogicalVolumes step, 11-7
CreateUsers step, 11-8

D

data center switch configuration, 2-12
dba group, 4-8, 11-8
dcli utility
 description, 1-3
 examples, 7-5
 option descriptions, 7-4
 return values, 7-5
 syntax, 7-3
 target servers, 7-2, 8-18
 using, 7-1
device mappings
 restoring original, 12-12
dimensions
 rack, 2-2
 shipping, 2-2
disablesm command, 8-14, 9-2, 9-3, 9-4, 12-27
disabling spanning tree, 8-9
disk configuration
 verifying, 12-23
disk controller configuration, 12-4
disk controllers, 12-4
disk drive mappings, 12-6
disk drives

identifying dedicated function, 12-12
spares, 1-2, B-4
disks
 repairing, 12-4
 replacing, 12-6
 software configuration, 12-10
dismounting HDFS partitions, 12-8
DNS network registration, 2-12
DNS server configuration, 8-7
door dimensions, 2-3
downgrading the factory image, 11-6

E

earth ground cable, 5-9
electric receptacles, 2-4
electrical access, 2-3
electrical requirements, 2-3, 5-2, 5-13
electrical storms, 2-4
electronic noise, 5-2
electrostatic discharge, 2-9
elevator requirements, 2-3
email accounts
 My Oracle Support, 10-2
email configuration, 4-8
emergency shutdown, 12-3
enablesm command, 8-14, 9-2, 9-3, 9-4
env_test command, 8-13
environmental specifications, 2-1
error messages, A-1
eth0 interface, 3-2
Ethernet cabling, 2-11, 3-3
Ethernet interface, 3-2
Ethernet over InfiniBand protocol, 3-4
Ethernet ports, 2-2, 3-1
Ethernet switch
 configuration, 8-4
 configuring host name, 8-6
 IP routing, 8-5
 port speed, 8-5
 replacement parts, B-9
 replacement switch, 1-4
 setting IP address, 4-6
 setting remote password, 8-6
 spare parts, 1-2
 time settings, 8-7
 validating and saving configuration, 8-9
Europe
 PDU specifications, 2-6, 2-7
Exadata Database Machine
 See Oracle Exadata Database Machine
ext4 file system, 12-17
external filer configuration, 11-8

F

Fabric Management, 8-13
factory image reinstallation, 8-25
factory patches
 installing, 11-7

factory software image, 8-20
failover connections, 3-1, 3-4
failover switch connections, 3-1, 3-2
Fibre Channel over Ethernet, 1-5
filler panels, 2-10
firewall configuration, 2-12
firmware upgrades
 KVM switch, 8-3
 PDUs, 8-15
 restrictions, 1-4
first server location, 8-21
floor cut dimensions, 5-7
floor load, 2-3
floor tiles, 2-10
flooring
 raised, 2-3
 raised pattern, 2-2
forklift, 2-2
FRUs, B-1

G

gateway addresses
 factory settings, 3-5
gateway configuration
 Ethernet switch, 8-5
gateway switches
 assigning to VNICs, 3-4
 cabling, 3-3
 configuring, 8-10
 connecting to Oracle ILOM, 6-5
 replacing, 12-26
 using splitter cables, 3-4
 See also InfiniBand switches
getmaster command, 9-3, 9-4
ground cable, 5-9
grounding resistance, 5-1
grounding to earth, 2-8
groups, 11-8
GRUB master boot records, 12-18

H

Hadoop
 cluster configuration, 1-5, 11-5
 installing, 11-9
 See also CDH
hadoop group, 11-8
Hadoop ports, 3-6
Hadoop services, starting, 11-9
hardware diagnostics, 6-1
hardware failures
 automatic reporting, 10-1
hardware profiles
 acquiring from servers, 8-19
HDFS disk
 configuring, 12-22
HDFS partition, 12-17
HDFS partitions, 12-7
 dismounting, 12-8

hdfs user, 11-8
health checks
 InfiniBand switches, 8-12
high-voltage PDUs, 2-6, 2-7, 8-16, 8-17
Hive service, starting, 11-9
hosts
 factory IP addresses, 3-5
hot swap repair category, B-1
HS repair category, B-1
humidity range, 2-8, 2-9

I

ibdiagnet command, 12-28
iblinkinfo command, 8-33
ibqueryerrors command, 12-28
ibstatus command, 8-18
ifcfg-bondib0 file, 12-24
ILOM
 See Oracle ILOM
ILOM host names, 4-5
ilom-admin
 default password, 8-10
ilom-admin user, 6-2
ilom-operator user, 6-2
image version
 downgrading, 11-6
 upgrading, 11-5
imagehistory utility, 8-33
imageinfo utility, 8-20, 8-33, 11-5
incline, maximum, 2-3
InfiniBand configuration, 8-10
InfiniBand network
 factory IP addresses, 3-5
 host names, 4-5
 verifying operations, 12-27
InfiniBand network configuration, 8-20
InfiniBand network interface card, 3-1
InfiniBand private network, 3-2
InfiniBand Subnet Managers, 12-29
InfiniBand switches
 changing IP addresses, 12-24
 checking status, 8-18
 configuring, 8-10
 health checks, 8-12
 LED color codes, B-3
 on Oracle Exadata Database Machine, 8-14
 on Oracle Exalogic, 8-14
 priority settings, 8-14
 replacing, 12-26
 setting IP addresses, 4-6
 setting time zone, 8-12
 spare parts, B-7, B-8
infrastructure repair category, B-1
installation errors, 11-4
Installation Template, 4-3
InstallBDAstep, 11-9
InstallHadoop step, 11-9
installing
 base image, 8-25

third-party software, 1-4
installing patches, 11-7
installing the software
 See also Mammoth Utility, 11-2
International Electrotechnical Commission (IEC), 2-8
IP address requirements, 2-2
IP addresses
 changing for InfiniBand switches, 12-24
 checking for conflicts, 4-4
 custom settings, 4-6
 factory settings, 3-5
IP routing, 8-5
ipmitool command, 8-18
IR repair category, B-1

J

Japan
 PDU specifications, 2-4, 2-5
jumper brackets, 5-3

K

KMM tray replacement parts, B-9
KVM switch
 accessing servers, 8-3
 cabling, C-2
 configuration procedures, 8-1
 configuring server names, 8-17
 connecting to the network, 8-2
 firmware upgrades, 8-3
 naming, 8-2
 replacement parts, B-9
 setting IP address, 4-6
 starting, 8-2
KVM switch, MAC address, 8-2

L

leaf switches
 configuring, 8-10
 See also InfiniBand switches
LED color codes, B-2, B-3
leveling bolts, 5-4
leveling feet, 5-5, 5-6
licenses, location, 11-7
licensing restrictions, 1-4
lift gate, 2-2
lightning arresters, 2-3
lights-out management, 6-1
listlinkup command, 8-34
loading dock, 2-2
locating
 first server, 8-21
location in the rack, C-2
locks, panel, 5-3
log files
 Mammoth Utility, 11-4
logical volume for Oracle NoSQL Database, 11-7
low-voltage PDUs, 2-4, 2-5, 8-16

M

MAC address of KVM switch, 8-2
maintenance access requirements, 2-3
Mammoth Reconfiguration Utility
 running, 11-10
 syntax, 11-10
Mammoth Utility
 description, 11-1
 execution steps, 11-6
 location, 11-3
 log files, 11-4
 reasons for failure, 11-5
 running, 11-2
 synchronizing clocks, 8-12
 syntax, 11-3
mammoth-rack_name.params file, 4-3
mammoth-saved.params configuration file, 11-10
management network
 cabling requirements, 2-11
 configuration, 8-20
 connecting the KVM, 8-2
 connections, 2-11
 factory IP addresses, 3-5
 host names, 4-5
 ports, 3-1
management tools
 Oracle ILOM, 6-2
management utilities, 1-3
mapred user, 11-8
MegaCli64 command, 8-19, 12-4
Microsoft Office Excel, 4-2
Middle East
 PDU specifications, 2-6, 2-7
minimum ceiling height, 2-2
monitoring PDUs, 8-16
mount point
 logical volume for Oracle NoSQL Database, 11-7
mount points, 12-7
mounting a device, 8-21
mounting brackets, 5-4, 5-6
 installation instructions, 5-7
mounting tool, 5-3
multiple-rack clusters
 client network connections, 3-5
multirack cabling
 procedures, 9-1
 three to eight, 9-3
 two racks, 9-2
multirack software configuration, 11-4
My Oracle Support, 4-2
 downloading PDU firmware from, 8-15
 email account, 10-2
My Oracle Support Information Center, 1-4, 4-2, 8-6, 8-13, 8-25, 11-2, 11-3, 11-6
MySQL Database installation, 11-9

N

name node
 external backup, 4-7

naming conventions, 4-5
NET0, 3-2
network
 connecting the KVM, 8-2
network configuration, 8-20
network configuration worksheet, 4-5
network connections, 2-11, 8-23
network diagram, 3-3, 6-2
network errors
 checking for, 4-4
network interfaces
 InfiniBand card, 3-1
network operations
 verifying InfiniBand, 12-27
network registration, 2-12
network requirements, 2-2
Network Time Protocol, 12-30
network topology, 12-29
networks
 administrative, 3-2
 client access, 2-11, 3-2
 descriptions, 3-2
 factory IP addresses, 3-5
 InfiniBand, 3-2
networksetup scripts, 8-22, 8-23
nodes
 location in the rack, C-2
noise, electronic, 5-2
North America
 PDU specifications, 2-4, 2-5
NoSQL database
 See Oracle NoSQL Database
NTP servers, 12-30
 configuring, 8-8
 setting on InfiniBand switches, 8-12
ntp.conf file, 12-30
nuts, 5-3

O

OASM package
 installing, 10-3
ODI
 See Oracle Data Integrator
oinstall group, 4-8, 11-8
operating cautions and warnings, 12-3
operating system
 configuring a new disk, 12-13
operating system disks, 12-13
Oracle Big Data Appliance
 general description, 1-1
 rack components, 1-1
 server components, 1-2
Oracle Big Data Appliance Configuration Worksheets, 2-12
Oracle Big Data Connectors
 installation, 11-5, 11-9
 installing server components, 4-7
 software installation, 11-5
Oracle Data Integrator, 4-8

installing the agent, 4-7
Oracle Exadata Database Machine, 4-8
InfiniBand firmware version, 8-14
network identifiers, 4-8
Oracle Exalogic
InfiniBand firmware version, 8-14
Oracle ILOM, 6-1
backing up configuration, 12-25
connecting to CLI, 6-3
connecting to web interface, 6-3
description, 6-1
Ethernet connections, 6-3
factory IP addresses, 3-5
opening serial connections, 6-5
restoring configuration, 12-26
user interfaces, 6-2
Oracle ILOM port, 3-1
Oracle Inventory Group, 4-8
Oracle Linux, 1-1
Oracle NoSQL Database, 11-7
creating logical volumes, 11-7
disk configuration, 12-19
installation, 4-7, 11-9
logical volume, 11-7
removing logical partitions, 12-19
stopping the service, 12-21
Oracle NoSQL Database disks, 12-19
Oracle Open Office Calc, 4-2
oracle user name, 4-8
orange cables, C-1

P

pallet jack, 2-2
panel locks, 5-3
parallel execution of commands, 7-1
partitioning
operating system disk, 12-14
partitioning a new disk, 12-14
passwordless SSH, 7-1
passwords (default)
Ethernet switch, 8-5
ilom-admin, 8-10
PDU metering unit, 8-15
servers, 8-18
patches
installing, 11-7
PatchFactoryImage step, 11-7
path to data center, 2-3
PDAs, location in the rack, C-2
PDU configuration, 8-15
PDU metering unit
default password, 8-15
PDU requirements, 2-4
PDU specifications
Americas, Japan, Taiwan, 2-4, 2-5
Europe, Middle East, Africa, Asia Pacific, 2-6, 2-7
PDU thresholds, 8-16
PDUs
circuit breakers, 5-13

firmware upgrade, 8-15
network connections, 8-15
powering on, 5-13
replacement parts, B-10
setting IP addresses, 4-7
single-phase, 5-2, 5-13
perforated tiles, 2-10
photocopiers, 5-2
physical drive configuration, 12-5
port assignments, 3-6
port speed
Ethernet switch, 8-5
ports
ASR Manager, 4-8
connecting to 40 GbE, 3-5
power connectors, 2-4
power cord length requirements, 2-8
power cords, 2-8
connecting to rack, 5-10
removing cable ties, 5-11
routing, 5-11
power LEDs
locations, 5-13
power off switch, 12-3
power outlet grounding, 2-8
power receptacle location, 2-8
power requirements, 2-1, 2-3
power supplies, uninterruptable, 2-8
powering off
emergency, 12-3
sequence, 12-2
powering on, 5-13, 12-2
preinstallation tasks, 2-11
preinstall-checkip script, 3-3
description, 4-3
running, 4-4
PreInstallChecks step, 11-6
private network
verifying operations, 12-27
problem reporting, 10-1
Puppet utility, 11-7

Q

QSFP ports, splitting, 3-4

R

rack components, 1-1
rack configuration file, 11-5
rack dimensions, 2-2
rack layout illustration, C-2
rack names, 4-5
rack weight, 2-1
rack-hosts-infiniband file, 8-18
racks
adding another, 9-1
connecting power cords, 5-10
inspecting, 5-10
installing

Oracle Big Data Appliance software, 11-1
moving instructions, 5-5
reinstalling base image, 8-25
RAID arrays
 repairing logical, 12-15
RAID configuration, 8-19
RAID logical drives, 12-6
raid0 logical drives, 12-6
red cables, C-1
registering ASR Manager, 10-4
reimaging the servers, 8-25
reinstalling the software, 11-2
relative humidity, 2-9, 2-10
remove-root-ssh script, 7-2
repair categories, B-1
replacement parts, B-1
replacing spine switch, 12-27
restart validation, 5-14
restarting servers, 12-9
restrictions on use, 1-4
root user
 passwordless SSH, 7-1
route requirements, 2-3
running-config output, 8-9

S

safety guidelines
 installation, 5-1
SCAN addresses, 3-2
scp file copy utility, 7-2
screwdrivers, 5-4
screws, 5-3
sensor monitoring, 6-1
server components, 1-2
server configuration, 8-17
server disks
 replacing, 12-6
 software configuration, 12-10
servers
 boot order, 8-19
 connecting to Oracle ILOM, 6-5
 default passwords, 8-18
 factory IP addresses, 3-5
 LED color codes, B-2
 location in the rack, C-2
 monitoring, 6-1
 setting IP addresses, 4-6
 spare parts, B-4
 upgrading base image, 8-25
service locations, restrictions on moving, 1-4
service tags bundle, 10-4
services, starting, 11-9
setsmpriority command, 8-14, 9-2, 9-3, 9-4
SetupASR step, 11-9
SetupInstall step, 11-6
SetupMountPoints step, 11-8
SetupMySQL step, 11-9
SetupPuppet step, 11-7
setup-root-ssh script, 7-1, 8-32

SetupSSHroot step, 11-6
shipping crate dimensions, 5-2
shipping dimensions, 2-2
shipping kit, 5-3
shipping weight, 2-2
show command, 8-9
showunhealthy command, 8-13
showusb command, 8-21
showvlan command, 8-34
showvnics command, 8-34
shutdown command, 12-2
single-phase PDUs, 2-4, 2-6, 5-2, 5-13, 8-16
site readiness tasks, 2-11
slot numbers, 12-7
SMTP alerts, 10-1
SMTP configuration, 4-8
SNMP alerts, 10-1
software configuration worksheet, 4-7
software image
 downgrading, 11-6
 upgrading, 11-5
 verifying, 8-20
software installation
 multiple racks, 11-4
 upgrading, 11-3
 using Mammoth, 11-1
software installation errors, 11-4
software installation,
 third-party, 1-4
software port assignments, 3-6
software profiles
 acquiring from servers, 8-19
software source code, third party, 11-7
sound levels, 2-1
South America
 PDU specifications, 2-4, 2-5
spanning tree
 disabling uplink port, 8-9
spare parts, 5-4
spares kit, 1-2
spine switch
 configuring, 8-10
spine switches
 cabling, D-2
 replacing, 12-26, 12-27
splitter cables, 3-4
spreadsheet
 configuration utility, 4-1
spreadsheet programs, 4-2
spring nuts, 5-3
SSH keys
 checking rack distribution, 8-18
stabilizing feet, 5-5
stabilizing the rack, 5-6
StartHadoopServices step, 11-9
StartHiveService step, 11-9
starting KVM, 8-2
startup order check, 8-19
state file, 11-5
static discharges, 2-9

static electricity, 5-2
Subnet Manager Master, 9-4
Subnet Managers, 12-29, 12-30
subnet mask requirements, 12-24
subnet masks
 factory settings, 3-5
subnet requirements, 12-24
Sun Network 10 GbE Switch 72p, 3-4
SUNWswasr package, 10-4
support accounts, 10-2
swap partition
 restoring, 12-18
switches, C-2
 LED color codes, B-3
 monitoring, 6-1
 naming conventions, 4-6
 replacement parts, B-9
 replacing, 12-26
 setting host names, 8-11
 setting IP addresses, 4-6
 spare parts, B-7, B-8
 third-party, 1-5

T

Taiwan
 PDU specifications, 2-4, 2-5
telnet access
 configuring Ethernet switch, 8-6
temperature
 acceptable range, 2-9
 operating range, 2-8
third-party licenses, location, 11-7
third-party software, 1-4
third-party software source code, 11-7
third-party switches
 restrictions on use, 1-5
three-phase PDUs, 2-5, 2-7, 8-16, 8-17
thresholds for PDUs, 8-16
time zone
 configuring Ethernet switch, 8-7
 configuring InfiniBand switches, 8-12
 setting summer time, 8-8
top clearance (rack), 2-2

U

umount command, 12-8
unpacking instructions, 5-4
UpdateEtcHosts step, 11-6
upgrading base image, 8-25
upgrading the factory image, 11-5
upgrading the software, 11-3
UPS, 2-8
USB devices
 restrictions on use, 1-5
USB flash drive, 8-21
users, 11-8

V

validation checks after reboot, 5-14
ventilation requirements, 2-9
verifying
 database server configuration, 12-4
verifying a disk configuration, 12-23
version command, 8-13
VIP addresses, 3-2
virtual drive configuration, 12-5
VLANs
 configuration example, 8-5
 default configuration, 3-1
VNICS
 assigning to gateway switches, 3-4
 default configuration, 3-1

W

warnings, 12-3
warranty restrictions, 1-4
washers, 5-7
weight
 rack, 2-1
 shipping, 2-2
wrenches, 5-3, 5-8
wrist strap, 5-4