

Oracle® Big Data Appliance

Owner's Guide



Release 4 (4.8)

E85370-07

August 2018

ORACLE®

Copyright © 2011, 2018, Oracle and/or its affiliates. All rights reserved.

Primary Author: Frederick Kush

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

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

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

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

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

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

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

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

Contents

Preface

Audience	xviii
Documentation Accessibility	xviii
Related Documentation	xviii
Conventions	xix
Backus-Naur Form Syntax	xx
Changes in Oracle Big Data Appliance Release 4 (4.8)	xx
Changes in Previous Releases of Oracle Big Data Appliance	xxiv

Part I Preinstallation

1 Introduction to Oracle Big Data Appliance

1.1 About Oracle Big Data Appliance	1-1
1.2 Oracle Big Data Appliance Rack Components	1-1
1.3 Server Components	1-2
1.4 Spare Parts Kit Components	1-4
1.5 Oracle Big Data Appliance Management Software	1-6
1.6 Supportability Guidelines	1-6

2 Site Requirements for Oracle Big Data Appliance

2.1 General Environmental Requirements	2-1
2.2 Space Requirements	2-4
2.2.1 Receiving and Unpacking Requirements	2-4
2.2.2 Maintenance Access Requirements	2-5
2.3 Flooring Requirements	2-6
2.4 Electrical Power Requirements	2-6
2.4.1 PDUs for Oracle Big Data Appliance	2-6
2.4.1.1 Single-Phase PDUs for North America, South America, Japan, and Taiwan	2-7
2.4.1.2 Three-Phase PDUs for North America, South America, Japan, and Taiwan	2-7

2.4.1.3	Single-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific	2-8
2.4.1.4	Three-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific	2-9
2.4.2	Facility Power Requirements	2-10
2.4.3	Circuit Breaker Requirements	2-11
2.4.4	Grounding Guidelines	2-11
2.5	Temperature and Humidity Requirements	2-11
2.6	Ventilation and Cooling Requirements	2-13
2.7	Network Connection Requirements	2-14
2.8	Ensuring That the Site Is Ready	2-14

3 Understanding the Network Requirements

3.1	Overview of Network Requirements	3-1
3.1.1	Default Network Connections	3-1
3.1.2	Minimum Requirements for the Networks	3-1
3.1.3	Network Diagram	3-3
3.2	Cabling the Client Network	3-5
3.2.1	How the Servers Connect to the Gateway Switches	3-5
3.2.2	Using Splitter Cables for Connecting to the 40 Gbps Physical Ports	3-7
3.2.3	What About Data Centers Without a 10 GbE Infrastructure?	3-7
3.2.4	Guidelines for Connecting to the Gateway Switches	3-7
3.2.4.1	Configuring for LACP/LAG Connections	3-8
3.2.5	Configuring Multiple Racks into a Single Hadoop Cluster	3-9
3.3	Factory Network Settings	3-9
3.4	Port Assignments for Oracle Big Data Appliance Software	3-10
3.5	Preparing the Network for Oracle Big Data Appliance	3-11

4 Using Oracle Big Data Appliance Configuration Generation Utility

4.1	Overview of Oracle Big Data Appliance Configuration Generation Utility	4-1
4.2	Generating the Configuration Files	4-3
4.3	About the Configuration Files	4-3
4.4	Validating the Network Settings	4-5
4.5	Customer Details Page	4-5
4.5.1	Using Standardized Host Names	4-6
4.5.2	Using Customized Host Names	4-6
4.5.3	Naming Multirack Clusters	4-6
4.5.4	Connecting to Oracle Big Data Appliance Over the Networks	4-6
4.6	Hardware Selection Page	4-7
4.7	Rack Details	4-7

4.8	Networking Page	4-8
4.9	Administration Network Page	4-8
4.10	General Network Properties Page	4-9
4.11	Review and Edit Details Page	4-10
4.12	Define Clusters Page	4-11
4.13	Cluster Page	4-12
4.13.1	A New Oracle NoSQL Database Cluster	4-12
4.13.1.1	User and Groups	4-12
4.13.1.2	Client Network	4-13
4.13.1.3	InfiniBand Network	4-14
4.13.1.4	MIT Kerberos	4-15
4.13.1.5	Installed Components	4-15
4.13.1.6	Oracle NoSQL Configuration	4-16
4.13.2	A New CDH Cluster	4-16
4.13.2.1	User/Groups	4-17
4.13.2.2	Client Network	4-17
4.13.2.3	InfiniBand Network	4-18
4.13.2.4	Big Data SQL	4-19
4.13.2.5	Big Data Connectors	4-20
4.13.2.6	MIT Kerberos	4-20
4.13.2.7	Active Directory Kerberos	4-21
4.13.2.8	HDFS Transparent Encryption	4-22
4.13.2.9	Audit Vault	4-22
4.13.2.10	Auto Service Request	4-23
4.13.2.11	Enterprise Management Cloud Control	4-24
4.13.2.12	Email Alerting	4-24
4.13.3	Adding to an Existing Cluster	4-25
4.14	Client Network Config Summary	4-25
4.15	Complete Page	4-25

5 Setting Up Auto Service Request

5.1	Understanding Auto Service Request	5-1
5.2	Getting Ready to Install ASR	5-2
5.3	Installing ASR Manager	5-3
5.4	Verifying the ASR Manager Installation	5-3
5.5	Configuring the Trap Destinations on Oracle Big Data Appliance	5-4
5.6	Verifying ASR Assets	5-4
5.7	Validating Auto Service Request on Oracle Big Data Appliance	5-5
5.8	Troubleshooting ASR	5-6

Part II Hardware Installation and Configuration

6 Installing Oracle Big Data Appliance at the Site

6.1	Reviewing Safety Guidelines	6-1
6.2	Unpacking Oracle Big Data Appliance	6-2
6.2.1	Contents of the Shipping Kit	6-3
6.2.2	Removing Oracle Big Data Appliance from the Shipping Crate	6-4
6.3	Placing Oracle Big Data Appliance in Its Allocated Space	6-6
6.3.1	Moving Oracle Big Data Appliance	6-6
6.3.2	Securing an Oracle Big Data Appliance Rack	6-7
6.3.2.1	Secure the Oracle Big Data Appliance Rack with Leveling Feet	6-8
6.3.3	Attaching a Ground Cable (Optional)	6-8
6.4	Powering On the System for the First Time	6-9
6.4.1	Inspecting the Oracle Big Data Appliance Rack After It Is in Place	6-9
6.4.2	Connecting Power Cords	6-10
6.4.3	Powering On Oracle Big Data Appliance	6-12

7 Configuring an Oracle Big Data Appliance Full Rack or Starter Rack

7.1	Configuring the KVM Switch	7-1
7.1.1	Starting the KVM Switch	7-2
7.1.2	Connecting the KVM Switch to the Management Network	7-2
7.1.3	Configuring the KVM Switch to Access the Servers	7-3
7.1.4	Accessing a Server by Using the KVM Switch	7-3
7.2	Configuring the Cisco Ethernet Switch	7-3
7.2.1	Scope of the Configuration	7-4
7.2.2	Prerequisites for Configuring the Ethernet Switch	7-4
7.2.3	Configuring the Ethernet Switch on the Customer Network	7-4
7.2.4	Setting Up Telnet Access on the Ethernet Switch	7-7
7.2.5	Configuring SSH on the Ethernet Switch	7-8
7.2.6	Setting the Clock and Time Zone on the Ethernet Switch	7-9
7.2.7	Configuring the NTP Servers for the Ethernet Switch	7-10
7.2.8	Disabling the Spanning Tree	7-11
7.2.9	Verifying the Ethernet Configuration	7-11
7.3	Configuring the InfiniBand Leaf and Spine Switches	7-12
7.3.1	Configuring an InfiniBand Switch	7-12
7.3.2	Setting the Time Zone and Clock on an InfiniBand Switch	7-14
7.3.3	Checking the Health of an InfiniBand Switch	7-15
7.3.4	Checking the Sun Network QDR InfiniBand Gateway Switches	7-17
7.4	Configuring the Power Distribution Units	7-18

7.4.1	Connecting the PDUs to the Network	7-18
7.4.2	Updating the PDU Firmware Version	7-19
7.5	Configuring the Oracle Big Data Appliance Servers	7-20
7.6	Configuring the Network	7-23
7.6.1	Copying the Configuration Files to Oracle Big Data Appliance	7-23
7.6.2	Starting the Network Configuration	7-24
7.6.3	Connecting to the Administrative Network	7-26
7.6.4	Connecting to the Client Network	7-27
7.6.5	Completing the Network Configuration	7-27
7.6.6	Checking the Virtual Interface Cards (VNICs)	7-29
7.6.7	Verifying Access to the Network Services	7-30
7.7	Configuring Multiple Oracle Big Data Appliance Racks	7-31
7.7.1	Configuring and Cabling Multiple Racks	7-31
7.7.2	Validating a Multirack InfiniBand Network	7-33

8 Expanding an Oracle Big Data Appliance Starter Rack

8.1	About Expanding an Oracle Big Data Appliance Starter Rack	8-1
8.1.1	Checking the Contents of a High Capacity Node plus Infiniband Infrastructure	8-2
8.1.2	Checking the Status of a Starter Rack	8-2
8.1.3	Expanding an Oracle Big Data Appliance X4-2 Starter Rack	8-2
8.2	Installing New Servers to an Oracle Big Data Appliance Starter Rack	8-3
8.3	Configuring the New Oracle Big Data Appliance Servers	8-4
8.4	Setting Up the Network	8-5
8.4.1	Copying the Configuration Files	8-5
8.4.2	Configuring the Network	8-6
8.4.3	Verifying the Network Configuration	8-7
8.5	Installing Oracle Big Data Appliance Software on the New Servers	8-7

9 Connecting Multiple Oracle Big Data Appliance Racks

9.1	Extending a Rack by Adding Another Rack	9-1
9.2	Prerequisites for Multirack Cabling	9-2
9.3	Cabling Two Racks Together	9-2
9.4	Cabling Several Racks Together	9-3

Part III Software Installation and Configuration

10 Installing the Oracle Big Data Appliance Software

10.1	About the Mammoth Utility	10-1
10.1.1	About Rolling Upgrades	10-2
10.2	Installation Prerequisites	10-3
10.3	Downloading the Mammoth Software Deployment Bundle	10-4
10.4	Installing the Software on a New Rack	10-6
10.4.1	Installing the Software	10-6
10.5	What If an Error Occurs During the Installation?	10-8
10.6	Adding Servers to a Cluster	10-8
10.7	Upgrading the Software on Oracle Big Data Appliance	10-10
10.7.1	About the Operating System Versions	10-11
10.7.2	Upgrading the Software	10-11
10.7.2.1	Prerequisites	10-11
10.7.2.2	Upgrading to the Current Software Version	10-11
10.7.2.3	Performing Parcel Upgrades Outside of the One-Off Patch Mechanism	10-13
10.8	Changing the Configuration of Optional Software	10-13
10.8.1	Changing Support for Oracle Big Data Connectors	10-13
10.8.1.1	Adding Oracle Big Data Connectors	10-13
10.8.1.2	Removing Oracle Big Data Connectors	10-14
10.8.2	Adding Support for Auto Service Request	10-16
10.8.3	Adding Support for Oracle Enterprise Manager Cloud Control	10-16
10.8.4	Adding Support for Oracle Audit Vault and Database Firewall	10-17
10.8.5	Adding Kerberos Authentication	10-18
10.9	Reinstalling the Base Image	10-19
10.9.1	Reimaging a Single Oracle Big Data Appliance Server	10-19
10.9.2	Reimaging an Oracle Big Data Appliance Rack	10-21
10.9.3	Reimaging an Oracle Big Data Appliance Cluster	10-23
10.10	Installing a One-Off Patch	10-26
10.11	Mammoth Software Installation and Configuration Utility	10-27
10.11.1	Mammoth Options	10-28
10.11.2	Mammoth Installation Steps	10-29
10.12	Oracle Big Data Appliance Base Imaging Utilities	10-32
10.12.1	makebdaimage	10-32
10.12.2	reimagecluster	10-33
10.12.3	reimagerack	10-34

Part IV Maintenance

11 Maintaining Oracle Big Data Appliance

11.1	Monitoring the Ambient Temperature of Servers	11-1
11.2	Powering On and Off Oracle Big Data Appliance	11-2
11.2.1	Nonemergency Power Procedures	11-2
11.2.1.1	Powering On Oracle Big Data Appliance	11-2
11.2.1.2	Powering Off Oracle Big Data Appliance	11-2
11.2.1.3	Powering On and Off Network Switches	11-3
11.2.2	Emergency Power-Off Considerations	11-3
11.2.2.1	Emergency Power-Off Procedure	11-3
11.2.2.2	Emergency Power-Off Switch	11-3
11.2.3	Cautions and Warnings	11-4
11.3	Adding Memory to the Servers	11-4
11.3.1	Adding Memory to an Oracle Server X6-2L	11-4
11.3.2	Adding Memory to an Oracle Server X5-2L, Sun Server X4-2L, or Sun Server X3-2L	11-5
11.3.3	Adding Memory to Sun Fire X4270 M2 Servers	11-5
11.4	Maintaining the Physical Disks of Servers	11-6
11.4.1	Verifying the Server Configuration	11-6
11.4.1.1	Verifying Disk Controller Configuration	11-6
11.4.1.2	Verifying Virtual Drive Configuration	11-6
11.4.1.3	Verifying Physical Drive Configuration	11-7
11.5	Maintaining the InfiniBand Network	11-7
11.5.1	Replacing a Failed InfiniBand Switch	11-7
11.5.2	Verifying InfiniBand Network Operation	11-9
11.5.3	Understanding the Network Subnet Manager Master	11-11
11.6	Changing the Number of Connections to a Gateway Switch	11-12
11.7	Changing the NTP Servers	11-12
11.8	Monitoring the PDU Current	11-13
11.9	Node Migration	11-15

12 Replacing a Server Disk

12.1	Overview of the Disk Replacement Process	12-1
12.2	About Disk Drive Identifiers	12-2
12.2.1	Standard Disk Drive Mappings	12-2
12.2.2	Standard Mount Points	12-3
12.2.3	Obtaining the Physical Slot Number of a Disk Drive	12-3
12.3	Prerequisites for Replacing a Failing Disk	12-4
12.4	What If a Server Fails to Restart?	12-6
12.5	Replacing a Disk Drive	12-7
12.6	Correcting a Mounted Partitions Error	12-9

12.7	Identifying the Function of a Disk Drive	12-10
12.8	Configuring an Operating System Disk	12-11
12.8.1	Partitioning the Operating System Disk	12-11
12.8.1.1	Partitioning for Oracle Linux 6	12-11
12.8.1.2	Partitioning for Oracle Linux 5	12-13
12.8.2	Repairing the RAID Arrays	12-14
12.8.3	Formatting the HDFS Partition of an Operating System Disk	12-16
12.8.4	Restoring the Swap Partition	12-16
12.8.5	Restoring the GRUB Master Boot Records and HBA Boot Order	12-17
12.9	Configuring an HDFS or Oracle NoSQL Database Disk	12-19
12.10	Verifying the Disk Configuration	12-21

13 Oracle Big Data Appliance Utilities

13.1.1	bdacheckcluster	13-2
13.1.2	bdacheckhw	13-3
13.1.3	bdacheckib	13-3
13.1.4	bdachecknet	13-5
13.1.5	bdachecksw	13-5
13.1.6	bdaclustersynctime	13-6
13.1.7	bdacli	13-7
13.1.8	bdaclustersynctime	13-16
13.1.9	bdadiag	13-17
13.1.10	bdadiagcluster	13-19
13.1.11	bdagetclientconfig	13-20
13.1.12	bdaid	13-21
13.1.13	bdaibstatus	13-21
13.1.14	bdaimagevalidate	13-22
13.1.15	bdanodedesc	13-23
13.1.16	bdaredoclientnet	13-23
13.1.17	bdaserials	13-25
13.1.18	bdaswapoff	13-26
13.1.19	bdaswapon	13-26
13.1.20	bdaupdatefw	13-27
13.1.21	iblinkinfo	13-28
13.1.22	imagehistory	13-29
13.1.23	imageinfo	13-29
13.1.24	listlinkup	13-30
13.1.25	mount_hadoop_dirs	13-30
13.1.26	remove-root-ssh	13-30
13.1.27	reset-boot-order	13-31

13.1.28	setup-root-ssh	13-32
13.1.29	showusb	13-33
13.1.30	showvlan	13-34
13.1.31	showwnics	13-34

14 Executing Commands Across a Cluster Using the dcli Utility

14.1	Overview of the dcli Utility	14-1
14.1.1	Setting Up Passwordless SSH	14-1
14.1.2	Basic Use of dcli	14-2
14.1.2.1	Getting Help	14-2
14.1.2.2	Identifying the Target Servers	14-2
14.1.2.3	Specifying the Commands	14-2
14.1.2.4	Controlling the Output Levels	14-3
14.2	dcli Syntax	14-3
14.3	dcli Return Values	14-5
14.4	dcli Examples	14-5

15 Using Oracle Integrated Lights Out Manager

15.1	Oracle ILOM Overview	15-1
15.1.1	Oracle ILOM Interfaces	15-2
15.1.2	Oracle ILOM Users	15-2
15.2	Administrative Network Diagram	15-2
15.3	Oracle ILOM IP Addresses for Oracle Big Data Appliance Components	15-3
15.4	Connecting to Oracle ILOM Using the Network	15-3
15.4.1	Connecting to the Command-Line Interface	15-4
15.4.2	Connecting to the Web GUI	15-4
15.5	Connecting to Oracle ILOM Using a Serial Connection	15-5
15.5.1	Connecting to the Oracle ILOM of a Server	15-5
15.5.2	Connecting to the Oracle ILOM of an InfiniBand Switch	15-6
15.6	Backing Up and Restoring Oracle ILOM Settings	15-7
15.6.1	Backing Up the Oracle ILOM Configuration Settings	15-7
15.6.2	Restoring Oracle ILOM Configuration Settings	15-8

Part V Appendixes

A Error Messages for Oracle Big Data Appliance

B Replacement Units

B.1	Understanding Repair Categories	B-1
B.2	Disk Controller Batteries	B-2
B.3	LED Status Descriptions	B-2
B.3.1	Sun Server LEDs	B-3
B.3.2	Sun InfiniBand Switch LEDs	B-4
B.4	Parts for Oracle Big Data Appliance Servers	B-5
B.4.1	Oracle Server X6-2L and X5-2L Parts	B-6
B.4.2	Sun Server X4-2L Parts	B-9
B.4.3	Sun Server X3-2L Parts	B-13
B.4.4	Sun Fire X4270 M2 Parts	B-17
B.5	Parts for Sun Datacenter InfiniBand Switch 36	B-22
B.6	Parts for Sun Network QDR InfiniBand Gateway Switch	B-22
B.7	Parts for the Cisco Ethernet Switch	B-23
B.8	Parts for Power Distribution Units	B-24
B.9	Cables and Cable Accessories for an Oracle Big Data Appliance Rack	B-25

C In-Rack Cabling Tables for Oracle Big Data Appliance X6–2 and X5-2

C.1	Cable Color Coding	C-1
C.2	Oracle Big Data Appliance Rack Layout	C-1
C.3	Administrative Gigabit Ethernet Cable Connections	C-4
C.4	Oracle Integrated Lights Out Manager Cable Connections	C-5
C.5	Single-Phase Power Distribution Unit Cable Connections	C-6
C.6	Three-Phase Power Distribution Unit Cable Connections	C-7
C.7	In-Rack InfiniBand Switch-to-Switch Cable Connections	C-8
C.8	In-Rack InfiniBand Switch-to-Server Cable Connections	C-9

D In-Rack Cabling Tables for Oracle Big Data Appliance X4-2 and X3-2

D.1	Cable Color Coding	D-1
D.2	Oracle Big Data Appliance Rack Layout	D-1
D.3	Administrative Gigabit Ethernet Cable Connections	D-4
D.4	Oracle Integrated Lights Out Manager Cable Connections	D-5
D.5	Single-Phase Power Distribution Unit Cable Connections	D-6
D.6	Three-Phase Power Distribution Unit Cable Connections	D-7
D.7	In-Rack InfiniBand Switch-to-Switch Cable Connections	D-8

E In-Rack Cabling Tables for Oracle Big Data Appliance Sun Fire X4270 M2

E.1	Cable Color Coding	E-1
E.2	Oracle Big Data Appliance Rack Layout	E-1
E.3	KVM Network Cable Connections	E-2
E.4	Administrative Gigabit Ethernet Cable Connections	E-3
E.5	Oracle Integrated Lights Out Manager Cable Connections	E-4
E.6	Single-Phase Power Distribution Unit Cable Connections	E-5
E.7	Three-Phase Power Distribution Unit Cable Connections	E-6
E.8	In-Rack InfiniBand Switch-to-Switch Cable Connections	E-7
E.9	In-Rack InfiniBand Switch-to-Server Cable Connections	E-8

F Multirack Cabling Tables

F.1	Understanding Multirack Cabling	F-1
F.1.1	Cabling Together Oracle Exadata Database Machine and Oracle Big Data Appliance	F-2
F.1.2	Cabling the Switches in One Rack	F-2
F.1.3	Cabling the Switches in Two Racks	F-3
F.1.4	Cabling the Switches in Up to Eight Racks	F-4
F.2	Key to Cabling Table Abbreviations	F-4
F.3	Two-Rack Cabling	F-5
F.4	Three-Rack Cabling	F-6
F.5	Four-Rack Cabling	F-8
F.6	Five-Rack Cabling	F-10
F.7	Six-Rack Cabling	F-13
F.8	Seven-Rack Cabling	F-17
F.9	Eight-Rack Cabling	F-21

Index

List of Tables

1-1	Oracle Big Data Appliance Rack Components	1-2
1-2	Oracle Server X6-2L Server Components	1-2
1-3	Oracle Server X5-2 Server Components	1-2
1-4	Sun Server X4-2L Server Components	1-3
1-5	Sun Server X3-2L Server Components	1-3
1-6	Sun Fire X4270 M2 Server Components	1-4
1-7	Bundled Components and Accessories	1-4
1-8	Bundled Cables	1-5
1-9	In-Rack Spare Cables	1-5
1-10	InfiniBand Inter-Switch Cables and Spares Provided	1-5
1-11	Additional Cables to Purchase (Based on Number of Racks)	1-5
2-1	Environmental Requirements for Oracle Big Data Appliance X6-2 and X5-2	2-1
2-2	Environmental Requirements for Oracle Big Data Appliance X6-2 and X5-2 High Capacity Node plus Infiniband Infrastructure	2-2
2-3	Environmental Requirements for Oracle Big Data Appliance X4-2 and X3-2	2-3
2-4	Environmental Requirements for Oracle Big Data Appliance Sun Fire X4270 M2-Based Rack	2-3
2-5	Shipping Measurements	2-5
2-6	Access Route Requirements for Oracle Big Data Appliance	2-5
2-7	Maintenance Access Requirements for Oracle Big Data Appliance	2-5
2-8	Floor Load Requirements for Oracle Big Data Appliance	2-6
2-9	Low-Voltage 15 kVA Single-Phase PDUs	2-7
2-10	Low-Voltage 15 kVA Three-Phase PDUs	2-8
2-11	High-voltage 22 kVA Single Phase PDUs	2-8
2-12	High-Voltage 15 kVA Three-Phase PDUs	2-9
2-13	Temperature, Humidity, and Altitude Requirements	2-12
3-1	Server/Gateway Connections Example	3-6
3-2	Default IP Addresses for Oracle Big Data Appliance	3-9
3-3	Oracle Big Data Appliance Port Numbers	3-10
3-4	Port Access Required by the Oracle Enterprise Manager Plug-In	3-11
4-1	Customer Details Page	4-5
4-2	Hardware Selection Page	4-7
4-3	Rack Details Page	4-8
4-4	Administration Network Page	4-9
4-5	General Network Properties Page	4-9

4-6	Advanced Network Properties	4-10
4-7	Define Clusters Page	4-11
4-8	User and Groups for a New Oracle NoSQL Database Cluster	4-13
4-9	Client Network	4-13
4-10	InfiniBand Network Page	4-14
4-11	MIT Kerberos	4-15
4-12	Installed Components	4-16
4-13	Oracle NoSQL Configuration	4-16
4-14	User /Groups for a New CDH Cluster	4-17
4-15	Client Network	4-18
4-16	InfiniBand Network Page	4-19
4-17	Big Data SQL	4-19
4-18	Installed Components	4-20
4-19	MIT Kerberos	4-20
4-20	Active Directory Kerberos	4-21
4-21	HDFS Transparent Encryption	4-22
4-22	Audit Vault	4-23
4-23	Auto Service Request	4-23
4-24	Enterprise Manager Cloud Control	4-24
4-25	Email Alerting Page	4-25
8-1	A High Capacity Node plus Infiniband Infrastructure Components and Accessories	8-2
8-2	A High Capacity Node plus Infiniband Infrastructure Cables	8-2
11-1	Threshold Values for Single-Phase, Low-Voltage PDU	11-14
11-2	Threshold Values for Three-Phase, Low-Voltage PDU	11-14
11-3	Threshold Values for Single-Phase, High-Voltage PDU	11-14
11-4	Threshold Values for Three-Phase, High-Voltage PDU	11-15
12-1	Disk Drive Identifiers	12-3
12-2	Mount Points	12-3
13-1	bdacli admin_cluster parameters	13-7
13-2	bdacli enable or disable Service Parameters	13-8
13-3	BDACLI start, stop, and restart Parameters	13-10
13-4	BDACLI reset Scope Parameters	13-10
13-5	BDACLI reset Network Parameters	13-10
13-6	InfiniBand Parameters	13-11
13-7	Rack Parameters	13-12
13-8	General Cluster Parameters	13-13
13-9	Oracle Big Data Connectors Status Parameters	13-13

13-10	Cluster Network Parameters	13-14
13-11	Cluster Security Parameters	13-14
13-12	Server Parameters	13-15
13-13	One-Off Patch Parameters	13-16
14-1	dcli Options	14-4
B-1	Oracle Big Data Appliance Repair Categories	B-1
B-2	Sun Server LED Status Descriptions	B-3
B-3	Sun InfiniBand Switch LED Status Descriptions	B-4
B-4	Replacement Parts for Oracle Server X6-2L and X5-2L Storage Servers	B-7
B-5	Replacement Parts for a Sun Server X4-2L Server	B-10
B-6	Replacement Parts for a Sun Server X3-2L Server	B-13
B-7	Replacement Parts for Sun Fire X4270 M2 Servers	B-17
B-8	Replacement Parts for the Sun Datacenter InfiniBand Switch 36	B-22
B-9	Replacement Parts for the Sun Network QDR InfiniBand Gateway Switch	B-23
B-10	Replacement Parts for Cisco Catalyst Ethernet	B-23
B-11	Replacement Parts for Power Distribution Units	B-24
B-12	Replacement Cables for an Oracle Big Data Appliance Rack	B-26
C-1	Administrative Gigabit Ethernet Cabling	C-5
C-2	Oracle ILOM Cabling	C-5
C-3	Single-Phase PDU Cabling	C-6
C-4	Three-Phase PDU Cabling	C-7
C-5	InfiniBand Switch-to-Switch Cabling	C-8
C-6	InfiniBand Switch-to-Server Cabling	C-9
D-1	Administrative Gigabit Ethernet Cabling	D-5
D-2	Oracle ILOM Cabling	D-5
D-3	Single-Phase PDU Cabling	D-6
D-4	Three-Phase PDU Cabling	D-7
D-5	InfiniBand Switch-to-Switch Cabling	D-8
D-6	InfiniBand Switch-to-Server Cabling	D-9
E-1	KVM Cabling	E-3
E-2	Administrative Gigabit Ethernet Cabling	E-3
E-3	Oracle ILOM Cabling	E-4
E-4	Single-Phase PDU Cabling	E-5
E-5	Three-Phase PDU Cabling	E-6
E-6	InfiniBand Switch-to-Switch Cabling	E-7
E-7	InfiniBand Switch-to-Server Cabling	E-8
F-1	Leaf Switch Connections for the First Rack in a Two-Rack System	F-5

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

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](#) (page xviii)
- [Related Documentation](#) (page xviii)
- [Conventions](#) (page xix)
- [Backus-Naur Form Syntax](#) (page xx)

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

Related Documentation

Oracle Big Data Appliance Documentation Library

In the [Big Data Documentation Portal](#) of the Oracle Help Center, you can find a link to the complete Oracle Big Data Appliance library for your release of the product. The library includes the following core documents as well as documents for products that are used in conjunction with Oracle Big Data Appliance:

- *Oracle Big Data Appliance Owner's Guide* (this guide)
- *Oracle Big Data Appliance Software User's Guide*

- *Oracle Big Data Connectors User's Guide*
- *Oracle Big Data Appliance Safety and Compliance Guide*
- *Oracle Big Data Appliance Licensing Information User Manual*

 Note:

The *Oracle Big Data Appliance Licensing Information User Manual* is the consolidated reference for licensing information for Oracle and third-party software included in the Oracle Big Data Appliance product. Refer to this manual or contact Oracle Support if you have questions about licensing.

Documentation for Affiliated Products

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

<https://docs.oracle.com/en/bigdata/>

- Oracle Server X5-2L library:
http://docs.oracle.com/cd/E41033_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) 3.1 library:
http://docs.oracle.com/cd/E24707_01/index.html

Conventions

The following text conventions are used in this document:

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

Backus-Naur Form Syntax

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

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

Changes in Oracle Big Data Appliance Release 4 (4.8)

Release 4.8 includes following feature changes and software updates.

Software Updates

- CDH (Cloudera's Distribution including Apache Hadoop) 5.10.1
- CM (Cloudera Manager) 5.10.1
- Oracle Big Data Connectors 4.8
- Big Data SQL 3.1 (earlier releases are not supported on Oracle Big Data Appliance 4.8)
- Oracle Big Data Spatial & Graph 2.2.0
- MySQL Enterprise Edition 5.7.17
- Perfect Balance 2.10.0
- Java JDK 8u121
- Oracle Linux 6.8 UEK 4 (Oracle Linux Unbreakable Kernel, Release 4)
Mammoth v4.8.0 for Oracle Linux 5 is available (only for upgrades of clusters based on Oracle Linux 5).
- Oracle Data Integrator Agent 12.2.1.1 (for Oracle Big Data Connectors)
- Oracle R Advanced Analytics for Hadoop (ORAAH) 2.7.0
- Oracle's R Distribution (ORD) 3.2.0
- Oracle NoSQL Community Edition or Oracle NoSQL Database Enterprise Edition 4.3.10. Both optional.

Note that support is no longer available for Oracle NoSQL Database Community Edition.

- Kudu 1.2
- Kafka 2.1
- Spark 2.0
- Cloudera Key Trustee Server 5.10.1

The Cloudera parcels for Spark, Kudu, Kafka, and Key Trustee Server are included for your convenience, but are not deployed or configured by default.

See the [Cloudera Documentation](#) for information about CDH and Cloudera Manager 5.10.1

Oracle Big Data Appliance Release 4.8 includes Oracle Big Data SQL 3.1 as a Mammoth installation option. You do not need to download this package from the [Oracle Software Delivery Cloud](#).

New Features

In addition to software upgrades and extra Cloudera parcels, Oracle Big Data Appliance provides these new features.

- **Automated LAG Setup**

The Oracle Big Data Appliance Configuration Utility provides an option to assign groups of connectors to LAGs (Link Aggregation Groups) for cluster client network connections. LAGs can provide faster uplink times and reduce the points of potential failure in the network. See [Configuring for LACP/LAG Connections](#) (page 3-8) in this guide for details.

- **MIT Kerberos Security Option for Oracle NoSQL Enterprise installations**

You can select MIT Kerberos security for an Oracle NoSQL cluster via the Oracle Big Data Appliance Configuration Utility.

- **Oracle Big Data SQL can Connect to the Exadata Database Machine Over Ethernet**

Ethernet is now supported for Oracle Big Data SQL networking between Oracle Exadata Database Machine and Oracle Big Data Appliance. This enables you to use Oracle Big Data SQL with these two Engineered Systems in environments where InfiniBand is not feasible, such as when the two systems are geographically distant from each other.

The Oracle Big Data Appliance Configuration Utility now provides an option to choose between Ethernet and InfiniBand for Oracle Big Data SQL when you are setting up the configuration for a cluster.

Ethernet connections for Oracle Big Data SQL are not supported for networking between Oracle Big Data Appliance and Oracle SPARC SuperCluster.

- **Oracle Big Data SQL Networking is Switchable Between Ethernet and InfiniBand**

If needed, you can switch between InfiniBand and Ethernet after the Mammoth installation and after the Oracle Big Data SQL installation has been completed. See [Choosing Between Ethernet and InfiniBand Connections For Oracle Big Data SQL](#) in this guide for details.

Oracle Big Data SQL 3.1 can be Enabled/Disabled Using bdaccli

In Release 4.8, the Oracle Big Data Appliance Configuration Generation Utility provides an option to install Oracle Big Data SQL 3.1. If you choose this option, then the Hadoop side of Oracle Big Data SQL is enabled automatically.

If you do not choose to enable Oracle Big Data SQL in the Mammoth installation, you can enable it later using the bdaccli command line interface:

```
# bdaccli enable big_data_sql
```

You can also use bdaccli to disable the Hadoop side Oracle Big Data SQL:

```
# bdaccli disable big_data_sql
```



Note:

For Oracle Big Data Appliance Release 4.8, the instructions in the *Oracle Big Data SQL 3.1 Installation Guide* for downloading, installing, and uninstalling the software on the Hadoop cluster do not apply.

- You do not need to download Oracle Big Data SQL 3.1 from the Oracle Software Delivery Cloud. The software is already included in the Mammoth bundle.
- Do not install Oracle Big Data SQL 3.1 on Release 4.8 using `./setup-bds install bds-config.json`.

Mammoth performs this part of the installation for you. Likewise, do not use `./setup-bds uninstall bds-config.json` to uninstall the software from the Hadoop cluster. To disable Oracle Big Data SQL after the Mammoth installation (or to enable it if you did not opt include it in the Mammoth installation), use the bdaccli commands as described above.

Except for the install and uninstall commands, the instructions on using setup-bds in the [Oracle Big Data SQL 3.1 Installation Guide](#) apply to Oracle Big Data Appliance and should be used as documented. This is also true for the database side of Oracle Big Data SQL. Create the database installation bundle and install it on the nodes of the Oracle Database system as described in the guide.



See Also:

[Oracle Big Data SQL Online Documentation Library Release 3.1](#)

In Release 4.8, Node Migration and the mammoth -e Command Require Extra Steps for Oracle Big Data SQL Support

- If Oracle Big Data SQL is enabled, then you must disable it on the entire cluster prior to running `bdaccli admin_cluster migrate <node>`. Re-enable it after the migration. Use the bdaccli enable and disable commands.
 1. Disable Oracle Big Data SQL across the cluster: `bdaccli disable big_data_sql`

2. Perform the migration.
 3. Re-enable Oracle Big Data SQL: `bdacli enable big_data_sql`
- The Mammoth command for extending an installation to additional servers, `mammoth -e`, does not make necessary Oracle Big Data SQL changes. Correct the problem by calling `./setup-bds extend bds-config.json` from the node where Cloudera Manager is installed after `mammoth -e` completes.

TLS 1.0 (TLSv1) Disabled for Cloudera Configuration Manager and Hue – May Affect Client Access

In order to improve security, Oracle Big Data Appliance has disabled TLS 1.0 for Cloudera Manager and Hue and also in the system-wide Java configuration. This can affect older browsers or operating systems.

It is recommended that you reconfigure or upgrade clients using TLS 1.0 to use newer encryption, but if necessary you can re-enable TLS 1.0 for Cloudera Manager and Hue.

For details, log on to [My Oracle Support](#) and search for *BDA 4.8 Disables TLSv1 by Default For Cloudera Manager/Hue/And in System-Wide Java Configurations (Doc ID 2250841.1)*.

Oracle Big Data Discovery Installation on Big Data Appliance 4.8 not Supported at This Time

A compatible build of Oracle Big Data Discovery is under development. This document will be updated when a compatible build is available.

Deprecated or Discontinued Features

- The Perfect Balance automatic invocation feature is not supported in this release. If you have been using automatic invocation, please switch to the Perfect Balance API.
- There is no longer an option to purchase support for Oracle NoSQL Database Community Edition.

Other Notices

- Mammoth, the Oracle Big Data Appliance installation package, now consists of four ZIP files, all of which must be downloaded and extracted before running the `BDAmammoth-4.8.0.run` file
- Mammoth cluster upgrades to v4.8.0 are only supported directly from Mammoth v4.4.0 or later. For upgrades from earlier Mammoth versions, it will be necessary to first upgrade to Mammoth v4.4.0.
- Configuring the network on new racks with an installed BDA Base Image lower than BDA v4.5.0 requires special configuration steps. Log on to [My Oracle Support](#) and search for this document :
Network Configuration Instructions for Shipped BDA racks with a BDA Base Image Less Than V4.5.0 (Doc ID 2135358.1).
- Oracle strongly recommends that customers using Oracle Big Data Appliance Clusters with Microsoft Active Directory configure their clusters to use cross-realm trust to Microsoft Active Directory and not use Direct-Active-Directory configuration for the Hadoop cluster.

- Before upgrading an Oracle Linux 5 cluster that has had HTTPS enabled for Cloudera Manager, it is necessary to follow these steps :

1. Confirm that HTTPS is enabled for Cloudera Manager .

2. Back up GridDefParams.pm:

```
# cp /opt/oracle/BDAMammoth/bdaconfig/GridDefParams.pm /opt/oracle/BDAMammoth/bdaconfig/GridDefParams.pm.orig
```

3. Update lines 141 and 142 of /opt/oracle/BDAMammoth/bdaconfig/GridDefParams.pm to these settings:

```
our $CM_ADMIN_PROTO =  
"https";
```

```
our $CM_ADMIN_PORT =  
"7183";
```

Changes in Previous Releases of Oracle Big Data Appliance

The following summarizes the change history of Oracle Big Data Appliance.

Changes in Oracle Big Data Appliance Release 4 (4.7)

Oracle Big Data Appliance 4.7 is focused on defect fixes and software stack version updates, including major updates to the UEK kernel, MySQL, and Cloudera (Cloudera Manager and CDH). There are no other customer-visible changes in this release.

Software Updates

- CDH (Cloudera's Distribution including Apache Hadoop) 5.9
- CM (Cloudera Manager) 5.9
- Oracle Big Data Connectors 4.7
- Big Data SQL 3.0.1 or 3.1 (both optional)
- MySQL Enterprise Edition 5.6.32
- Perfect Balance 2.9
- Java JDK 8u111
- Oracle Linux 6.8 UEK 4 (Oracle Linux Unbreakable Kernel, Release 4)
- Oracle Data Integrator Agent 12.2.1.1 (for Oracle Big Data Connectors)
- Oracle R Advanced Analytics for Hadoop (ORAAH) 2.7.0
- Oracle's R Distribution (ORD) 3.2.0
- Oracle Big Data Discovery 1.3.2 or 1.4 (1.4.0.37.1388 or greater). This is optional.
- Oracle NoSQL Database Community Edition or Enterprise Edition 4.2.9 (optional)

See the [Cloudera Documentation](#) for information about CDH and Cloudera Manager 5.9

Oracle Big Data Appliance Release 4.7 includes Oracle Big Data SQL 3.0.1 as a Mammoth installation option. You do not need to download this package from the

[Oracle Software Delivery Cloud](#). Oracle Big Data SQL 3.1 is also compatible with Release 4.7 and is available for download from Oracle Software Delivery Cloud.

Changes in Oracle Big Data Appliance Release 4 (4.6)

The following are changes in Oracle Big Data Appliance Release 4 (4.6):

Software Updates

- CDH (Cloudera's Distribution including Apache Hadoop) 5.8
- CM (Cloudera Manager) 5.8.1
- Oracle Big Data Connectors 4.6
- Oracle Data Integrator Agent 12.2.1.1 (for Oracle Big Data Connectors)
- Oracle R Advanced Analytics for Hadoop (ORAAH) 2.6.0
- Oracle's R Distribution (ORD) 3.2.0
- Perfect Balance 2.8
- Oracle Linux 6.8
- Big Data Discovery 1.2.2 or 1.3.x (optional)
- Big Data SQL 3.0.1 (optional)
- Java JDK 8u101

New Features

- **Cloudera CDH 5.8 and Cloudera Manager 5.8.1**

See the [Cloudera Enterprise 5.8.x Documentation](#) for information about CDH and Cloudera Manager 5.8.x

- **Oracle Big Data SQL Updates**

Oracle Big Data Appliance Release 4.6 includes Oracle Big Data SQL 3.0.1 as a Mammoth installation option. You do not need to download it from the [Oracle Software Delivery Cloud](#) to install it on Oracle Big Data Appliance 4.6.

Oracle Big Data SQL 3.1 will be downloadable from the [Oracle Software Delivery Cloud](#) when available.

- **Networking Changes for Greater Configuration Flexibility**

The release provides more modular control over Oracle Big Data Appliance networks by separately storing the network configuration settings for each rack and cluster: `<rack_name>-rack-network.json` and `<cluster_name>cluster-network.json`. When you are using the Oracle Big Data Appliance Configuration Generation Utility, these changes enable you to reconfigure the client network or private network of a cluster without affecting the configuration of other servers on the rack. They also allow you to expand a cluster or rack without affecting servers that are not part of the configuration.

In previous releases, all such information was stored in a single `network.json` file. This file still exists for backward compatibility with some scripts.

- **Encryption for Data Spills and some Intermediate Files**

Data spills to disk outside of HDFS during the following memory-intensive processes can now be protected by encryption.

- Spark shuffle.
- Creation of intermediate files in MapReduce encrypted shuffle and spillage during map and reduce operations.
- Impala SQL queries that generate extremely large result sets.

When you enable Hadoop Network Encryption in Mammoth during a full Oracle Big Data Appliance installation, or later, via **bdaccli enable hadoop_network_encryption**, then encryption is also enabled for Spark, Impala, as well MapReduce intermediate files and data spills.

 Note:

In the case of an Oracle Big Data Appliance upgrade to Release 4.6, the upgrade does not automatically enable this new extension of encryption to data spills, regardless of whether or not Hadoop Network Encryption is already enabled. If you want this feature on a system that has been upgraded to Release 4.6, run **bdaccli enable hadoop_network_encryption** after the upgrade.

- **New reset Command in bdaccli**

The `bdaccli reset` command selectively reconfigures Oracle Big Data Appliance networks. It pulls the new settings from the network configuration files generated by Oracle Big Data Appliance Configuration Generation Utility. The user controls the scope of the reset (server, network, cluster) and which networks are reset.

See [bdaccli reset](#) (page 13-10) in the [Oracle Big Data Appliance Utilities](#) (page 13-1) section of this guide.

- **Support for Oracle NoSQL Database 4.0.9**

Other Changes

- **Mammoth Installation Step Changes**

Some Mammoth installation steps have been reorganized and renamed. An important change to note is that the Kerberos installation now consists of two separate pre- and post-cluster configuration steps in order to enable additional security setups on the cluster. See [Mammoth Installation Steps](#) (page 10-29).

- X6-2 servers are shipped with an Oracle Big Data Appliance v4.5.0 base image.

Changes in Oracle Big Data Appliance Release 4 (4.5)

The following are changes in Oracle Big Data Appliance Release 4 (4.5):

Software Updates

- CDH (Cloudera's Distribution including Apache Hadoop) 5.7
- CDM (Cloudera Manager) 5.7
- Cloudera Navigator 2.4.1
- Java JDK 8u92
- Oracle Big Data Connectors 4.5

- Oracle Data Integrator Agent 12.2.1 (for Oracle Big Data Connectors)
- Oracle NoSQL Database 4.0.5
- Perfect Balance 2.7
- Big Data Discovery 1.2.2 (optional)
- Big Data SQL 3.0.1 (optional)

Hardware Updates

- Oracle Big Data Appliance X6-2 server

The X6-2 provides substantial increases in processing power and memory over X5-2 servers:

- 2 x 22-core (2.2GHz) Intel® Xeon® E5-2699 v4 processors.
- 8 x 32 GB DDR4-2400 memory (expandable to maximum of 768 GB per node).

X6-2 servers are shipped with an Oracle Big Data Appliance v4.4.0 base image.

X6-2 nodes can be mixed with X5-2 nodes (and older Release 4.5-compatible nodes) in a CDH or NoSQL cluster. The X6-2 server is not compatible as a node of an Oracle Big Data Appliance cluster in releases prior to Oracle Big Data Appliance Release 4.4.

See the [Oracle Big Data Appliance X6-2 Data Sheet](#) for more details.

New Features

- **Cloudera CDH 5.7 and Cloudera Manager 5.7**

See the [Cloudera Enterprise 5.7.x Documentation](#) for information about CDH 5.7 and Cloudera Manager 5.7

- **Support for Either Local or Remote Key Trustee Servers**

Oracle Big Data Appliance supports both local and remote Key Trustee Servers for HDFS Transparent Encryption. The Oracle Big Data Configuration Utility includes HDFS Transparent Encryption as a configuration option. You can either click a checkbox to automatically install and configure active and passive Key Trustee Servers locally on the Oracle Big Data Appliance or define an “off-board” configuration, including the address of the active and passive servers, the Key Trustee organization, and the authorization code. You can also enable HDFS Transparent Encryption via the `bdacli` utility at any time after Mammoth installation and will be prompted to make the same choice between remote or local key trustee services.

Although Oracle Big Data Appliance supports local Key Trustee Servers, remote servers are still the recommended choice.

- **Support for Oracle Big Data SQL 3.0.1**

Oracle Big Data Appliance Release 4.5 includes Oracle Big Data SQL 3.0.1 as a Mammoth installation option. See the [Oracle Big Data SQL User's Guide](#) for Release 3.0.1 installation instructions.

- **Enhanced Networking**

Release 4.5 provides more flexibility in the configuration of networks on the Oracle Big Data Appliance. This includes support for the following options:

- Separate networks for each cluster in a rack (both client and private networks).
- Multiple client networks on the same BDA cluster.
- VLAN tagging for client networks.
- Partition keys for private InfiniBand networks.

- **Lower Minimum Size for CDH Clusters**

The minimum recommended CDH cluster size for a production environment is now five nodes. For development purposes, the Oracle Big Data Appliance Configuration Generation Utility now enables you to create three-node CDH clusters.

Note that Oracle Big Data Appliance Starter Rack is still sold with six servers.

Changes in Oracle Big Data Appliance Release 4 (4.4)

The following are changes in Oracle Big Data Appliance Release 4 (4.4):

Software Updates

- CDH (Cloudera's Distribution including Apache Hadoop) 5.5.1
- CDM (Cloudera Manager) 5.5.1
- Cloudera Navigator 2.4.1
- MySQL Database Enterprise Server - Advanced Edition 5.6
- Oracle Big Data Connectors 4.4
- Oracle Data Integrator Agent 12.2.1 (for Oracle Big Data Connectors)
- Oracle NoSQL Database 3.5.2
- Perfect Balance 2.6

Hardware Updates

- Oracle Big Data Appliance X6-2 server

The X6-2 provides substantial increases in processing power and memory over X5-2 servers:

- 2 x 22-core (2.2GHz) Intel® Xeon® E5-2699 v4 processors.
- 8 x 32 GB DDR4-2400 memory (expandable to maximum of 768 GB per node).

X6-2 servers are shipped with an Oracle Big Data Appliance v4.4.0 base image.

In Oracle Big Data Appliance Release 4.4 or greater, X6-2 nodes can be mixed with X5-2 nodes (and older Release 4.4-compatible nodes) in a CDH or NoSQL cluster. The X6-2 server is not compatible as a node of an Oracle Big Data Appliance cluster in releases previous to 4.4.

See the [Oracle Big Data Appliance X6-2 Data Sheet](#) for more details.

New Features

- **Cloudera CDH 5.5.1 and Cloudera Manager 5.5.1**

CDH 5.5.1 is a maintenance release on top of CDH 5.5. See the [Cloudera CDH 5.5 Release Notes](#)

For information on Cloudera Manager 5.5 and 5.5.1, see [New Features and Changes in Cloudera Manager 5](#)

- **Automated Installation for Cloudera Navigator**

Mammoth now provides an automated installation for Cloudera Navigator in both a full Mammoth installation and Mammoth upgrade. No user intervention is required and the installation occurs transparently. If Cloudera Navigator is not already installed, Mammoth installs the software on node 3 of the cluster, which is where other Cloudera Management services are hosted. If Cloudera Navigator is already installed, Mammoth skips this step and does not overwrite the existing installation.

The Cloudera Navigator Metadata Server and Audit Server are automatically added to Cloudera Manager and auditing is enabled. Mammoth also enables Web UI encryption for the Audit Server.

Mammoth does not enable the Cloudera Navigator key management components.

- **Support for Oracle Big Data SQL 3.0**

Oracle Big Data Appliance Release 4.4 includes Oracle Big Data SQL 2.0 as a Mammoth installation option. Oracle Big Data SQL 3.0 is also available for Release 4.4, as a patch. See the [Oracle Big Data SQL User's Guide](#) for Release 3.0 installation instructions.



Note:

If you want to install Oracle Big Data SQL 3.0, do not select Oracle Big Data SQL 2.0 in the Mammoth installation. If Oracle Big Data SQL 2.0 is installed, you must uninstall it prior to installing the 3.0 patch. The patch README file includes steps for removing 2.0 if you have previously installed it.

Release 4.4 as an Update to a Earlier Base Image

Mammoth 4.4.0 can run on top of any earlier Oracle Big Data Appliance 4.x base image and will update the base image software as needed.

Changes in Oracle Big Data Appliance Release 4 (4.3)

The following are changes in Oracle Big Data Appliance release 4 (4.3):

Software Updates

- CDH (Cloudera's Distribution including Apache Hadoop) 5.4.7
- CDM (Cloudera Manager) 5.4.7
- Oracle Big Data Connectors 4.3
- Oracle Big Data Discovery 1.1.1
- Oracle Big Data SQL 2.0
- Oracle NoSQL Database 1.3.4.7 (Community and Enterprise Edition)
- Oracle Table Access for Hadoop and Spark
- Perfect Balance 2.5

- JDK 8u60

See *Oracle Big Data Appliance Software User's Guide*.

New Features

- **Automatic Installation for Oracle Big Data Discovery**

Customers can download Big Data Discovery 1.1.1 and then use the `bdacli` command line utility to install the software on a designated node of the primary CDH cluster.

See [Expanding an Oracle Big Data Appliance Starter Rack](#) (page 8-1).

- **Oracle Table Access for Hadoop and Spark**

Oracle Table Access for Hadoop and Spark is an Oracle Big Data Appliance feature that converts Oracle Database tables into Hadoop or Spark data sources. This feature enables fast and secure access to data in the Oracle Database.

- **HDFS Transparent Encryption**

Oracle Big Data Appliance 4.3 provides the option to use HDFS Transparent Encryption. This replaces the eCryptfs on-disk encryption software provided with previous releases. Customers can enable HDFS Transparent Encryption for both new and pre-existing CDH clusters. When enabled, HDFS Transparent Encryption secures Hadoop operations running on the cluster (including HDFS, MapReduce on YARN, Spark on YARN, Hive, and Hbase tasks).

- The Oracle Big Data Appliance Configuration Generation Utility provides an option to include HDFS Transparent Encryption when a new cluster is created.
- HDFS Transparent Encryption can be enabled or disabled on a cluster via the `bdacli` command line interface.

- **HTTPS / Network Encryption**

- Provides HTTPS for Cloudera Manager, Hue, Oozie, and Hadoop Web UIs.
- Enables network encryption for other internal Hadoop data transfers, such as those made through YARN shuffle and RPC.

Like HDFS Transparent Encryption, HTTPS/ Network Encryption is an option in the Oracle Big Data Appliance Configuration Generation Utility, and can also be enabled via `bdacli`.

- **Zero Downtime for Upgrades, One-Off Patches, and Cluster Extensions**

In Release 4.3, Oracle Big Data Appliance leverages Cloudera's Rolling Upgrades functionality to keep clusters operational during Mammoth upgrades, patches, and cluster extensions. This is an installation option that allows certain services on a cluster to remain continuously available while each node in the cluster is upgraded and rebooted. Zero Downtime is an option for the following tasks:

- Upgrades of the Mammoth software (including Cloudera's Distribution Including Apache Hadoop, Cloudera Manager, and the Mammoth software itself).
- One-off patches of Mammoth-installed software.
- Cluster extensions. (For cluster extensions within a single rack, rolling upgrades are not optional. These extensions are always done as rolling upgrades.)

Deprecated Features

The following features are deprecated in this release, and may be desupported in a future release:

- **MapReduce 1 (MRv1)**

YARN (MRv2) supersedes MRv1. Users who want to continue to use MRv1 on Oracle Big Data Appliance versions 3.x and 4.x should contact Oracle Support before using Mammoth to patch or upgrade the software.

Desupported Features

The following features are no longer supported as of this release:

- **eCryptfs On-Disk Encryption**

This has been replaced by HDFS Transparent Encryption.

Changes in Oracle Big Data Appliance Release 4 (4.2)

The following are changes in Oracle Big Data Appliance release 4 (4.2):

New Features

- **Software Upgrades**

- Cloudera's Distribution including Apache Hadoop 5.4.0
- Cloudera Manager 5.4.0
- Perfect Balance 2.4.0
- Oracle Big Data SQL 1.1
- No SQL Database 3.2.5
- Oracle Linux 6.6 and 5.11
- JDK 8u45

See *Oracle Big Data Appliance Software User's Guide*.

- **Hardware Upgrades**

- Oracle Big Data Appliance is now shipped with 8 TB disk drives

- **Elastic Configuration**

- Oracle Big Data Appliance now provides the flexibility of adding one or more servers on a starter rack using Big Data Appliance X5-2 High Capacity Nodes plus InfiniBand Infrastructure. You can add up to 12 additional servers on a starter rack.

See [Expanding an Oracle Big Data Appliance Starter Rack](#) (page 8-1).

- **Automatic Installation Support**

- Spark-on-YARN is deployed automatically
- Oracle Spatial and Graph is installed and configured automatically

- **Oracle Big Data SQL 1.1**

- Copy to BDA

This utility enables you to copy relatively static tables from an Oracle database into Hadoop, with the purpose of improving query times.

See *Oracle Big Data Appliance Software User's Guide*.

- Oracle NoSQL Database Support

Oracle databases on Oracle Exadata Database Machine can use Oracle Big Data SQL to connect to clusters running Oracle NoSQL Database.

- Parquet Support

CDH 5.2 and later versions include Hive 0.13, which supports the Apache Parquet file format. This file format is used by Cloudera Impala and other Hadoop software.

See *Oracle Big Data Appliance Software User's Guide*.

Other Changes

- **Oracle Big Data Appliance X5-2**

Oracle Big Data Appliance 4.2 software supports Oracle Big Data Appliance X5-2 and earlier version server hardware.

See "[Server Components](#) (page 1-2)".

- **Oracle Big Data Appliance Configuration Generation Utility**

This utility generates two new configuration files:

- `network.json`: Supersedes `BdaDeploy.json`. For software upgrades, Mammoth converts the existing `BdaDeploy.json` to `network.json`. New installations must have `network.json`.
- `networkexpansion.json`: Supersedes `BdaExpansion.json`.

See "[About the Configuration Files](#) (page 4-3)".

- **CDH Deployment**

Mammoth uses parcels instead of RPMs to deploy CDH.

- **Apache Sentry**

Installation of Apache Sentry does not require `sentry-provider.ini` as a prerequisite.

- **Microsoft Active Directory Server in Mammoth**

Support for directly using Microsoft Active Directory named as Active Directory Kerberos in Mammoth.

- **Oracle Linux Support**

Oracle Linux 5 support for Oracle Big Data Appliance X5-2 servers.

- **Cloudera Navigator Trustee Server**

Cloudera Navigator Trustee Server installer package and documentation are now shipped in Mammoth. It must be manually installed on a separate server.

Deprecated Features

The following features are deprecated in this release, and may be desupported in a future release:

- **Mammoth Reconfiguration Utility**

The `bdaccli` utility supersedes `mammoth-reconfig`. The `mammoth-reconfig` utility is only needed to change the disk encryption password.

See "[bdaccli](#) (page 13-7)".

- **MapReduce 1 (MRv1)**

YARN (MRv2) supersedes MRv1. Users who want to continue to use MRv1 on Oracle Big Data Appliance versions 3.x and 4.x should contact Oracle Support before using Mammoth to patch or upgrade the software.

- **Disk Encryption**

A new encryption system that is more flexible and robust will replace the current system in an upcoming release.

Changes in Oracle Big Data Appliance Release 4 (4.1)

The following are changes in Oracle Big Data Appliance release 4 (4.1):

New Features

- **Software Upgrades**

- Cloudera's Distribution including Apache Hadoop 5.3.0
- Cloudera Manager 5.3.0
- Perfect Balance 2.3.0
- Oracle Big Data SQL 1.1
- Oracle Big Data Connectors 4.1
- Oracle Linux 6.5

See *Oracle Big Data Appliance Software User's Guide*.

- **Oracle Big Data SQL 1.1**

- Copy to BDA

This utility enables you to copy relatively static tables from an Oracle database into Hadoop, with the purpose of improving query times.

See *Oracle Big Data Appliance Software User's Guide*.

- Oracle NoSQL Database Support

Oracle databases on Oracle Exadata Database Machine can use Oracle Big Data SQL to connect to clusters running Oracle NoSQL Database.

- Parquet Support

CDH 5.2 and later versions include Hive 0.13, which supports the Apache Parquet file format. This file format is used by Cloudera Impala and other Hadoop software.

See *Oracle Big Data Appliance Software User's Guide*.

- **Oracle NoSQL Database**

The `bdaccli admin_cluster` command supports Oracle NoSQL Database nodes that require repair or replacement.

See *Oracle Big Data Appliance Software User's Guide*.

Other Changes

- **Oracle Big Data Appliance X5-2**

Oracle Big Data Appliance 4.1 software supports the Oracle Big Data Appliance X5-2 server hardware.

See "[Server Components](#) (page 1-2)".

- **Oracle Big Data Appliance Configuration Generation Utility**

This utility generates two new configuration files:

- `network.json`: Supersedes `BdaDeploy.json`. For software upgrades, Mammoth converts the existing `BdaDeploy.json` to `network.json`. New installations must have `network.json`.
- `networkexpansion.json`: Supersedes `BdaExpansion.json`.

See "[About the Configuration Files](#) (page 4-3)".

- **CDH Deployment**

Mammoth uses parcels instead of RPMs to deploy CDH.

- **Apache Sentry**

Installation of Apache Sentry does not require `sentry-provider.ini` as a prerequisite.

Deprecated Features

The following features are deprecated in this release, and may be desupported in a future release:

- **Mammoth Reconfiguration Utility**

The `bdaccli` utility supersedes `mammoth-reconfig`. The `mammoth-reconfig` utility is only needed to change the disk encryption password.

See "[bdaccli](#) (page 13-7)".

- **MapReduce 1 (MRv1)**

YARN (MRv2) supersedes MRv1. Users who want to continue to use MRv1 on Oracle Big Data Appliance versions 3.x and 4.x should contact Oracle Support before using Mammoth to patch or upgrade the software.

- **Disk Encryption**

A new encryption system that is more flexible and robust will replace the current system in an upcoming release.

Changes in Oracle Big Data Appliance Release 4 (4.0)

The following are changes in Oracle Big Data Appliance release 4 (4.0):

New Features

- **Oracle Big Data SQL 1.0.0**

Oracle Big Data SQL supports queries against vast amounts of big data stored in multiple data sources, including HDFS and Hive. You can view and analyze data

from various data stores together, as if it were all stored in an Oracle database. Support for Oracle Big Data SQL includes the following new features in Oracle Database:

- DBMS_HADOOP PL/SQL package
- Hive static data dictionary views
- Access drivers for Hadoop and Hive

Oracle Big Data SQL is an installation option, which you can specify using the Oracle Big Data Appliance Configuration Generation Utility.

You can monitor and manage Oracle Big Data SQL using the `bdaccli` command and Cloudera Manager.

See "[bdaccli](#) (page 13-7)" and *Oracle Big Data Appliance Software User's Guide*.

- **Service Migration**

The `bdaccli` utility can migrate services from a failing critical node to a healthy noncritical node. It can also remove failing critical and noncritical nodes from a cluster, and restore them to the cluster after repairs. See "[bdaccli](#) (page 13-7)" and *Oracle Big Data Appliance Software User's Guide*.

- **Software Upgrades**

- Cloudera's Distribution including Apache Hadoop 5.1.0
- Cloudera Manager 5.1.1
- Perfect Balance 2.2.0
- Oracle Data Integrator Agent 12.1.3.0 (for Oracle Big Data Connectors)

See *Oracle Big Data Appliance Software User's Guide*.

- **Oracle NoSQL Database Zone Support**

The Oracle Big Data Appliance Configuration Generation Utility and the `mammoth -e` command support multiple zones on Oracle NoSQL Database clusters. You can add nodes to an existing zone, or create a new primary or secondary zones.

See "[Oracle NoSQL Configuration](#) (page 4-16)" and "[Mammoth Software Installation and Configuration Utility](#) (page 10-27)".

- **Multiple Rack Clusters**

You can now install a cluster on multiple racks using one `cluster_name-config.json` file.

Part I

Preinstallation

You must complete the preinstallation procedures before receiving shipment of Oracle Big Data Appliance, to avoid delays with the installation.

Part I contains the following chapters:

- [Introduction to Oracle Big Data Appliance](#) (page 1-1)
- [Site Requirements for Oracle Big Data Appliance](#) (page 2-1)
- [Understanding the Network Requirements](#) (page 3-1)
- [Using Oracle Big Data Appliance Configuration Generation Utility](#) (page 4-1)
- [Setting Up Auto Service Request](#) (page 5-1)

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](#) (page 1-1)
- [Oracle Big Data Appliance Rack Components](#) (page 1-1)
- [Server Components](#) (page 1-2)
- [Spare Parts Kit Components](#) (page 1-4)
- [Oracle Big Data Appliance Management Software](#) (page 1-6)

1.1 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 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

1.2 Oracle Big Data Appliance Rack Components

Oracle Big Data Appliance consists of a rack and the components to connect to your network. [Table 1-1](#) (page 1-2) lists the components in an Oracle Big Data Appliance rack.

Table 1-1 Oracle Big Data Appliance Rack Components

Full Rack	Starter Rack	Description
18	6	Oracle Server X6-2L, X5-2L, Sun Server X4-2L, Sun Server X3-2L, or Sun Fire X4270 M2 (full rack only)
1	1	Sun Rack II 1242 base
2	2	Sun Network QDR InfiniBand Gateway NM2-GW switches
1	1	Sun Datacenter InfiniBand Switch 36 NM2-36P switches
1	1	Cisco Catalyst 4948-E-F Ethernet switch (4948 in a Sun Fire X4270 M2-based rack)
2	2	Power distribution units
0 or 1	0	KVM (only in Sun Fire X4270 M2-based racks)

1.3 Server Components

[Table 1-2](#) (page 1-2) lists the factory-shipped components of an Oracle Server X6-2L.

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

Quantity	Description
1	Oracle Server X6-2L server base
2	22-core Intel® Xeon® E5-2699 v4 processors (2.2GHz)
8	32 GB DDR4-2400 memory (expandable to maximum of 768 GB)
12	8 TB 3.5-inch 7200 RPM drives high capacity SAS (hot swappable)
1	Oracle Integrated Lights Out Manager v3.1 with an Ethernet port for remote management
2	InfiniBand 4 x QDR (40 Gbps) ports (One dual-port PCIe 3.0 Host Channel Adapter network interface card)
4	Embedded 10 gigabit Ethernet port
1	Host bus adapter (disk controller) with 8 internal ports and 512 MB battery-backed write cache. 12 Gbps SAS III PCIe 8P RAID-INT
2	Redundant hot-swappable power supplies and fans
1	USB-2 4 GB flash drive

[Table 1-3](#) (page 1-2) lists the factory-shipped components of an Oracle Server X5-2.

Table 1-3 Oracle Server X5-2 Server Components

Quantity	Description
1	Oracle Server X5-2L server base
2	18-core Intel Xeon E5-2699 v3 processors (2.3 GHz)
8	16 GB DDR4 Memory

Table 1-3 (Cont.) Oracle Server X5-2 Server Components

Quantity	Description
12	8 TB 3.5-inch 7200 RPM drives high capacity SAS (hot swappable)
1	Oracle Integrated Lights Out Manager v3.1 with an Ethernet port for remote management
2	InfiniBand 4 x QDR (40 Gbps) ports (One dual-port PCIe 3.0 Host Channel Adapter network interface card)
4	Embedded 10 gigabit Ethernet port
1	Host bus adapter (disk controller) with 8 internal ports and 512 MB battery-backed write cache. 12 Gbps SAS III PCIe 8P RAID-INT
2	Redundant hot-swappable power supplies and fans
1	USB-2 4 GB flash drive

[Table 1-4](#) (page 1-3) lists the components of a Sun Server X4-2L.

Table 1-4 Sun Server X4-2L Server Components

Quantity	Description
1	Sun Server X4-2L server base
2	Eight-core Intel Xeon E5-2650 v2 processors (2.6 GHz)
8	8 GB DDR3 2RX4 1600 MHz DIMMs (64 GB RAM expandable up to 512 GB)
12	8 TB 3.5-inch 7200 RPM drives high capacity SAS (hot swappable)
1	Oracle Integrated Lights Out Manager v3.1 with an Ethernet port for remote management
2	InfiniBand 4 x QDR (40 Gbps) ports (One dual-port PCIe 3.0 Host Channel Adapter network interface card)
4	Embedded 10 gigabit Ethernet port
1	Host bus adapter (disk controller) with 8 internal ports and 512 MB battery-backed write cache. SAS III 6 Gbps
2	Redundant hot-swappable power supplies and fans
1	USB-2 4 GB flash drive

[Table 1-5](#) (page 1-3) lists the components of a Sun Server X3-2L.

Table 1-5 Sun Server X3-2L Server Components

Quantity	Description
1	Sun Server X3-2L server base
2	Eight-core Intel Xeon E5-2660 processors (2.2 GHz)
8	8 GB DDR3 2RX4 1600 MHz DIMMs (64 GB RAM expandable up to 512 GB)
12	3 TB 3.5-inch 7200 RPM drives high capacity SAS (hot swap)
1	Oracle Integrated Lights Out Manager v3.1 with an Ethernet port for remote management

Table 1-5 (Cont.) Sun Server X3-2L Server Components

Quantity	Description
2	InfiniBand 4 x QDR (4 Gbps) ports (One dual-port PCIe 2.0 Host Channel Adapter network interface card)
4	Embedded 10 gigabit Ethernet port
1	Host bus adapter (disk controller) with 8 internal ports and 512 MB battery-backed write cache. SAS III 6 GBPS
2	Redundant hot-swappable power supplies and fans
1	USB-2 4 GB flash drive

[Table 1-6](#) (page 1-4) lists the components of a Sun Fire X4270 M2.

Table 1-6 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
4	Embedded 1 gigabit Ethernet port
1	USB-2 4 GB flash drive

1.4 Spare Parts Kit Components

The Oracle Big Data Appliance spare parts kit includes disk drives, accessory kits, cables, and documentation. [Table 1-7](#) (page 1-4) and [Table 1-8](#) (page 1-5) identify the bundled spares.

Table 1-7 Bundled Components and Accessories

Part Number	Quantity	Description
7066826	2	8 TB 3.5-inch 7200 RPM drives
590-896-501B	1	LCD Accessory Kit (Sun Fire X4270 M2-based racks only)
-	1	Oracle Storage Server Documentation Kit
-	1	Sun Rack II Accessory Kit
53-2332-xx	1	Cisco Switch Documents and Accessory Kit

Table 1-8 Bundled Cables

Part Number	Length	Quantity	Description
530-4446-01	5 meters	10	QSFP passive copper cable for multitracking
530-4445-01	3 meters	8	QSFP passive copper cable for multitracking
350-1519-01	-	3	ASSY, NM2 serial cable sets for the InfiniBand switches

[Table 1-9](#) (page 1-5) lists the spare cables that are secured to the center section of an Oracle Big Data Appliance rack.

Table 1-9 In-Rack Spare Cables

Part Number	Length	Quantity	Description
530-4445-01	3 meters	2	QSFP passive copper replacement cables
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

Table 1-10 InfiniBand Inter-Switch Cables and Spares Provided

Cable	Inter-switch	Spares	Total
InfiniBand 2 meter cable	7	0	7
InfiniBand 3 meter cable	2	6	8
InfiniBand 5 meter cable	0	10	10

Table 1-11 Additional Cables to Purchase (Based on Number of Racks)

Number of Racks	3 Meter Cables Needed	5 Meter Cables Needed	10 Meter Cables Needed
2	16	16	0
3	18	30	0
4	16	40	8
5	20	42	18
6	24	46	26
7	28	44	40
8	16	52	60

1.5 Oracle Big Data Appliance Management Software

The software components are installed on all 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 [Oracle Big Data Appliance Utilities](#) (page 13-1) 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 [Using Oracle Integrated Lights Out Manager](#) (page 15-1).
- 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](#) (page 14-1)".
- The `dcli` utility executes commands across a group of servers on Oracle Big Data Appliance and returns the output. See [Executing Commands Across a Cluster Using the dcli Utility](#) (page 14-1).
- Various utilities check the health of the cluster and the networks. See [Oracle Big Data Appliance Utilities](#) (page 13-1).

All end-user software is installed on site using the Mammoth utility. See [Installing the Oracle Big Data Appliance Software](#) (page 10-1).



See Also:

Oracle Big Data Appliance Software User's Guide for information about using Oracle Enterprise Manager Cloud Control and Cloudera Manager to monitor Oracle Big Data Appliance.

1.6 Supportability Guidelines

Oracle sets guidelines for supportable customer modifications to the Oracle Big Data Appliance hardware and software. Oracle support is contingent upon customer compliance with these guidelines.

See Restrictions on Use as Condition of Receiving Support in the *Oracle Big Data Appliance License Information User Manual*

2

Site Requirements for Oracle Big Data Appliance

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](#) (page 2-1)
- [Space Requirements](#) (page 2-4)
- [Flooring Requirements](#) (page 2-6)
- [Electrical Power Requirements](#) (page 2-6)
- [Temperature and Humidity Requirements](#) (page 2-11)
- [Ventilation and Cooling Requirements](#) (page 2-13)
- [Network Connection Requirements](#) (page 2-14)
- [Ensuring That the Site Is Ready](#) (page 2-14)

2.1 General Environmental Requirements

The environmental requirements for Oracle Big Data Appliance depend on the model and the size of the system. The following tables provide an overview of the rack requirements. Measurements are approximate. The other sections in this chapter provide detailed information.

[Table 2-1](#) (page 2-1) identifies the general environmental requirements for an Oracle Big Data Appliance X6-2 or X5-2 rack.

Table 2-1 Environmental Requirements for Oracle Big Data Appliance X6-2 and X5-2

Environmental Component	Starter Rack	Full Rack
Net Weight	415 kg (915 lb)	836 kg (1843 lb)
See Also: " Flooring Requirements (page 2-6)"		
Acoustic levels	8.4 Bel	8.5 Bel
Power	Maximum: 5.2 kW	Maximum: 13.0 kW
See Also: " Electrical Power Requirements (page 2-6)"	Typical: 3.6 kW	Typical: 9.1 kW
Cooling	Maximum: 17,521 BTU/hour (18,485 kJ/hour)	Maximum: 44,136 BTU/hour (46,564 kJ/hour)
See Also: " Temperature and Humidity Requirements (page 2-11)" and " Ventilation and Cooling Requirements (page 2-13)"	Typical: 12,265 BTU/hour (12,940 kJ/hour)	Typical: 30,895 BTU/hour 32,594 kJ/hour

Table 2-1 (Cont.) Environmental Requirements for Oracle Big Data Appliance X6-2 and X5-2

Environmental Component	Starter Rack	Full Rack
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements (page 2-11)" , and "Ventilation and Cooling Requirements (page 2-13)"	Maximum: 811 CFM Typical: 568 CFM	Maximum: 2043 CFM Typical: 1430 CFM
IP addresses See Also: Understanding the Network Requirements (page 3-1) .	12 for Ethernet network 6 for InfiniBand network	24 for Ethernet network 18 for InfiniBand network
Network drops See Also: Understanding the Network Requirements (page 3-1) .	One network connection	One network connection
External connectivity See Also: Understanding the Network Requirements (page 3-1) .	1 x 1 Gbps Ethernet port 6 x 10 Gbps Ethernet ports	1 x 1 Gbps Ethernet ports 18 x 10 Gbps Ethernet ports

[Table 2-2](#) (page 2-2) identifies the general environmental requirements for an Oracle Big Data Appliance X5-2 High Capacity Node plus Infiniband Infrastructure.

Table 2-2 Environmental Requirements for Oracle Big Data Appliance X6-2 and X5-2 High Capacity Node plus Infiniband Infrastructure

Environmental Component	High Capacity Node plus Infiniband Infrastructure
Height	3.5 inches (87.6 mm)
Width	17.5 inches (445.0 mm)
Depth	29.0 inches (737.0 mm)
Net Weight See Also: "Flooring Requirements (page 2-6)"	33.1 kg (73.0 lb)
Power See Also: "Electrical Power Requirements (page 2-6)"	Maximum: 0.7 kW Typical: 0.5 kW (Typical power usage varies by application workload)
Cooling See Also: "Temperature and Humidity Requirements (page 2-11)" and "Ventilation and Cooling Requirements (page 2-13)"	Maximum: 2,481 BTU/hour (2617 kJ/hour) Typical: 1,736 BTU/hour (1831 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements (page 2-11)" , and "Ventilation and Cooling Requirements (page 2-13)"	Maximum: 115 CFM Typical: 80 CFM

[Table 2-3](#) (page 2-3) shows the general environmental requirements for Oracle Big Data Appliance X4-2 and Oracle Big Data Appliance X3-2.

Table 2-3 Environmental Requirements for Oracle Big Data Appliance X4-2 and X3-2

Environmental Component	Starter Rack	Full Rack
Net Weight See Also: "Flooring Requirements (page 2-6)"	418 kg (922 lb)	829 kg (1828 lb)
Acoustic levels	76 Bel	8.3 Bel
Power See Also: "Electrical Power Requirements (page 2-6)"	Maximum: 4.2 kW Typical: 3 kW	Maximum: 10 kW (10.2 kVA) Typical: 7 kW (7.15 kVA)
Cooling See Also: "Temperature and Humidity Requirements (page 2-11)" and "Ventilation and Cooling Requirements (page 2-13)"	Maximum: 14,052 BTU/hour (14,825.6 kJ/hour) Typical: 9,836 BTU/hour (10,377.5 kJ/hour)	Maximum: 34,142 BTU/hour (34,838 kJ/hour) Typical: 23,940 BTU/hour (24,420 kJ/hour)
Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements (page 2-11)", and "Ventilation and Cooling Requirements (page 2-13)"	Maximum: 676 CFM Typical: 473 CFM	Maximum: 1,573 CFM Typical: 1,103 CFM
IP addresses See Also: Understanding the Network Requirements (page 3-1).	12 for Ethernet network 6 for InfiniBand network	24 for Ethernet network 18 for InfiniBand network
Network drops See Also: Understanding the Network Requirements (page 3-1).	One network connection	One network connection
External connectivity See Also: Understanding the Network Requirements (page 3-1).	1 x 1 Gbps Ethernet port 6 x 10 Gbps Ethernet ports	1 x 1 Gbps Ethernet ports 18 x 10 Gbps Ethernet ports

[Table 2-4](#) (page 2-3) identifies the general environmental requirements for the Oracle Big Data Appliance Sun Fire X4270 M2-based rack.

Table 2-4 Environmental Requirements for Oracle Big Data Appliance Sun Fire X4270 M2-Based Rack

Environmental Component	Full Rack
Net Weight See Also: "Flooring Requirements (page 2-6)"	816.5 kg (1800 lb)
Acoustic levels	8.3 Bel

Table 2-4 (Cont.) Environmental Requirements for Oracle Big Data Appliance Sun Fire X4270 M2-Based Rack

Environmental Component	Full Rack
Power	Maximum: 12.0 kW (12.25 kVA)
See Also: "Electrical Power Requirements (page 2-6)"	Typical: 8.4 kW (8.6 kVA) (Typical power usage varies by application workload)
Cooling	Maximum: 40,971 BTU/hour (41,807 kJ/hour)
See Also: "Temperature and Humidity Requirements (page 2-11)" and "Ventilation and Cooling Requirements (page 2-13)"	Typical: 28,680 BTU/hour (29,265 kJ/hour)
Air flow front-to-back (subject to actual data center environment)	Maximum: Approximately 1,886 CFM
See Also: "Temperature and Humidity Requirements (page 2-11)" , and "Ventilation and Cooling Requirements (page 2-13)"	Typical: Approximately 1,340 CFM
IP addresses	24 for Ethernet network
See Also: Understanding the Network Requirements (page 3-1) .	18 for InfiniBand network
Network drops	Minimum of two network connections
See Also: Understanding the Network Requirements (page 3-1) .	
External connectivity	2 x 1 Gbps Ethernet ports
See Also: Understanding the Network Requirements (page 3-1) .	18 x 10 Gbps Ethernet ports

2.2 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 required above the rack height in the front and rear aisle space for maintenance access. The space surrounding the cabinet must not restrict the movement of cool air between the air conditioner and the front of the systems within the cabinet, or the movement of hot air coming out of the rear of the cabinet.

2.2.1 Receiving and Unpacking Requirements

Before Oracle Big Data Appliance arrives, ensure that the receiving area is large enough for the package. [Table 2-5 \(page 2-5\)](#) lists the dimensions and weight of the shipped items.

Table 2-5 Shipping Measurements

Measurement	Full Rack	Starter Rack
Height	216 cm (85 inches)	216 cm (85 inches)
Width	122 cm (48 inches)	122 cm (48 inches)
Depth	157.5 cm (62 inches)	157.5 cm (62 inches)
Weight	979 kg (2158 lb)	546 kg (1203 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 2-6](#) (page 2-5) lists the access route requirements.

Table 2-6 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)

2.2.2 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 2-7](#) (page 2-5) lists the maintenance access requirements.

Open tiles are required for electrical access.

Table 2-7 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)

2.3 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 2-8](#) (page 2-6) lists the floor load requirements.

Table 2-8 Floor Load Requirements for Oracle Big Data Appliance

Description	Requirement
Maximum allowable weight of installed rack equipment	952.5 kg (2100 lb)
Maximum allowable weight of installed power distribution units	52 kg (115 lb)
Maximum dynamic load (maximum allowable weight of installed equipment including PDUs)	1050 kg (2315 lb)

2.4 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.

2.4.1 PDUs for Oracle Big Data Appliance

The type of PDU depends on the location where Oracle Big Data Appliance is installed. Refer to the appropriate sections for your location:

- [Single-Phase PDUs for North America, South America, Japan, and Taiwan](#) (page 2-7)
- [Three-Phase PDUs for North America, South America, Japan, and Taiwan](#) (page 2-7)
- [Single-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific](#) (page 2-8)
- [Three-Phase PDUs for Europe, the Middle East, Africa, and Asia Pacific](#) (page 2-9)

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

[Table 2-9](#) (page 2-7) lists the requirements for low-voltage single-phase PDUs for North America, South America, Japan, and Taiwan. There are two PDUs per rack.

Table 2-9 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](#) (page 2-7) shows the low-voltage, single-phase PDU power connector for North America, South America, Japan, and Taiwan.

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



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

[Table 2-10](#) (page 2-8) lists the requirements for low-voltage three-phase PDUs for North America, South America, Japan, and Taiwan. There are two PDUs per rack.

Table 2-10 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](#) (page 2-8) shows the low-voltage three-phase PDU power connector for North America, South America, Japan, and Taiwan.

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



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

[Table 2-11](#) (page 2-8) lists the requirements for a 15 kVA high-voltage single-phase PDU for Europe, the Middle East, Africa, and Asia Pacific, except Japan and Taiwan. There are two PDUs per rack.

Table 2-11 High-voltage 22 kVA Single Phase PDUs

Option	Requirement per PDU
Voltage	220 to 240 VAC
Frequency	50/60 Hz

Table 2-11 (Cont.) High-voltage 22 kVA Single Phase PDUs

Option	Requirement per PDU
Current	25A maximum per input
Power rating	15 kVA
Output current	72A (3 x 32A)
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 1-phase voltage
- 6 receptacles to connect the PDUs to 2 IEC 60309 32A 3-pin 250 VAC IP44 data center receptacles, and that require 22 kVA, 1-phase

[Figure 2-3](#) (page 2-9) shows the high-voltage single-phase PDU power connector for EMEA and APAC, except for Japan and Taiwan.

Figure 2-3 High-voltage Single Phase Power Connector

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

[Table 2-12](#) (page 2-9) 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 2-12 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

Table 2-12 (Cont.) High-Voltage 15 kVA Three-Phase PDUs

Option	Requirement per PDU
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 (page 2-10) 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



2.4.2 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.

2.4.3 Circuit Breaker Requirements

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

In addition to 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.

2.4.4 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.

2.5 Temperature and Humidity Requirements

Airflow through Oracle Big Data Appliance is from front to back. See [Table 2-4](#) (page 2-3) 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 2-13](#) (page 2-12) lists the temperature, humidity, and altitude requirements for operating and nonoperating systems.

Table 2-13 Temperature, Humidity, and Altitude Requirements

Condition	Operating Requirement	Nonoperating Requirement	Optimum
Temperature	5 to 35 degrees Celsius (40 to 95 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	3,000 meters (9,840 feet) maximum	12,000 meters (39,400 feet) maximum	Ambient temperature is reduced by 1 degree Celsius for each 300 meters above 900 meters altitude above sea level

Set conditions to the optimal temperature and humidity ranges to minimize the chance of downtime 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.

2.6 Ventilation and Cooling Requirements

Always provide adequate space in front and behind the rack to allow for proper ventilation. Do not obstruct the front or rear of the rack with equipment or objects that might prevent air from flowing through the rack. Rack-mountable servers and equipment typically draw cool air in through the front of the rack and let out warm air through the rear of the rack. There is no air flow requirement for the left and right sides 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 rack has adequate air flow.
- Ensure that the rack 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 rack and 91.4 cm (36 inches) at the rear of the rack for ventilation.

Use perforated tiles, rated for 400 cubic feet per minute (CFM) per tile air flow, in front of the rack for cold air intake. The tiles can be arranged in any order in front of the rack, if cold air from the tiles can flow into the rack. Inadequate cold air flow could result in a higher inlet temperature in the servers due to exhaust air recirculation. Oracle recommends four floor tiles for Oracle Big Data Appliance.

[Figure 2-5](#) (page 2-14) 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



2.7 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 (Sun Fire X4270 M2-based racks only)
- Client access network connection requirements
 - 2 (minimum) to 16 (maximum) 10 Gbps Ethernet connections split between the two Sun Network QDR InfiniBand Gateway switches in the rack. The exact number of connections depends on your bandwidth requirements.



See Also:

[Understanding the Network Requirements](#) (page 3-1)

2.8 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](#) (page 2-15)
- [Task 2, Generate the Configuration Files](#) (page 2-15)
- [Task 3, Configure the Network](#) (page 2-15)
- [Task 4, Prepare the Site Based on Requirements](#) (page 2-15)

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 Generate the Configuration Files

Run the Oracle Big Data Appliance Configuration Generation Utility. Enter the network configuration settings, software installation choices, and credentials in the fields provided by this utility. When you are done, generate the configuration files.

See Also:

- [Understanding the Network Requirements](#) (page 3-1)
- [Using Oracle Big Data Appliance Configuration Generation Utility](#) (page 4-1)

Task 3 Configure the Network

Configure the existing network. This includes registering the networks in the Domain Name System (DNS), assigning IP addresses, and configuring data center switches and firewalls. Run the `bda-preinstall-checkip.sh` script and fix any errors.

See Also:

["Validating the Network Settings](#) (page 4-5)"

Task 4 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](#) (page 6-1)"

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](#) (page 2-4)"

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](#) (page 2-6)"

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

 See Also:

- "[Electrical Power Requirements](#) (page 2-6)"
- *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. Ensure that the installation site provides adequate air conditioning.

 See Also:

"[Ventilation and Cooling Requirements](#) (page 2-13)"

6. Operate the air conditioning system for 48 hours to bring the room temperature to the appropriate level.

 See Also:

"[Temperature and Humidity Requirements](#) (page 2-11)"

7. Install the network cables for Oracle Big Data Appliance.

 See Also:

"[Network Connection Requirements](#) (page 2-14)"

8. Clean and vacuum the area thoroughly in preparation for installation.

3

Understanding the Network Requirements

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

- [Overview of Network Requirements](#) (page 3-1)
- [Cabling the Client Network](#) (page 3-5)
- [Factory Network Settings](#) (page 3-9)
- [Port Assignments for Oracle Big Data Appliance Software](#) (page 3-10)
- [Preparing the Network for Oracle Big Data Appliance](#) (page 3-11)

3.1 Overview of Network Requirements

Oracle Big Data Appliance includes five or more servers and the equipment to connect the servers to your network. The network connections enable the servers to be administered remotely and enable 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 v3.1 for remote management (v3.0 for Sun Fire X4270 M2 servers)
- 4 10-Gigabit Ethernet ports (1-Gigabit Ethernet ports for Sun Fire X4270 M2)

3.1.1 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.

3.1.2 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.

**Note:**

Each network must be on a distinct and separate subnet from the other networks.

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 InfiniBand switches connected to the Ethernet switch in the rack. This network is also called the management network.

There are two uplinks to the administrative network:

- From the Ethernet switch in the rack
- From the KVM switch in the rack (Sun Fire X4270 M2-based racks only)

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.

Domain Name System (DNS) servers are not required on the administrative network, although Oracle Big Data Appliance uses them if they are available. At least one Network Time Protocol (NTP) server must also be available. The NTP server for the administrative network can be different from the NTP server for the client network.

All Oracle Big Data Appliance components on the administrative network must be on the same subnet; that is, the eth0 IP addresses for all servers, ILOMs, PDUs, InfiniBand switches, and the Cisco Ethernet switch must be on the same subnet.

- **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 the client network host names of the servers.

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.

At least one DNS server must be accessible on the client network. At least one NTP server must also be available. The NTP server for the client network can be different from the NTP server for the administrative network.

- **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 within Oracle Big Data Appliance and any other Oracle engineered systems that are connected to it with InfiniBand cables. This network does not connect to your existing network. It is automatically configured during installation.

3.1.3 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.
- `ILOM`: Provides access to Oracle Integrated Lights Out Manager (ILOM) using the administrative network.

Use the `bda-preinstall-checkip.sh` script to verify the readiness of the network environment before cabling Oracle Big Data Appliance to the existing network. See ["Preparing the Network for Oracle Big Data Appliance" \(page 3-11\)](#).

[Figure 3-1](#) (page 3-4) shows the network connections for Oracle Server X5–2L, Sun Server X4-2L, and Sun Server X3-2L-based racks.

Figure 3-1 Network Diagram for Oracle Big Data Appliance X5-2L, X4-2L, X3-2L





Figure 3-2 (page 3-5) shows the network connections for a Sun Fire X4270 M2 - based rack, which includes a KVM.

Figure 3-2 Network Diagram for Oracle Big Data Appliance (with KVM)



3.2 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. Oracle recommends that the two switches have the same number of active 10 GbE connections, so that failover does not result in a loss of available client network bandwidth.

3.2.1 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-passive 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-passive 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, suppose a customer creates three 10 GbE uplinks from each gateway switch in order to connect the servers for two clusters, each of five nodes. In this scenario, it is known that the majority of the network I/O will be on Cluster1. Therefore, when the customer runs the [Oracle Big Data Appliance Configuration Utility](#) (page 4-1), she or he can choose to allocate more ports to Cluster1 than to Cluster2. Based on that customer-defined configuration, Mammoth determines which ports are assigned to the cluster. For clusters using multiple ports, it round-robins the server assignments among those ports, also alternating which switch is active as it loops through the servers. This is demonstrated in the table below.

Table 3-1 Server/Gateway Connections Example

Cluster, Node	Active Link	Backup Link
Cluster1, Node1	Switch 1, Link 1 ¹	Switch 2, Link 1
Cluster1, Node2	Switch 1, Link 2	Switch 2, Link 2
Cluster1, Node3	Switch 2, Link 1	Switch 1, Link 1
Cluster1, Node4	Switch 2, Link 2	Switch 1, Link 2
Cluster1, Node5	Switch 1, Link 1	Switch 2, Link1
Cluster2, Node1	Switch 1, Link 3	Switch 1, Link 3
Cluster2, Node2	Switch 1, Link 3	Switch 1, Link 3
Cluster2, Node3	Switch 1, Link 3	Switch 1, Link 3
Cluster2, Node4	Switch 1, Link 3	Switch 1, Link 3
Cluster2, Node5	Switch 1, Link 3	Switch 1, Link 3

¹ All switches in this table are Sun Network QDR InfiniBand Gateway Switches

Depending on the number of 10 GbE connections, the client network traffic for multiple servers can 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.

3.2.2 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.

3.2.3 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.

3.2.4 Guidelines for Connecting to the Gateway Switches

Oracle recommends that you follow these guidelines for making 10 GbE connections to the two gateway switches in Oracle Big Data Appliance:

- Make the same number of 10 GbE connections to both gateway switches.
- To make 1 to 4 10-GbE connections to each switch, use a single QSFP splitter cable to the 0A-ETH port of both switches.
- To make 5 to 8 10-GbE connections to each switch, use two QSFP splitter cables to both 0A-ETH and 1A-ETH ports of both switches. Divide the connections as evenly as possible between the 2 splitter cables.
- To make multiple 10 GbE connections to a single QSFP splitter cable, start with the lowest number port and work upward. For example, for two connections, use 0A-ETH1 and 0A-ETH2.
- Make the 10 GbE connections to the same ports on both NM2-GW switches. If you make connections to 0A-ETH1 and 0A-ETH2 on one switch, then make connections only to 0A-ETH1 and 0A-ETH2 on the other switch.

3.2.4.1 Configuring for LACP/LAG Connections

A Link Aggregation Group (LAG) bonds two or more Ethernet connections together for increased uplink speed and greater resiliency.

LAGs impose the following constraints:

- A connector cannot belong to more than one LAG.
- A connector in a LAG cannot be used outside of the LAG.
- A connector in a LAG cannot be used in a VLAN or VNIC.
- A LAG can service only one rack.
- LAGs cannot be used with multirack Oracle Big Data Appliance clusters. You can use a LAG with single-rack clusters only. Within a rack, a LAG may be shared by multiple clusters that occupy the same rack.

When you pre-configure the networking for a cluster in the Oracle Big Data Appliance Configuration Utility, you can define a LAG by clicking the **Use LACP/LAG** checkbox and then selecting the connectors to include in the LAG.

If you define a LAG in the configuration utility, then related configuration changes to the cluster and the InfiniBand switches are done automatically when you run the network setup scripts. However, you must manually enable LAG mode on external switches connected to the InfiniBand switch.



Note:

If you set up LAGs in the configuration utility, then configure external switches to enable LAG before running the network setup scripts that are generated by the configuration utility.

The steps required to enable LAN will differ for various Ethernet switch models. The following is an example of how to use 802.3ad-based LACP on a Cisco 4948 switch. Refer to the product documentation provided with your own switch for the equivalent steps.

The steps in the example also use VLAN ID 1 and 2. Please adjust and modify these according to your own datacenter network design. This example assumes that Cisco ports 1/4 and 1/5 are connected to InfiniBand switch 10GbE uplink ports 1A-ETH-1 and 1A-ETH-2, respectively.

Example 3-1 LAG Configuration for External Cisco 4948 Ethernet Switch Connections to InfiniBand

The first set of commands in the example must be run in the order provided for Cisco ports 1/4 and 1/5, which are connected to the InfiniBand switch.

```
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,2
switchport mode trunk
no shutdown
channel-group 3 mode active
```

Below are commands you need to run in the order shown for Cisco Port-Chan 3.

```
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1,2
switchport mode trunk
```

See Also:

- [Oracle Big Data Appliance Configuration Utility](#) (page 4-16)

3.2.5 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 support a configuration in which the 10 GbE connections to the switches in a single rack are used for all servers in all Oracle Big Data Appliance racks.

3.3 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 3-2](#) (page 3-9) lists the default IP addresses for Oracle Big Data Appliance.

Table 3-2 Default IP Addresses for Oracle Big Data Appliance

Host	Administrative IP Addresses	Oracle ILOM IP Addresses	InfiniBand Bonded IP Addresses
bda18 ¹	192.168.1.18	192.168.1.118	192.168.10.18
bda17 ¹	192.168.1.17	192.168.1.117	192.168.10.17
bda16 ¹	192.168.1.16	192.168.1.116	192.168.10.16
bda15 ¹	192.168.1.15	192.168.1.115	192.168.10.15
bda14 ¹	192.168.1.14	192.168.1.114	192.168.10.14
bda13 ¹	192.168.1.13	192.168.1.113	192.168.10.13
bda12 ²	192.168.1.12	192.168.1.112	192.168.10.12
bda11 ²	192.168.1.11	192.168.1.111	192.168.10.11
bda10 ²	192.168.1.10	192.168.1.110	192.168.10.10
bdasw-ib3	192.168.1.203		
Cisco Switch	192.168.1.200		

Table 3-2 (Cont.) Default IP Addresses for Oracle Big Data Appliance

Host	Administrative IP Addresses	Oracle ILOM IP Addresses	InfiniBand Bonded IP Addresses
bdasw-ib2	192.168.1.202		
bda9 ²	192.168.1.9	192.168.1.109	192.168.10.9
bda8 ²	192.168.1.8	192.168.1.108	192.168.10.8
bda7 ²	192.168.1.7	192.168.1.107	192.168.10.7
bda6	192.168.1.6	192.168.1.106	192.168.10.6
bda5	192.168.1.5	192.168.1.105	192.168.10.5
bda4	192.168.1.4	192.168.1.104	192.168.10.4
bda3	192.168.1.3	192.168.1.103	192.168.10.3
bda2	192.168.1.2	192.168.1.102	192.168.10.2
bda1	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		

¹ Full racks only² Full racks or starter racks with additional servers only

3.4 Port Assignments for Oracle Big Data Appliance Software

[Table 3-3](#) (page 3-10) identifies the port numbers used by Oracle Big Data Appliance software.

For port numbers used by Cloudera CDH, CM, and other Cloudera products, see the [Cloudera documentation library](#).

Table 3-3 Oracle Big Data Appliance Port Numbers

Port	Used by
22	ssh
80	yumrepos (only during installation)
111	portmap
162	Auto Service Request Manager (optional)
443	Auto Service Request Manager (optional)
668	rpc.statd
3306	MySQL Database
5000	Oracle NoSQL Database registration
5001	Oracle NoSQL Database administration
5010 to 5020	Oracle NoSQL Database processes

Table 3-3 (Cont.) Oracle Big Data Appliance Port Numbers

Port	Used by
5042	Oracle Big Data SQL
6481	Auto Service Request Manager service tags listener (optional)
8139	Puppet nodes
8140	Puppet parent
20910	Oracle Data Integrator agent
30920	Automated Service Monitor (ASM)

The following table describes Oracle Big Data Appliance port access required by the Oracle Management Service (OMS) and Management Agent in order to use the Oracle Enterprise Manager plug-in.

Table 3-4 Port Access Required by the Oracle Enterprise Manager Plug-In

Port Number	On Which Node	Use
7183	Cloudera Manager Server host.	For discovery, if HTTPS access is enabled on Cloudera Manager.
7180	Cloudera Manager Server host.	For discovery and monitoring, if HTTPS access is not enabled on Cloudera Manager.
7187	Cloudera Manager Server host.	For monitoring, if HTTPS access is enabled on Cloudera Manager.
22	All nodes	SSH for bdacli execution during discovery.

3.5 Preparing the Network for Oracle Big Data Appliance

Oracle Big Data Appliance Configuration Generation Utility generates these files for you to use before installation:

- **bda-install-review.html**: Lists the host names and custom IP addresses for Oracle Big Data Appliance. Verify that the information is correct, and add the host names and IP addresses to a name server on the network. You can view this file in a browser.
- **bda-preinstall-checkip.sh**: Performs validation checks on network configuration values, including host names and IP addresses.

See "[Generating the Configuration Files](#) (page 4-3)."

Run the `bda-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.



Note:

Running `bda-preinstall-checkip.sh` is a critical pre-installation step. Network errors can cause extensive delays in the installation procedure.

To check the network for conflicts:

1. Copy the `bda-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 bda-preinstall-checkip.sh`
3. Execute the script:
`./bda-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 all reported errors before rerunning the script.

[Example 3-2](#) (page 3-12) shows success and error messages from the `bda-preinstall-checkip.sh` script.

Example 3-2 Output from `bda-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
```

4

Using Oracle Big Data Appliance Configuration Generation Utility

This chapter describes how to use Oracle Big Data Appliance Configuration Generation Utility.

This chapter contains the following topics:

- [Overview of Oracle Big Data Appliance Configuration Generation Utility](#) (page 4-1)
- [Generating the Configuration Files](#) (page 4-3)
- [About the Configuration Files](#) (page 4-3)
- [Validating the Network Settings](#) (page 4-5)
- [Customer Details Page](#) (page 4-5)
- [Hardware Selection Page](#) (page 4-7)
- [Rack Details](#) (page 4-7)
- [Networking Page](#) (page 4-8)
- [Administration Network Page](#) (page 4-8)
- [General Network Properties Page](#) (page 4-9)
- [Review and Edit Details Page](#) (page 4-10)
- [Define Clusters Page](#) (page 4-11)
- [Cluster Page](#) (page 4-12)
- [Complete Page](#) (page 4-25)

4.1 Overview of Oracle Big Data Appliance Configuration Generation Utility

Oracle Big Data Appliance Configuration Generation Utility acquires information from you, such as IP addresses and software preferences, that are required for deploying Oracle Big Data Appliance. After guiding you through a series of pages, the utility generates a set of configuration files. These files help automate the deployment process and ensure that Oracle Big Data Appliance is configured to your specifications.

Choose the option that describes the type of hardware installation you are configuring:

- Click **One or more new Big Data Appliance racks being installed** if you do not have a previously-generated master .xml configuration file to start with.

The other three choices activate the **Import** button, so that you can select the master.xml file from a previous configuration. After selecting one of these options, import the master .xml,

- Click **One or more Big Data Appliance racks being added to an existing group of Big Data Appliances** if you need to configure new rack hardware and do not intend to add servers or make any cluster changes at this time.
- Click **One or more nodes being added to a Big Data Appliance starter rack or partially filled rack** for in-rack or cluster expansion – adding new nodes to the existing racks and/or to existing clusters.
- Click **An in-process configuration using a saved master.xml configuration file** for general editing purposes.

Click **Next** to go to the Customer Details page. If you are working off of an existing master.xml file, fields are pre-populated and selectors are preset throughout the pages of the utility based on the data in the file.

Figure 4-1 (page 4-2) shows the Customer Details page of the Oracle Big Data Appliance Configuration Generation Utility.

Figure 4-1 Oracle Big Data Appliance Configuration Generation Utility

Customer Details

ORACLE
BIG DATA APPLIANCE

Welcome
Customer Details
Hardware Selection
Rack Details
Networking
Client Ethernet Network
Administration Network
InfiniBand Network
General Network Properties
Review and Edit Details
Define Clusters
Cluster 1
Complete

Customer Details

Customer Name: Example Inc.
Region: America
Time Zone: America/New_York

Base Name and Suffixes

Rack Base Name: bda Rack Start Index: 1
Server Base Name: node
Admin Name Suffix: -adm
Private Name Suffix: -priv
ILOM Name Suffix: -ilom
Switch Base Name: sw
Domain Name: example.com

Generated Names

Server Admin: bda1node01-adm Server Client: bda1node01
Server ILOM: bda1node01-ilom Server Private: bda1node01-priv
Ethernet Switch: bda1sw-ip InfiniBand Switch: bda1sw-ib

Help Import... < Back Next > Create Files Cancel

Note:

- Oracle Big Data Appliance uses Cloudera's Distribution including Apache Hadoop (CDH). A Hadoop cluster on Oracle Big Data Appliance is called a **CDH cluster**.
- The terms **appliance** and **rack** refer to Oracle Big Data Appliance, Oracle Big Data Appliance X3-2, Oracle Big Data Appliance X4-2, Oracle Big Data Appliance X5-2, and Oracle Big Data Appliance X6-2.

4.2 Generating the Configuration Files

The following procedures explain how to install Oracle Big Data Appliance Configuration Generation Utility and generate the configuration files.

To configure Oracle Big Data Appliance:

1. Download Oracle Big Data Appliance Configuration Generation Utility from Oracle Technology Network at

<http://www.oracle.com/technetwork/database/bigdata-appliance/downloads/index.html>

The file is named `BDAConfigurator-version.zip`. The utility requires Oracle JRE 1.7 or later.

2. Extract the files in `BDAConfigurator-version.zip`. For example:

```
$ unzip BDAConfigurator-<version>.zip
Archive:  BDAConfigurator-<version>.zip
  creating: BDAConfigurator-<version>/
  inflating: BDAConfigurator-<version>/exagen.jar
  inflating: BDAConfigurator-<version>/oracle_ice.jar
  inflating: BDAConfigurator-<version>/passwd.jar
  inflating: BDAConfigurator-<version>/orai18n-utility.jar
  .
  .
  .
```

3. Change to the `BDAConfigurator-version` directory.
4. Run Oracle Big Data Appliance Configuration Generation Utility.
 - On Linux:

```
$ sh bdaconf.sh
```
 - On Microsoft Windows, double-click `bdaconf.cmd` in Windows Explorer, or run the file from the command line:

```
C:\> bdaconf.cmd
```
5. On the Welcome page, select a configuration type.
6. Click **Import** if the button is activated, and select a previously saved `master.xml` configuration file.
7. Follow the steps of the wizard. On the Complete page, click **Create Files**.
8. Validate the network configuration.
See "[Validating the Network Settings](#) (page 4-5)."
9. Send the generated `bda.zip` file to your Oracle representative.

4.3 About the Configuration Files

Oracle Big Data Appliance Configuration Generation Utility generates the following files to use when you configure the system. You can select the directory where they are saved.

This is the basic structure of the directory:

```
company_name /
  bda-timestamp.zip
  bda-install-preview.html
  bda-preinstall-checkip.sh
rack_name /
  rack_name-network.json
  rack_name-rack-network.json
cluster_name /
  cluster_name-config.json
  cluster_name-cluster-network.json
master.xml
```

bda-timestamp.zip

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

bda-install-preview.html

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

bda-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.

rack_name-network.json

Contains the network configuration for a full rack, a starter rack, or a starter rack with additional servers. It contains information about all the servers, switches, and PDUs.

cluster_name-config.json

Contains all the information for a cluster, including the network configuration, port numbers, user names, and passwords. The configuration utility creates a separate parameter file for each cluster. If several clusters are being configured, then each parameter file is located in a separate subdirectory.

If additional servers are configured over a starter rack as an addition to an existing cluster, then the configuration utility does not generate a parameter file; the Mammoth utility generates it.

rack_name-rack-network.json

Contains the Administrative Network information for the entire rack.

cluster_name-cluster-network.json

Contains the client and private InfiniBand network information for the cluster created.

master.xml

Contains all the configuration settings in XML format so that Oracle Big Data Appliance Configuration Generation Utility can read it. To alter the configuration of an Oracle Big Data Appliance deployment, you can load this file, enter the changes, and regenerate the configuration files.

This file is used only by Oracle Big Data Appliance Configuration Generation Utility. It is not used for the actual configuration of Oracle Big Data Appliance.

4.4 Validating the Network Settings

Validating the network settings before the Oracle Big Data Appliance hardware arrives at your site is a critical step. Network problems can cause extended delays in the installation.

To validate the network configuration settings:

1. Copy the `bda-preinstall-checkip.sh` file generated by Oracle Big Data Appliance Configuration Generation Utility to a Linux host on the same network that Oracle Big Data Appliance will use.

2. Log in to the Linux host and run `bda-preinstall-checkip.sh`:

```
$ sh bda-preinstall-checkip.sh
```

This script checks your existing network for conflicts with the Oracle Big Data Appliance IP address pool.

3. Correct any network problems discovered by the script before installation begins. Network problems during the installation can cause extensive delays.
4. Run the network connections to the planned location for Oracle Big Data Appliance.
5. Inform your Oracle representative when you have completed these steps.

4.5 Customer Details Page

The following table describes the customer details fields. The Generated Names shows the names created from the entered values.

Table 4-1 Customer Details Page

Customer Details Field	Description
Customer Name	The name of your enterprise. Required.
Region	The geographic area where Oracle Big Data Appliance will be installed.
Time Zone	The time zone for your installation. You must select the appropriate region before selecting the time zone.
Rack Base Name	A maximum of 10 alphanumeric characters for the name of the Oracle Big Data Appliance rack.
Rack Start Index	A digit that uniquely identifies the rack. It is a suffix of the rack base name.
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: <code>eth0</code> , <code>bondib0</code> , <code>bondeth0</code> , and Oracle ILOM. For example, a rack base name of <code>bda</code> , a rack start index of <code>1</code> , and a server base name of <code>node</code> results in host names of <code>bda1node01</code> , <code>bda1node02</code> , and so forth.
Admin Name Suffix	Suffix to the basic host name to form the <code>eth0</code> host names
Private Name Suffix	Suffix to the basic host name to form the <code>bondib0</code> host name

Table 4-1 (Cont.) Customer Details Page

Customer Details Field	Description
ILOM Name Suffix	Suffix to the basic host name to form the Oracle ILOM name
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 bda1sw-ip, bda1sw-ib1, and so forth. Important: Ensure that the fully constructed switch names (such as bda1sw-ib1) do not exceed 20 characters.
Domain Name	Name of the domain in which Oracle Big Data Appliance operates. Required.

4.5.1 Using Standardized Host Names

The rack name is used in the assignment of standardized host names for all Oracle Big Data Appliance servers. The host name for all servers on the client network is in this format:

racknameNodeNN.domain

In this syntax:

- *NN* is the position number of the server node in the rack (01 to 18).
- *domain* is the domain name.

Host names must have fewer than 16 characters, which can be ASCII lower-case letters (a to z), numbers (0 to 9), and hyphens (-) only. Do not begin or end the name with a hyphen.

4.5.2 Using Customized Host Names

You can change the suffixes used for the different network interfaces.

You can enter server host names that do not follow the naming conventions on the [Review and Edit Details Page](#) (page 4-10).

4.5.3 Naming Multirack Clusters

Oracle recommends that for a cluster of multiple racks, you use the cluster name as the rack name.

For example, in a three-rack cluster, if the cluster name is `cluster1` and the domain name is `example.com`, then the fully qualified host name of the server at the bottom of the first rack is `cluster101node01.example.com`. For the top server in the third rack of this cluster, the host name is `cluster103node18.example.com`.

4.5.4 Connecting to Oracle Big Data Appliance Over the Networks

The host names on the other networks have a short extension that follows the unit number. If you retained the default extensions, then use these formats to connect to Oracle Big Data Appliance after it is connected to the network:

- For short host names over the administration network:
racknameNodeNN-adm
- For the private InfiniBand network host names:
racknameNodeNN-priv
- For the Oracle Integrated Lights Out Manager (ILOM) host names:
racknameNodeNN-ilom
- For the switch host names:
racknamesw-ibM

In this syntax:

- *NN* is the position number of the server in the rack (01 to 18).
- *M* is 1, 2, or 3, depending on the switch location in the rack.

4.6 Hardware Selection Page

The hardware page identifies one or more racks that you want to deploy at the same time. The racks must be cabled together.

For example, if you are deploying three full racks, then add Full Rack three times to your deployment.

The following table describes the hardware selection choices.

Table 4-2 Hardware Selection Page

Hardware Selection Field	Description
Select interconnected hardware to deploy	Lists the available hardware configurations. Choose one or more racks. You can choose the same type of rack multiple times. • Full rack: Contains 18 servers. • Starter rack: Contains 6 servers. • Partially filled rack: Contains any number of servers between 7 and 17.
This is your deployment	Lists the hardware selected for your site.
Will you use non-Oracle PDUs?	Oracle strongly recommends that the PDUs that are shipped with Oracle Big Data Appliance be used to supply power to its servers and switches. However, if your data center has specific requirements that the Oracle-supplied PDUs do not meet, then you can use other PDUs to supply power to the Oracle Big Data Appliance servers and switches. Oracle does not support customer-supplied PDUs, and Oracle Enterprise Manager Cloud Control does not monitor them.

4.7 Rack Details

The rack details page identifies the optional network connections for an Oracle Big Data Appliance rack.

The following table describes the rack detail fields.

Table 4-3 Rack Details Page

Rack Detail Field	Description
Rack Name	Enter the name of the rack.
Number of 10 GbE Connections	The two Sun Network QDR InfiniBand Gateway switches in every Oracle Big Data Appliance must have an equal number of 10 GbE links to the client network. Each gateway switch supports up to eight 10-GbE links, for a total of 16 links. Each server is assigned to one 10-Gbe link on each gateway switch. Therefore, an Oracle Big Data Appliance full rack or starter rack with an extension kit can use up to 16 10-GbE links. A starter rack can use up to 12 10-GbE links. Oracle recommends using as many 10 GbE links to the Oracle Big Data Appliance rack as the data center can support. The additional links increase the network bandwidth available between Oracle Big Data Appliance and the client network, and reduces the impact if a link fails. Oracle uses this information to ensure that the correct number of cables is ordered or available at the installation site. It also lets the Oracle field engineer plan how many 10 GbE links to connect to the Oracle Big Data Appliance rack.
BDA will be connected via InfiniBand to any Oracle engineered systems	Select this option if you are connecting this rack to another rack through the InfiniBand fabric. When connecting multiple racks, ensure the following: • The InfiniBand IP addresses of all servers are unique, including the servers in other Oracle engineered systems. • All InfiniBand IP addresses are on the same network. For example, if you connect an Oracle Big Data Appliance rack to Oracle Exadata Database Machine, then you must use the same netmask on the InfiniBand networks for both systems. Moreover, after you apply this netmask to the InfiniBand IP addresses of the Exadata database servers, the Exadata storage servers, and the Oracle Big Data Appliance servers, all IP addresses are in the same subnet."

4.8 Networking Page

The networking page identifies the number of IP addresses required for each network. The administration, client Ethernet, and InfiniBand networks are required. You must allocate the specified number of IP addresses for them.

Each IP address pool initially contains a range of consecutive IP addresses. If some IP addresses in the range are not available, then you can change individual addresses on the [Review and Edit Details Page](#) (page 4-10).

4.9 Administration Network Page

The values that you enter on this page are used to add the Oracle Big Data Appliance servers, switches, and PDUs to your existing administration network. Each server has two network interfaces for administration. One interface provides access to the operating system, and the other provides access to Oracle Integrated Lights Out Manager (ILOM).

 Note:

Each network must be on a separate subnet from the other networks.

The following table describes the administration network fields.

Table 4-4 Administration Network Page

Administration Network Field	Description
Starting IP Address for Pool	The first IP address on the administration network available for use by Oracle Big Data Appliance.
Pool Size	The required number of IP addresses on the administration network. The IP addresses for a rack are assigned in this order: Oracle Big Data Appliance servers, Oracle ILOMs, Ethernet switch, spine switch, leaf switches (2), and PDUs (2).
Ending IP Address for Pool	The last IP address on the administration network assigned to Oracle Big Data Appliance. The value in this field is automatically calculated from the starting IP address and the pool size. Ensure that all of the IP addresses are available for use in the pool defined by the starting and ending addresses. If they are not available, then you can either assign a different range or manually change individual IP addresses on the Review and Edit Details Page (page 4-10).
Subnet Mask	The subnet mask for the administration network.
Gateway	The IP address for the gateway. The gateway IP address is generated automatically, so verify that it is correct.

4.10 General Network Properties Page

The client and administration networks typically use the same Domain Name System (DNS) and Network Time Protocol (NTP) servers. If they are different on your networks, then enter the values for the client network first.

The following table describes the general network properties.

Table 4-5 General Network Properties Page

General Network Properties Field	Description
DNS Servers	Up to six IP addresses for the DNS servers. At least one DNS server must be accessible on the client network. DNS servers are not required on the administration network, although Oracle Big Data Appliance uses them if they are available.

Table 4-5 (Cont.) General Network Properties Page

General Network Properties Field	Description
NTP Servers	Up to six IP addresses for the NTP servers. Both the client network and the administration network must have access to at least one NTP server. The NTP servers for the administration network can be different from the NTP servers for the client network. If they are, then identify the NTP servers for the client network in this field.
Search Domains	Up to six domain names in which Oracle Big Data Appliance operates, such as <code>example.com</code> and <code>us.example.com</code>
Are administration host name entries in DNS?	Select Yes or No: • Yes: Administration host name DNS entries are validated during preinstall checks and network configuration. • No: Administration host name DNS entries are not validated during preinstall checks and network configuration.

The next table describes the network properties for the administration network if they are different from the client network. Typically, these properties are the same for both networks. They default to the values you entered for the general network properties.

Table 4-6 Advanced Network Properties

Advanced Network Properties Field	Description
Advanced Network Configuration	Select this option if the client and administration networks are isolated on your system and use different DNS and NTP servers, different domains, or both. You can then complete the fields for the administration network.
Admin Domain Name	The name of the administration domain.
Admin DNS Servers	Up to six IP addresses for the administration Domain Name System server, if they are different from the client network.
Admin NTP Servers	Up to six IP addresses for the administration Network Time Protocol server, if they are different from the client network.
Admin Search Domains	Up to six domain names in which Oracle Big Data Appliance administration network operates, if they are different from the client network.

4.11 Review and Edit Details Page

Use this page to review and modify the network configuration settings.

If you specified a range of IP addresses for any of the networks that includes addresses already in use, then replace those IP addresses on this page. Otherwise, the network configuration of Oracle Big Data Appliance will fail, causing unnecessary delays. When you are done making changes, click **Regenerate using changed base values**.

4.12 Define Clusters Page

Use the define clusters page to identify the number of clusters to create and the servers that compose each cluster. You can configure clusters for either CDH or Oracle NoSQL Database on this page. A cluster can occupy a single rack or span multiple racks.

Note:

In a multi-cluster environment, it is not required that all clusters uniformly run the same version of Oracle Big Data Appliance. Multiple clusters running different versions of the software on different racks or the same rack are supported configurations.

The minimum cluster size for each type is as follows:

- CDH cluster – five servers minimum is recommended for production in order to support high availability. Clusters as small as three servers can be created. These are generally recommended for development only. However the five server minimum for clusters in production is a guideline, not a rule, and has no implications for Oracle support.
- NoSQL Database cluster – three servers minimum

For each rack contributing servers to a cluster of either type, the minimum number of servers is three.

The following table describes the cluster definition choices.

Table 4-7 Define Clusters Page

Define Clusters Field	Description
Number of clusters to create	Select the number of clusters. For each cluster, a new tab appears on the page, and a new page appears in Oracle Big Data Appliance Configuration Generation Utility. Be sure to complete all tabs before continuing to the next page.
Cluster Name	Enter a unique name for the cluster. The name must begin with a letter and can consist of alphanumeric characters and dashes (-). Underscores, () and other non-alphanumeric characters are not accepted.
Cluster Type	Choose the type of cluster: • CDH cluster: Installs Cloudera's Distribution including Apache Hadoop and optional software on cluster of new servers • NoSQL DB cluster: Installs Oracle NoSQL Database on a cluster of new servers • Adding to existing cluster: Installs the same software on the new servers as the rest of the cluster.
Unassigned Servers	From the list on the left, select the servers for the cluster and move them to the list of assigned servers on the right.

Table 4-7 (Cont.) Define Clusters Page

Define Clusters Field	Description
<i>Assigned Servers</i>	Lists the servers selected for the cluster. A CDH cluster must consist of at least three servers. A minimum of five is recommended for a production environment. Oracle NoSQL Database clusters must consist of a minimum of three servers for development or production.

4.13 Cluster Page

Select the software to install on this cluster. The fields displayed on this field depend on the type of cluster being configured:

- [A New Oracle NoSQL Database Cluster](#) (page 4-12)
- [A New CDH Cluster](#) (page 4-16)

4.13.1 A New Oracle NoSQL Database Cluster

For each new Oracle NoSQL Database cluster, the Configuration Utility provides a page where you configure the networking for the cluster. On this page, provide the following general information about the cluster's client network:

- Cluster Name
- DNS Servers
- NTP Servers
- Search Domains
- Domain Name
- Region
- Time Zone

Also provide the configuration details requested in the following subsections of the page:

- [User and Groups](#) (page 4-12)
- [Client Network](#) (page 4-17)
- [InfiniBand Network](#) (page 4-18)
- [MIT Kerberos](#) (page 4-15)
- [Auto Service Request](#) (page 4-23)
- [Installed Components](#) (page 4-15)
- [Oracle NoSQL Configuration](#) (page 4-16)

4.13.1.1 User and Groups

The following table describes the user name, groups, and password fields for a new Oracle NoSQL Database cluster. Passwords are optional, but you must enter them during the software installation if you do not provide them here.

Table 4-8 User and Groups for a New Oracle NoSQL Database Cluster

User/Groups Field	Description
OS password for root user	The root password on all servers in the cluster.
OS password for oracle user	The oracle password on all servers in the cluster. Oracle applications run under this identity.
oracle user ID	The ID number of the oracle user. It must match the oracle UID of a connected Oracle Exadata Database Machine. The UID ensures that the oracle user can log in from Oracle Big Data Appliance to the correct account in Oracle Database. Required.
oinstall group ID	The ID number of the Oracle Inventory Group (oinstall). It must match the oinstall group ID of a connected Oracle Exadata Database Machine.
dba group ID	The ID number of the dba group. It must match the dba group ID of a connected Oracle Exadata Database Machine.

4.13.1.2 Client Network

The following table describes the fields of the Client Network panel on the Cluster Page of the Oracle Big Data Configuration Utility.

Note that Client and private InfiniBand networks within a cluster may not use the same subnet. However, subnets may be shared by client networks on different clusters.

Table 4-9 Client Network

Client Network Field	Description
Starting IP Address for Pool	The first IP address on the client network you are configuring.
Pool Size	The required number of IP addresses (one for each node in the current cluster).
Ending IP Address for Pool	The last IP address on the client network. This is calculated automatically based on the starting address and pool size. Ensure that all of the IP addresses are available for use in the pool defined by the starting and ending addresses. If not all are available, then you can either assign a different range or manually change individual IP addresses on the Client and InfiniBand Network page.
Subnet Mask	The subnet mask of the client network on the current cluster.
Gateway	The IP address of the client network on the current cluster.
VLAN ID	Note that the default VLAN ID value is NO.
Domain	FQDN of the domain hosting the network.

Table 4-9 (Cont.) Client Network

Client Network Field	Description
Use LACP/LAG	Click this checkbox to create a Link Aggregation Group for client networks. Note that after a LAG is created, the connectors in the LAG are not available for individual selection, nor inclusion in other LAGs or for use in a VLAN or VNIC. LAG setup on the Oracle Big Data Appliance side is automated if you select this option and assign connectors to the LAG. However, you must manually enable LAG mode on external switches that will communicate with the InfiniBand switch over the LAG. See Configuring for LACP/LAG Connections (page 3-8)
Connectors	You must check at least one connector from the eight connectors for each cluster's client network. This provides each cluster with least two connectors. (One connector is dedicated to failover and is located on the secondary InfiniBand switch.) Make the 10 GbE connections to the same ports on both switches.
Add a Client to this Cluster	Add another client network to this cluster. Click this button to create another set of input fields. Add the necessary identifiers to these fields in order to configure the additional client network.
Remove a Client from this Cluster	Removes the last additional client from the list of client networks.

4.13.1.3 InfiniBand Network

The InfiniBand network connects the Oracle Big Data Appliance servers within a rack. It can also connect multiple racks to form a multirack Hadoop cluster, or to provide access to Oracle Big Data Appliance from Oracle Exadata Database Machine.

Although you can configure multiple client networks within a cluster, each cluster can support only one InfiniBand network.

The following table describes the InfiniBand network fields.

Table 4-10 InfiniBand Network Page

InfiniBand Network Field	Description
Starting IP Address for Pool	The first IP address on the private InfiniBand network available for use by the Oracle Big Data Appliance servers. The default is 192.168.10.1.
Pool Size	The required number of IP addresses. All Oracle Big Data Appliance servers require an IP address on the InfiniBand network. The pool size is calculated for the deployment that you identified on the Hardware Selection Page (page 4-7).

Table 4-10 (Cont.) InfiniBand Network Page

InfiniBand Network Field	Description
Ending IP Address for Pool	The last IP address assigned to the InfiniBand network for this deployment. This address is automatically calculated from the starting IP address and the pool size. Ensure that all of the IP addresses are available for use in the pool defined by the starting and ending addresses. If they are not available, then you can either assign a different range or manually change individual IP addresses on the Review and Edit Details Page (page 4-10).
Subnet Mask	The subnet mask for the InfiniBand network. The default is 255.255.252.0.
PKey	The unique ID of the InfiniBand partition used by this network.
PName	The partition name.
PType	The partition membership type (FULL or LIMITED).
MTU	Maximum Transmission Unit for this network.

4.13.1.4 MIT Kerberos

MIT Kerberos authentication is a security option for NoSQL clusters.

 Note:

The MIT Kerberos option is only available if you choose to install Oracle NoSQL Enterprise Edition. This option is not available for Oracle NoSQL Community Edition.

The following table describes the MIT Kerberos fields.

Table 4-11 MIT Kerberos

Kerberos Field	Description
Enable MIT Kerberos-based authentication?	Select this option to support MIT Kerberos on Oracle Big Data Appliance.
Kerberos KDC database password	A password for the KDC database, if it is being created on Oracle Big Data Appliance.
Kerberos realm	Enter the name of the realm for Oracle Big Data Appliance, such as EXAMPLE.COM.

If MIT Kerberos is selected, the configurator will create a default `admin` user with administrator privileges. The same password is used for the Kerberos KDC. The password may not include the dollar sign or backslash characters (`$`, `\`).

4.13.1.5 Installed Components

You can install either Community Edition or Enterprise Edition of Oracle NoSQL Database.

The following table describes these choices.

Table 4-12 Installed Components

Component Field	Description
Oracle NoSQL Database Edition	Choose between Community Edition and Enterprise Edition: • Community Edition is included in the license for Oracle Big Data Appliance. • Enterprise Edition requires a separate license. You must have this license to install Enterprise Edition on Oracle Big Data Appliance.

4.13.1.6 Oracle NoSQL Configuration

Oracle NoSQL Database 12c Release 1.3.0.5 and later versions support secondary zones, which are composed of nodes that function only as replicas. You can use the secondary zones on Oracle Big Data Appliance to maintain extra copies of the data for increased redundancy and read capacity, or to provide low latency, read access to data at a distant location.

The following table describes the modifiable configuration settings.

Table 4-13 Oracle NoSQL Configuration

Configuration Field	Description
Oracle NoSQL store name	A unique name for the KV store on Oracle Big Data Appliance. The default value is BDAKV1.
Oracle NoSQL primary zone name	The name of the primary zone for all nodes in the cluster. The default value is BDAKV1_PRIMARY_ZN.
Oracle NoSQL primary zone replication factor	A replication factor of 1 or more for the primary zone. The default value is 3.

4.13.2 A New CDH Cluster

For each new cluster, the Configuration Utility provides a page where you configure the networking for the cluster. On this page, provide the following general information about the cluster's client network:

- Cluster Name
- DNS Servers
- NTP Servers
- Search Domains
- Domain Name
- Region
- Time Zone

Also provide the configuration details requested in the following subsections of the page:

- [User/Groups](#) (page 4-17)

- [Client Network](#) (page 4-17)
- [InfiniBand Network](#) (page 4-18)
- [Big Data SQL](#) (page 4-19)
- [Big Data Connectors](#) (page 4-20)
- [MIT Kerberos](#) (page 4-20)
- [Active Directory Kerberos](#) (page 4-21)
- [HDFS Transparent Encryption](#) (page 4-22)
- [Audit Vault](#) (page 4-22)
- [Auto Service Request](#) (page 4-23)
- [Enterprise Management Cloud Control](#) (page 4-24)
- [Email Alerting](#) (page 4-24)

4.13.2.1 User/Groups

The following table describes the user name, groups, and password fields for a new CDH cluster. Passwords are optional, but you must enter them during the software installation if you do not provide them here.

Table 4-14 User /Groups for a New CDH Cluster

User/Groups Field	Description
OS password for root user	The root password on all servers in the cluster.
OS password for oracle user	The oracle password on all servers in the cluster. Oracle applications run under this identity.
oracle user ID	The ID number of the oracle user. It must match the oracle UID of a connected Oracle Exadata Database Machine. The UID ensures that the oracle user can log in from Oracle Big Data Appliance to the correct account in Oracle Database. Required.
oinstall group ID	The ID number of the Oracle Inventory Group (oinstall). It must match the oinstall group ID of a connected Oracle Exadata Database Machine.
dba group ID	The ID number of the dba group. It must match the dba group ID of a connected Oracle Exadata Database Machine.
Cloudera Manager admin password	The password for the admin user for Cloudera Manager. Available only for CDH clusters.
MySQL admin password	The password for the MySQL Database administration user. Available only for CDH clusters.

4.13.2.2 Client Network

The following table describes the fields of the Client Network panel on the Cluster Page of the Oracle Big Data Configuration Utility.

Note that Client and private InfiniBand networks within a cluster may not use the same subnet. However, subnets may be shared by client networks on different clusters.

Table 4-15 Client Network

Client Network Field	Description
Starting IP Address for Pool	The first IP address on the client network you are configuring.
Pool Size	The required number of IP addresses (one for each node in the current cluster).
Ending IP Address for Pool	The last IP address on the client network. This is calculated automatically based on the starting address and pool size. Ensure that all of the IP addresses are available for use in the pool defined by the starting and ending addresses. If not all are available, then you can either assign a different range or manually change individual IP addresses on the Client and InfiniBand Network page.
Subnet Mask	The subnet mask of the client network on the current cluster.
Gateway	The IP address of the client network on the current cluster.
VLAN ID	Note that the default VLAN ID value is N0.
Domain	FQDN of the domain hosting the network.
Use LACP/LAG	Click this checkbox to create a Link Aggregation Group for client networks. Note that after a LAG is created, the connectors in the LAG are not available for individual selection, nor inclusion in other LAGs or for use in a VLAN or VNIC. LAG setup on the Oracle Big Data Appliance side is automated if you select this option and assign connectors to the LAG. However, you must manually enable LAG mode on external switches that will communicate with the InfiniBand switch over the LAG. See Configuring for LACP/LAG Connections (page 3-8)
Connectors	You must check at least one connector from the eight connectors for each cluster's client network. This provides each cluster with least two connectors. (One connector is dedicated to failover and is located on the secondary InfiniBand switch.) Make the 10 GbE connections to the same ports on both switches.
Add a Client to this Cluster	Add another client network to this cluster. Click this button to create another set of input fields. Add the necessary identifiers to these fields in order to configure the additional client network.
Remove a Client from this Cluster	Removes the last additional client from the list of client networks.

4.13.2.3 InfiniBand Network

The InfiniBand network connects the Oracle Big Data Appliance servers within a rack. It can also connect multiple racks to form a multirack Hadoop cluster, or to provide access to Oracle Big Data Appliance from Oracle Exadata Database Machine.

Although you can configure multiple client networks within a cluster, each cluster can support only one InfiniBand network.

The following table describes the InfiniBand network fields.

Table 4-16 InfiniBand Network Page

InfiniBand Network Field	Description
Starting IP Address for Pool	The first IP address on the private InfiniBand network available for use by the Oracle Big Data Appliance servers. The default is 192.168.10.1.
Pool Size	The required number of IP addresses. All Oracle Big Data Appliance servers require an IP address on the InfiniBand network. The pool size is calculated for the deployment that you identified on the Hardware Selection Page (page 4-7).
Ending IP Address for Pool	The last IP address assigned to the InfiniBand network for this deployment. This address is automatically calculated from the starting IP address and the pool size. Ensure that all of the IP addresses are available for use in the pool defined by the starting and ending addresses. If they are not available, then you can either assign a different range or manually change individual IP addresses on the Review and Edit Details Page (page 4-10).
Subnet Mask	The subnet mask for the InfiniBand network. The default is 255.255.252.0.
PKey	The unique ID of the InfiniBand partition used by this network.
PName	The partition name.
PType	The partition membership type (FULL or LIMITED).
MTU	Maximum Transmission Unit for this network.

4.13.2.4 Big Data SQL

You can install Oracle Big Data SQL on a CDH cluster. You must have a separate license for this product. The following table describes the Oracle Big Data SQL fields.

Table 4-17 Big Data SQL

Big Data SQL Field	Description
Is Big Data SQL licensed?	Oracle Big Data SQL supports queries from Oracle Database against data stored in Hive and HDFS on Oracle Big Data Appliance. This technology requires a separate license. If you have a license, choose Yes .
Install Big Data SQL?	Choose Yes to install Oracle Big Data SQL 2.0 during the initial software release. Oracle Big Data Appliance must be connected over the InfiniBand network to Oracle Exadata Database Machine. Choose No to defer this installation to a later date. If you opt to install Big Data SQL 3.0.1 instead of 2.0, then choose No . Big Data SQL 3.0.1 can be installed via patch after the installation. It is not a Mammoth installation option. See the Oracle Big Data SQL Master Compatibility Matrix (Doc ID 2119369.1) in My Oracle Support.
Is Big Data SQL using InfiniBand?	Choose Yes to use the InfiniBand network. Choose No to use the Client network. The default is Yes . Note: Ethernet networking for Oracle Big Data SQL between Oracle SPARC SuperCluster and Oracle Big Data Appliance is not supported at this time.

4.13.2.5 Big Data Connectors

You can install Oracle Big Data Connectors on a CDH cluster. You must have a separate license for this product. The following table describes the installed components field.

Table 4-18 Installed Components

Component 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 . This choice activates the other fields.
Install Oracle Data Integrator Agent?	The agent supports Oracle Data Integrator, which is one of the Oracle Big Data Connectors. To configure Oracle Data Integrator for use immediately, choose Yes . You must have a license for Oracle Big Data Connectors.
MySQL password for Oracle Data Integrator Agent?	The password for the Oracle Data Integrator user in MySQL Database.

4.13.2.6 MIT Kerberos

MIT Kerberos authentication is a security option for CDH clusters. It is included with your Oracle Big Data Appliance license.

To use a key distribution center (KDC) elsewhere on the network (that is, not on Oracle Big Data Appliance), you must complete several steps before installing the software. See "[Installation Prerequisites](#) (page 10-3)."

The following table describes the MIT Kerberos fields.

Table 4-19 MIT Kerberos

Kerberos Field	Description
Enable MIT Kerberos-based authentication?	Select this option to support MIT Kerberos on Oracle Big Data Appliance.
Set up key distribution center on BDA	Choose Yes to set up a key distribution center (KDC) on Oracle Big Data Appliance. Otherwise, a KDC must always be available on the network to all clients.
Kerberos KDC database password	A password for the KDC database, if it is being created on Oracle Big Data Appliance.

Table 4-19 (Cont.) MIT Kerberos

Kerberos Field	Description
Non-BDA key distribution center hosts	List the fully qualified names or the IP addresses of the KDCs, available on the same network, that can serve as either the primary or backup KDC for Oracle Big Data Appliance.
 Important: The first external KDC server in the list (the one at the top of the list), must be the same server identified as the Kerberos Admin Server in <code>Krb5.conf</code>.	
Kerberos realm	Enter the name of the realm for Oracle Big Data Appliance, such as <code>EXAMPLE.COM</code> .
Encrypt HDFS Data Transport	Enables encryption of data transfer between DataNodes and clients, and among DataNodes.
Encrypt Hadoop Services	Enables SSL encryption for HDFS, MapReduce, and YARN web interfaces, as well as encrypted shuffle for MapReduce and YARN. It also enables authentication for access to the web console for the HDFS, MapReduce, and YARN roles.
Enable Sentry Authorization	Select this option to use Apache Sentry to provide fine-grained authorization to data stored in Hadoop.

4.13.2.7 Active Directory Kerberos

Active Directory Kerberos authentication is a Windows-based security option for CDH clusters. It is included with your Oracle Big Data Appliance license.

The following table describes the Active Directory Kerberos fields.

Table 4-20 Active Directory Kerberos

Kerberos Field	Description
Enable Active Directory Kerberos-based authentication?	Select this option to support Active Directory Kerberos on Oracle Big Data Appliance.
Active Directory Host	The fully qualified name or IP address of the Active Directory host.
Kerberos realm	Enter the name of the realm for Oracle Big Data Appliance, such as <code>EXAMPLE.COM</code> .
Active Directory Domain	Domain distinguished name where all the accounts are created on the Active Directory, such as <code>ou=bda,DC=example,DC=com</code> .
Active Directory Account Prefix	Add an optional prefix for all Active Directory account names in this domain. 10 characters maximum length. Restricted to <code>aA-zZ, 0-9, -, _</code> .

Table 4-20 (Cont.) Active Directory Kerberos

Kerberos Field	Description
Active Directory admin user	Name of the Active Director Admin user that can create, delete, and modify accounts on the Active Directory.
Active Directory admin password	Password for the Active Director Admin user.
Encrypt HDFS Data Transport	Enables encryption of data transfer between DataNodes and clients, and among DataNodes.
Encrypt Hadoop Services	Enables SSL encryption for HDFS, MapReduce, and YARN web interfaces, as well as encrypted shuffle for MapReduce and YARN. It also enables authentication for access to the web console for the HDFS, MapReduce, and YARN roles.
Enable Sentry authorization	Select this option to use Apache Sentry to provide fine-grained authorization to data stored in Hadoop.

4.13.2.8 HDFS Transparent Encryption

HDFS Transparent Encryption is a security option for data at-rest in CDH clusters.

HDFS Transparent Encryption requires that you enable Kerberos. It works with either Active Directory Kerberos or MIT Kerberos. (See "[Installation Prerequisites](#) (page 10-3).")

The following table describes the HDFS Transparent Encryption fields.

Table 4-21 HDFS Transparent Encryption

HDFS Transparent Field	Description
Enable HDFS Transparent Encryption.	Select this option to enable encryption of Hadoop data at rest on the Oracle Big Data Appliance.
Set up Key Trustee Servers on the BDA	Select this option to set up Key Trustee Servers internally on BDA cluster. If you enable this option, the remaining key trustee setup parameters are not needed. Mammoth will configure the installation.
Active Navigator Key Trustee Server	Manages encryption keys, certificates, and passwords. In the current implementation, this must be a server external to the Oracle Big Data Appliance.
Passive Navigator Key Trustee Server	The passive Navigator Key Trustee Server is a backup to ensure that this service does not become a single point of failure.
Key Trustee Organization	Key Trustee organization name configured by the Key Trustee Server administrator. This is the organization that keys are stored against.
Key Trustee Authorization Code	The authorization code that corresponds to the Navigator Key Trustee organization. This code is used by Cloudera Manager to authenticate against the Key Trustee Server.

4.13.2.9 Audit Vault

You can configure CDH clusters on Oracle Big Data Appliance as secured targets for Oracle Audit Vault and Database Firewall. The Audit Vault plug-in on Oracle Big Data

Appliance collects audit and logging data from MapReduce, HDFS, and Oozie services. You can then use Audit Vault Server to monitor these services on Oracle Big Data Appliance.

Oracle Audit Vault and Database Firewall Server Release 12.1.1 must be up and running on a separate server on the same network as Oracle Big Data Appliance before you perform the actual configuration. No other version of Oracle Audit Vault and Database Firewall Server is supported on Oracle Big Data Appliance.

The following table describes the Audit Vault fields.

Table 4-22 Audit Vault

Audit Vault Field	Description
Enable Audit Vault	Select this option to support Oracle Audit Vault and Database Firewall on Oracle Big Data Appliance.
Audit Vault server	The IP address of the Audit Vault server.
Audit Vault port	The port number that Audit Vault Server listens on.
Audit Vault database service name	The database service name for Audit Vault Server.
Audit Vault admin user	The name of the Audit Vault administration user.
Audit Vault admin user password	The password for the administration user.

4.13.2.10 Auto Service Request

ASR monitors the health of Oracle Big Data Appliance hardware and automatically submits a service request when it detects a fault. Although you can opt out of this program, Oracle recommends that you enable ASR.

ASR Manager must be installed and configured to run on a separate server outside of Oracle Big Data Appliance before the software is installed and configured on Oracle Big Data Appliance. The software installation fails with an error if Enable Auto Service Request is selected, but ASR Manager is not accessible using the specified host address and port number. The Mammoth utility does not install ASR Manager.

The software on Oracle Big Data Appliance must be able to connect to ASR Manager. ASR Manager must be able to route to the Internet, either directly or through a proxy, to send event information that automatically opens service requests.

The following table describes the Auto Service Request fields.

Table 4-23 Auto Service Request

ASR Field	Description
Enable Auto Service Request	Select this option to support Auto Service Request.
ASR Manager Host Name	The fully qualified name or the IP address 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 Root Password	Password for root on the ASR Manager host

4.13.2.11 Enterprise Management Cloud Control

The Mammoth utility deploys and validates agents on Oracle Big Data Appliance that Enterprise Manager uses to monitor the appliance. Mammoth does not install Oracle Enterprise Manager Cloud Control.

Before you can configure Oracle Big Data Appliance for the Enterprise Manager system monitoring plugin, you must install and configure Enterprise Manager to run on a separate server outside of Oracle Big Data Appliance. The Oracle Big Data Appliance software installation fails with an error if you choose the Enterprise Manager option, but Enterprise Manager is not installed and accessible using the specified host address, port numbers, and so forth.

The following table describes the Enterprise Manager Cloud Control fields.

Table 4-24 Enterprise Manager Cloud Control

Cloud Control Field	Description
Enable Oracle Enterprise Manager Cloud Control Agent	Select this option to use the Oracle Enterprise Manager system monitoring plugin.
OMS Host Name	The fully qualified name or the IP address of the server where Oracle Management Server (OMS) is installed with the plugin for Oracle Big Data Appliance.
OMS HTTPS Console Port	The port number for the Oracle Enterprise Manager Cloud Control web interface. To obtain the HTTPS port numbers, use an <code>emctl status oms - details</code> command from the Enterprise Manager host.
OMS HTTPS Upload Port	The HTTP upload port number for Oracle Enterprise Manager Cloud Control web interface.
EM Super Admin User	A Cloud Control user with super-administration privileges to perform administration <code>emcli</code> commands. Typically, this user is <code>sysman</code> .
EM Super Admin Password	Password for the Cloud Control user name.
EM Agent Registration Password	The password for validating the Oracle Management agents on Oracle Big Data Appliance. The Agent Registration password is part of the security setup of Enterprise Manager. To obtain the password in Enterprise Manager, click Setup at the top right of the window, Security , and then Registration Passwords .
Cloud Control SYS Password	The SYS password for the Cloud Control repository.
Inventory Location	The full path of the <code>oraInventory</code> directory for the system where Oracle Enterprise Manager is installed.

4.13.2.12 Email Alerting

Cloudera Manager sends email alerts when it detects a problem in the CDH cluster.

The following table describes the email alert fields.

Table 4-25 Email Alerting Page

Email Alerting Field	Description
SMTP Server	The fully qualified name or the IP address of the existing SMTP server that the company uses on its internal network. Required.
Uses SSL	Select Yes if a Secure Sockets Layer (SSL) connection is required.
SMTP Port	The port number used by the email server
Requires Authentication	Select this option if your SMTP server requires authentication. You can then enter a user name and a password.
SMTP User Name	User name for Cloudera Manager to log in to the SMTP server. This field is hidden when authentication not selected.
SMTP Password	Password for the user name. This field is hidden when authentication is not selected.
Recipient Addresses	The valid email addresses of users who need to get alerts from Cloudera Manager. Enter each email address on a separate line. Required. The field to the right indicates the number of email addresses entered in the dialog box.

4.13.3 Adding to an Existing Cluster

You cannot use the Oracle Big Data Appliance Configuration Utility to add servers to an existing cluster. Use the Mammoth Utility to do this.

See "[Adding Servers to a Cluster](#) (page 10-8)."

4.14 Client Network Config Summary

This page provides an opportunity to view all of the client network IP addresses on the rack and make any final changes before generating the configuration files.

You can also click **Reset all IPs using default values** to void all client IP addresses and re-generate IP address assignments from the default client network IP address pool.

4.15 Complete Page

You have now set all the installation and configuration options. Click **Back** to return to a page and change its settings. The **Back** button does not clear the pages; your settings remain unless you change them.

The text boxes on this page provide a place for you to identify the Hadoop components and any third-party applications that you plan to use on the cluster, and to record any notes that might be useful at a later date. This information is saved in a file named `master.xml`, which you can use to reload these configuration settings into Oracle Big Data Appliance Configuration Generation Utility.

To generate the configuration files, click **Create Files** and click **Yes** in response to the prompt. An operating system window automatically opens in the directory where the files are saved.

5

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](#) (page 5-1)
- [Getting Ready to Install ASR](#) (page 5-2)
- [Installing ASR Manager](#) (page 5-3)
- [Verifying the ASR Manager Installation](#) (page 5-3)
- [Configuring the Trap Destinations on Oracle Big Data Appliance](#) (page 5-4)
- [Verifying ASR Assets](#) (page 5-4)
- [Validating Auto Service Request on Oracle Big Data Appliance](#) (page 5-5)
- [Troubleshooting ASR](#) (page 5-6)

5.1 Understanding Auto Service Request

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

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

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

An email message is sent to both the My Oracle Support email account and the technical contact for 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>

5.2 Getting Ready to Install ASR

Before installing ASR, complete these prerequisites:

1. Create a My Oracle Support account at <http://support.oracle.com>.
2. Ensure that the following are set up correctly:
 - 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 host 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. ASR Manager uses the following ports for normal operations. Ensure that the Oracle Big Data Appliance and ASR Manager server ports are open:
 - All Oracle Big Data Appliance servers must have port 6481 open for HTTP-based asset activation requests that they receive from ASR Manager.
 - The server running ASR Manager must have port 162 open for SNMP-based telemetry messages that the Oracle Big Data Appliance servers send to it.
 - The Oracle `transport.oracle.com` web server has port 443 open for HTTPS-based telemetry messages that ASR Manager sends to it.

See [Figure 5-1](#) (page 5-3) for the port locations.

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

```
java -version
```


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

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

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

Figure 5-1 (page 5-3) shows the network connections between ASR and Oracle Big Data Appliance.

Figure 5-1 Auto Service Request Network Connections



5.3 Installing ASR Manager

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

http://docs.oracle.com/cd/E37710_01/index.htm

Do not set up the assets.

5.4 Verifying the ASR Manager Installation

As root, perform these checks on ASR Manager to ensure that it is installed properly:

- 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
```

5.5 Configuring the Trap Destinations on Oracle Big Data Appliance

ASR is an optional software component of Oracle Big Data Appliance. When configuring the software in Oracle Big Data Appliance Configuration Generation Utility, you must complete the section about ASR. Then during the software installation, the Mammoth utility configures the servers to trap the component fault indicators. See ["Mammoth Installation Steps" \(page 10-29\)](#).

You can reconfigure your appliance at any time to add or drop support for ASR. See ["Changing the Configuration of Optional Software" \(page 10-13\)](#).

5.6 Verifying ASR Assets

To verify the ASR assets:

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

```
asr list_asset -i asset_ip
```

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

```
asr list_asset
```

The following is an example of the output. You should see a host name and an Oracle ILOM name for each of the servers, for a total of 12, 24, or 36 activations. If you see fewer, 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	1143FMM023	Enabled	SUN FIRE X4270 M2 SERVER
203.0.114.45	bda1node16	1143FMM021	Enabled	SUN FIRE X4270 M2 SERVER
203.0.115.140	bda1node16-c	1143FMM021	Enabled	SUN FIRE X4270 M2 SERVER
203.0.114.46	bda1node17	1143FMM022	Enabled	SUN FIRE X4270 M2 SERVER
203.0.115.141	bda1node17-c	1143FMM022	Enabled	SUN FIRE X4270 M2 SERVER

2. If no assets are listed, then verify that ASR is configured on Oracle Big Data Appliance. Log in to the first server and enter this command:

```
# cat /opt/oracle/BDAMammoth/cluster_name-config.json | grep ASR
```

The following output shows that ASR has not been configured.

```
ASR_ENABLED=false
SETUP_ASR_SERVER=false
ASR_HOST=
```

```
ASR_PORT=162
ASR_SERVER_PWD=
```

To configure ASR, see "[Changing the Configuration of Optional Software](#) (page 10-13)."

 Note:

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

5.7 Validating Auto Service Request on Oracle Big Data Appliance

To validate the ASR installation, you can 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 = disable
destination = 0.0.0.0
destination_port = 0
community_or_username = public
snmp_version = 1
testrule = (Cannot show property)
```

Commands:

```
cd
set
show
```

4. Ensure that the properties have valid values, not blank or set to zeros as shown in Step 3 (page 5-5). 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. Log in to any server on Oracle Big Data Appliance.
2. Enter this command to validate the operating system.

```
# ./opt/oracle/bda/compmon/bda_mon_hw_asr.pl -validate_snmp_subscriber -type asr
```

The sample output shows 10.10.10.123:162 as the ASR Manager address:

```
Sending test trap to destination - 10.10.10.123:162
```

3. 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.

If 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](#) (page 5-6)".

If the emails state that there are contract issues, contact your installation coordinator for further assistance.

5.8 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.

Part II

Hardware Installation and Configuration

An Oracle field engineer typically installs and configures the hardware components of Oracle Big Data Appliance. You must follow these instructions only if you plan to install and configure Oracle Big Data Appliance yourself.

Part II contains the following chapters:

- [Installing Oracle Big Data Appliance at the Site](#) (page 6-1)
- [Configuring an Oracle Big Data Appliance Full Rack or Starter Rack](#) (page 7-1)
- [Expanding an Oracle Big Data Appliance Starter Rack](#) (page 8-1)
- [Connecting Multiple Oracle Big Data Appliance Racks](#) (page 9-1)

6

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](#) (page 6-1)
- [Unpacking Oracle Big Data Appliance](#) (page 6-2)
- [Placing Oracle Big Data Appliance in Its Allocated Space](#) (page 6-6)
- [Powering On the System for the First Time](#) (page 6-9)

6.1 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 before 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 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/>.

6.2 Unpacking Oracle Big Data Appliance

You should determine the unpacking location during the site planning process. See [Site Requirements for Oracle Big Data Appliance](#) (page 2-1).

Figure 6-1 (page 6-3) shows the Oracle Big Data Appliance shipping crate.

Figure 6-1 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>

6.2.1 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

6.2.2 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:

Important:

1. When unpacking, verify that the leveling feet are **up** before moving rack.
2. When moving rack to a different location, including repacking, verify that leveling feet are up before moving rack.

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 shipping brackets from the pallet.
 - b. Unfasten the interior shipping brackets from the pallet. Use care when removing the shipping brackets from underneath Oracle Big Data Appliance. Access to the inside shipping brackets might be limited.
 - c. Roll the cabinet down the ramps to the level floor. Oracle recommends having three people available to move the rack down the ramp: one person on each side to help guide the rack and one person in back.
5. Save the shipping brackets that secure the rack to the shipping pallet. You will need them if you later need to relocate the rack.

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.

6.3 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](#) (page 6-6)
- [Securing an Oracle Big Data Appliance Rack](#) (page 6-7)
- [Attaching a Ground Cable \(Optional\)](#) (page 6-8)

6.3.1 Moving Oracle Big Data Appliance



Note:

If you are moving an appliance that has already been installed and in use, see [Steps to Gracefully Shutting Down and Powering On Oracle Big Data Appliance \(Doc ID 1607802.1\)](#) in My Oracle Support (support.oracle.com). If there is no change in network configuration, then a relocation has no impact from a software perspective. Verify that any multi-racking InfiniBand cables are long enough for the new location. Also ensure that all external cables (including power cables) are removed and secured in the rack while wheeling it


To move Oracle Big Data Appliance:

1. Ensure that the doors are closed and secured.
2. **Important:** Check to 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 6-2](#) (page 6-7). You can safely maneuver Oracle Big Data Appliance by carefully pushing it.

[Figure 6-2](#) (page 6-7) shows the correct way to maneuver Oracle Big Data Appliance.

Figure 6-2 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.

⚠ WARNING:

- 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.

6.3.2 Securing an Oracle Big Data Appliance Rack

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

⚠ Caution:

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


6.3.2.1 Secure the Oracle Big Data Appliance Rack with Leveling Feet

The rack contains four leveling feet that you can lower to secure the rack.

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 the figure below, using the SW 12 mm wrench. When lowered correctly, the four leveling feet should support the full weight of the rack.

Figure 6-3 Securing the Oracle Big Data Appliance Rack Using Leveling Feet



6.3.3 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 6-4](#) (page 6-9). The attachment point is an adjustable bolt that is inside the rear of the cabinet.

Figure 6-4 Earth Ground Attachment Bolt Location



6.4 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](#) (page 6-9)
- [Connecting Power Cords](#) (page 6-10)
- [Powering On Oracle Big Data Appliance](#) (page 6-12)

6.4.1 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](#) (page 2-13)" for more information

6.4.2 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 6-5](#) (page 6-11).
4. Route the power cords to the facility receptacles as shown in [Figure 6-6](#) (page 6-11) or [Figure 6-7](#) (page 6-12). 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 6-5](#) (page 6-11) shows the location of the power cord cable ties in a rack.

Figure 6-5 Location of Power Cord Cable Ties




[Figure 6-6](#) (page 6-11) shows the power cords routed out of the bottom of the rack.

Figure 6-6 Power Cord Routing Out the Bottom



[Figure 6-7](#) (page 6-12) shows the power cords routed out the top of the rack.

Figure 6-7 Power Cord Routing Out the Top



6.4.3 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 up 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 6-8](#) (page 6-13). Press the ON (I) side of the toggle switch. When the breaker is ON, the switch is flush with the side of the PDU.

Figure 6-8 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 (Sun Fire X4270 M2-based rack only)
 - KMM switch: Only supplied by PDU B (Sun Fire X4270 M2-based rack only)
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](#) (page 6-1)" 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.

9. If a system does not go into standby mode:
 - a. Connect a laptop to the service processor SET MGT port with baud settings of 115200,8,N,1.
 - b. At the pre-boot prompt, type boot.

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 [Configuring an Oracle Big Data Appliance Full Rack or Starter Rack](#) (page 7-1).

7

Configuring an Oracle Big Data Appliance Full Rack or Starter Rack

This chapter describes how to configure the system, accounts, and software for an Oracle Big Data Appliance full rack or starter rack.

This chapter contains the following sections:

- [Configuring the KVM Switch](#) (page 7-1)
- [Configuring the Cisco Ethernet Switch](#) (page 7-3)
- [Configuring the InfiniBand Leaf and Spine Switches](#) (page 7-12)
- [Configuring the Power Distribution Units](#) (page 7-18)
- [Configuring the Oracle Big Data Appliance Servers](#) (page 7-20)
- [Configuring the Network](#) (page 7-23)
- [Configuring Multiple Oracle Big Data Appliance Racks](#) (page 7-31)

See Also:

- [Expanding an Oracle Big Data Appliance Starter Rack](#) (page 8-1)
- [Executing Commands Across a Cluster Using the dcli Utility](#) (page 14-1)
- [Using Oracle Integrated Lights Out Manager](#) (page 15-1)

Note:

- Oracle recommends that you have the `bda-install-preview.html` file available while completing these procedures. See "[About the Configuration Files](#) (page 4-3)."
- Oracle recommends that an Oracle field engineer performs the configuration steps described in this chapter.

7.1 Configuring the KVM Switch

A KVM switch is provided only in Oracle Big Data Appliance racks with Sun Fire X4270 M2 servers.

The KVM configuration consists of these procedures:

- [Starting the KVM Switch](#) (page 7-2)

- [Connecting the KVM Switch to the Management Network](#) (page 7-2)
- [Configuring the KVM Switch to Access the Servers](#) (page 7-3)
- [Accessing a Server by Using the KVM Switch](#) (page 7-3)

7.1.1 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.

7.1.2 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 restart the KVM switch, choose **Reboot** under Overview and **Yes** to confirm.

7.1.3 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. Start up 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 (page 7-3) through 5 (page 7-3) for each server in the rack. Each server boots through BIOS, and boots the operating system with the default factory IP configuration.

7.1.4 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.

7.2 Configuring the Cisco Ethernet Switch

The Cisco Catalyst 4948 Ethernet switch supplied with Oracle Big Data Appliance has IPBASEK9-MZ firmware. The switch is minimally configured during installation. These procedures configure the Cisco Ethernet switch into one large virtual LAN.

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

- [Scope of the Configuration](#) (page 7-4)
- [Prerequisites for Configuring the Ethernet Switch](#) (page 7-4)
- [Configuring the Ethernet Switch on the Customer Network](#) (page 7-4)
- [Setting Up Telnet Access on the Ethernet Switch](#) (page 7-7)
- [Configuring SSH on the Ethernet Switch](#) (page 7-8)
- [Setting the Clock and Time Zone on the Ethernet Switch](#) (page 7-9)
- [Configuring the NTP Servers for the Ethernet Switch](#) (page 7-10)
- [Disabling the Spanning Tree](#) (page 7-11)

- [Verifying the Ethernet Configuration](#) (page 7-11)

7.2.1 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

7.2.2 Prerequisites for Configuring the Ethernet Switch

To avoid disrupting the customer network, observe these prerequisites:

- Do not connect the Cisco Ethernet switch until the network administrator has verified the running configuration and 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:

- [Understanding the Network Requirements](#) (page 3-1)
- *Oracle Big Data Appliance Site Checklists*

7.2.3 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.
The status lights turn on and this message is displayed:

```
Switch con0 is now available  
Press RETURN to get started.
```
2. If the Ethernet switch displays status lights and a welcome message, then continue. If it does not, then go to "[Correcting the Autoboot Settings](#) (page 7-6)."

3. Ensure that the terminal session is recorded on the laptop by logging the output. You can use the output 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.

4. Change to enable mode using the following command. The default password is welcome1.

```
Switch> enable
Password:
```

5. 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 ]
```

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

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

7. 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 ]
```

8. Set the host name of the switch to the name specified in the installation preview. This example uses the name bda1sw-ip.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# hostname bda1sw-ip
bda1sw-ip(config)# end
*Jan 23 15:57:50.886: %SYS-5-CONFIG_I: Configured from console by console
bda1sw-ip# write memory
Building configuration...
Compressed configuration from 3604 bytes to 1308 bytes[OK]
bda1sw-ip#
```

The system host name appears in the prompt.

9. Configure up to three DNS servers, using the IP addresses specified in the installation preview instead of the examples shown here. If no DNS service is available, then set the domain name anyway, so that the SSH keys are configurable.

```
bda1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bda1sw-ip(config)# ip domain-name mydomain.com
bda1sw-ip(config)# ip name-server 10.7.7.3
bda1sw-ip(config)# ip name-server 129.148.5.5
bda1sw-ip(config)# ip name-server 10.8.160.1
bda1sw-ip(config)# end
*Sep 15 14:26:37.045: %SYS-5-CONFIG_I: Configured from console by console
bda1sw-ip# write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes[OK]
```

Correcting the Autoboot Settings

If the Cisco Ethernet switch does not display any link or status lights, then the incorrect autoboot settings may prevent it from starting. If the serial console port shows the rommon prompt, then complete the following procedure.

Correcting the autoboot settings of the Cisco switch:

1. Verify that the IPBASE and IPBASEK9 firmware reboot files are available:

```
rommon 1 >dir bootflash:
File Size (Bytes) File Name
-----
25213107 cat4500e-ipbase-mz.150-2.SG2.bin
32288280 cat4500e-ipbasek9-mz.151-1.SG.bin
Total space = 128165888 bytes, Available = 60643328 bytes
```

2. Restart the switch manually using the default image:

```
rommon 2 >boot
Rommon reg: 0x00004780
Reset2Reg: 0x00009EFF
##
Fortooin controller 0x0B5A753B..0x0B70B6CD original
size:0x001C3EA5
#####
diagsk5 version 5.0.1
.
.
.
```

3. Press Enter to get the Switch prompt.
4. Enter enable mode. Enter welcome1 for the password or just press Enter:

```
Switch>enable
Password: welcome1
Switch#
```

5. Verify that the boot variable is not set:

```
Switch# show bootvar
BOOT variable =
CONFIG_FILE variable does not exist
```

```
BOOTLDR variable does not exist
Configuration register is 0x2102
```

6. Verify that IPBASEK9 is available in the boot file system:

```
Switch# dir bootflash:
Directory of bootflash:/
 6 -rw- 25213107 Dec 3 2012 17:27:38 +00:00 cat4500eipbase-mz.150-2.SG2.bin
 7 -rw- 32288280 Apr 1 2013 11:41:54 +00:00 cat4500eipbasek9-mz.151-1.SG.bin
128165888 bytes total (60643328 bytes free)
```

7. If IPBASEK9 is not listed in the bootflash directory, then you must install it before continuing. See My Oracle Support Document ID 1415044.1.

8. Set the BOOT variable as shown in this example:

```
Switch# config term
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# config-register 0x2102
Switch(config)# no boot system
Switch(config)# boot system bootflash:cat4500e-ipbasek9-mz.151-1.SG.bin
Switch(config)# end
```

9. Verify the new BOOT setting:

```
Switch# show running-config | include boot
boot-start-marker
boot system bootflash:cat4500e-ipbasek9-mz.151-1.SG.bin
boot-end-marker
```

10. Save the changes to the startup configuration:

```
Switch# copy running-config startup-config all
Destination filename [startup-config]? Enter
Switch# write memory
Building configuration...
Compressed configuration from 3628 bytes to 1281 bytes[OK]
```

11. Reload the firmware.

```
Switch# reload
Proceed with reload? [confirm] y
*Apr 24 09:31:32.887: %SYS-5-RELOAD: Reload requested by console. Reload reason:
Reload command
```

The switch automatically restarts and then continues the reload process.

7.2.4 Setting Up Telnet Access on the Ethernet Switch

Oracle Big Data Appliance ships with a version of the Cisco Ethernet switch software that supports both telnet and SSH. Telnet access is optional. The following procedure describes how to enable and disable remote telnet access.

Note:

Telnet access is not secure.

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.

```
bda1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bda1sw-ip(config)# enable password welcome1
bda1sw-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.
bda1sw-ip(config)# end
bda1sw-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.
bda1sw-ip(config)# line vty 0 15
bda1sw-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
...
bda1sw-ip(config-line)# password welcome1
bda1sw-ip(config-line)# login
bda1sw-ip(config-line)# end
bda1sw-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.
bda1sw-ip(config)# line vty 0 15
bda1sw-ip(config-line)# no password
bda1sw-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
...
bda1sw-ip(config-line)# end
bda1sw-ip# write memory
*Jan 23 15:58:53.630: %SYS-5-CONFIG_I: Configured from console by console
Building configuration...
Compressed configuration from 3786 bytes to 1468 bytes[OK]
```

7.2.5 Configuring SSH on the Ethernet Switch

The current firmware supports remote connections using SSH. To set up SSH on the Cisco Ethernet switch, follow this example:

```
bda1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```

```

bda1sw-ip(config)# crypto key generate rsa
% You already have RSA keys defined named bda1sw-ip.mydomain.com.
% Do you really want to replace them? [yes/no]: yes
Choose the size of the key modulus in the range of 360 to 2048 for
your General Purpose Keys. Choosing a key modulus greater than 512
may take a few minutes.
How many bits in the modulus [512]: 768
% Generating 768 bit RSA keys, keys will be non-exportable...[OK]
bda1sw-ip(config)# username admin password 0 welcome1
bda1sw-ip(config)# line vty 0 15
bda1sw-ip(config-line)# transport input ssh
bda1sw-ip(config-line)# exit
bda1sw-ip(config)# aaa new-model
bda1sw-ip(config)# ip ssh time-out 60
bda1sw-ip(config)# ip ssh authentication-retries 3
Oracle Internal and Approved Partners Only Page 11 of 47 V 1.0 Created: 21. Dec. 2012
Task Comment Check
bda1sw-ip(config)# ip ssh version 2
bda1sw-ip(config)# end
*Sep 15 14:26:37.045: %SYS-5-CONFIG_I: Configured from console by console
bda1sw-ip# write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes[OK]

```

7.2.6 Setting the Clock and Time Zone on the Ethernet Switch

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:

```

bda1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bda1sw-ip(config)# clock timezone EST -5
bda1sw-ip(config)# clock summer-time EDT recurring
bda1sw-ip(config)# end
bda1sw-ip# clock set 15:00:00 January 23 2012
bda1sw-ip# write memory
Building configuration...
Compressed configuration from 3778 bytes to 1433 bytes[OK]
bda1sw-ip# show clock
15:00:18.819 EST Mon Jan 23 2012

```

Clock Commands

The 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. Summer time is disabled by default.

- 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

7.2.7 Configuring the NTP Servers for the Ethernet Switch

Configure up to two NTP servers. Use IP addresses instead of names, which fail when the switch is not connected to the network. If you must use names, then configure the NTP servers after the switch is connected to the network and DNS is resolving names correctly.

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.

```
bda1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bda1sw-ip(config)# ntp server 10.196.23.254 prefer
bda1sw-ip(config)# ntp server 192.168.9.19
bda1sw-ip(config)# end
Jan 23 20:00:41.235: %SYS-5-CONFIG_I:Configured from console by console
bda1sw-ip# write memory
Building configuration...
Compressed configuration from 3870 bytes to 1487 bytes [OK ]
bda1sw-ip# show ntp status
output varies by network
bda1sw-ip# show clock
15:00:57.919 EST Mon Jan 23 2012
```

A symbol may precede the output of the show clock command:

- * (star): The time is not authoritative.
- . (dot) The time is authoritative, but NTP is not synchronized.
- (blank) The time is authoritative, as shown in the example.

7.2.8 Disabling the Spanning Tree

By default, spanning tree is enabled on Cisco switches, but it is disabled in the factory on all the server connection ports 1 to 47.

Adding a switch with spanning tree enabled can cause problems on networks where it is disabled intentionally. As a precaution, disable spanning tree on the uplink port's VLAN before connecting to the customer network.

If spanning tree must be turned on before or after the connection, then the network administrator can specify the appropriate protocol settings, as described in the Cisco switch configuration documentation.

To disable the spanning tree:

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

```
bda1sw-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
bda1sw-ip# write memory
Building configuration...
Compressed configuration from 2654 bytes to 1163 bytes[OK]
```

2. To verify the disabling of the spanning tree:

```
bda1sw-ip# show spanning-tree vlan 1
Spanning tree instance(s) for vlan 1 does not exist.
```

3. To re-enable the default protocol settings, if necessary:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# spanning-tree vlan 1
Switch(config)# end
Switch# write memory
```

7.2.9 Verifying the Ethernet Configuration

To verify the Cisco Ethernet switch configuration:

1. Verify the configuration by entering the following command:

```
bda1sw-ip# show running-config
```

The following is an example of the output:

```
Building configuration...
Current configuration :2654 bytes
!
version 15.0
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
.
```

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

```
bda1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bda1sw-ip# no ip default-gateway 10.7.7.1
bda1sw-ip# end
bda1sw-ip#
*Sep 15 14:13:26.013: %SYS-5-CONFIG_I: Configured from console by console
bda1sw-ip# write memory
```

3. Save the current configuration by entering this command:

```
bda1sw-ip# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
Compressed configuration from 2654 bytes to 1189 bytes[OK]
```

4. Exit from the session with this command:

```
bda1sw-ip#exit
bda1sw-ip con0 is now available
Press RETURN to get started.
```

5. Disconnect the cable from the Cisco console.
6. 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.

7.3 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](#) (page 7-12)
- [Setting the Time Zone and Clock on an InfiniBand Switch](#) (page 7-14)
- [Checking the Health of an InfiniBand Switch](#) (page 7-15)
- [Checking the Sun Network QDR InfiniBand Gateway Switches](#) (page 7-17)

7.3.1 Configuring an InfiniBand Switch

To configure an InfiniBand switch:

1. Connect to the switch using a serial or an Ethernet connection.

- For a remote Ethernet connection to 192.168.1.201, 192.168.1.202, or 192.168.1.203, see "[Connecting to Oracle ILOM Using the Network](#) (page 15-3)".
- For a local serial connection, see "[Connecting to Oracle ILOM Using a Serial Connection](#) (page 15-5)".

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. Check the firmware version, which should be 2.1.8-1 or later:

```
-> version
SP firmware 2.1.8-1
SP firmware build number: 47111
SP firmware date: Thu Feb 19 12:16:50 IST 2015
SP filesystem version: 0.1.22
```

4. Change to the /SP/network directory.

```
cd /SP/network
```

5. 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.

6. 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
```

->

7. 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-iba
```

8. 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.

9. 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
```

7.3.2 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 usntpserver=enabled



Note:

Properly synchronized clocks are required for the Mammoth software installation to succeed. If no NTP server is available on the administrative 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

7.3.3 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 for gateway switches should be 2.0.5-2 or later, and for the spine switch should be 2.1.8-1 or later. Check My Oracle Support ID 1474875.1 for the current version.

```
FabMan@bda1sw-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@bda1sw-02->
```

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

```
FabMan@bda1sw-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@bda1sw-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 = 3.28
Measured 3.3V Standby = 3.39 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@bda1sw-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@bda1sw-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@bda1sw-ib2-> setsmpriority 5
```

- c. Restart the InfiniBand Subnet Manager:

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

7. Exit the Fabric Management shell:

```
FabMan@bda1sw-ib2-> exit
->
```

8. Exit the Oracle ILOM shell:

```
-> exit
```

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

```
# reboot
```

10. Repeat these steps for the other InfiniBand switches.

7.3.4 Checking the Sun Network QDR InfiniBand Gateway Switches

You must check all gateway switch instance numbers (NM2-GW) in the fabric to ensure that the automatic setup assigned unique instance numbers to each switch, both within each rack and among all racks in a multirack environment.

If two switches in the fabric share the same instance number, then the IB-10GbE ports might not work correctly.

The switches must not use consecutive numbers. Oracle recommends even numbers only. From the factory, the Oracle Big Data Appliance gateway switches use the numbers 10 and 20.

If Oracle Big Data Appliance racks are connected to Oracle Exalogic Elastic Cloud Machine racks, then verify that the instance numbers are even and unique to each switch in the fabric.

Note:

All Sun Network QDR InfiniBand Gateway switches in the fabric must run firmware version 2.1.3-4 NM2 or later. Update the firmware as necessary.

To check the gateway switches:

1. Login as user `ilom-admin` on rack 1 gateway switch 1:

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

2. Enter the Fabric Management shell:

```
-> show /SYS/Fabric_Mgmt
NOTE: show on Fabric_Mgmt will launch a restricted Linux shell.
User can execute switch diagnosis, SM Configuration and IB
monitoring commands in the shell. To view the list of commands,
use "help" at rsh prompt.
Use exit command at rsh prompt to revert back to
ILOM shell.
FabMan@hostname->
```

3. Display the current number:

```
FabMan@bda1sw-iba->setgwinstance --list  
No BXM system name set, using 6 last bits of the ip-address, value: 24
```

If you see No BXM system name set (as shown here), or a system name of 0, then you must set the instance number manually.

4. If rack 1 gateway switch 1 is not set to 10:

```
FabMan@bda1sw-iba->setgwinstance 10  
Stopping Bridge Manager... [ OK ]  
Starting Bridge Manager. [ OK ]  
FabMan@bda1sw-iba->setgwinstance --list  
BXM system name set to 10
```

5. Repeat Steps 1 (page 7-17) to 4 (page 7-18) for gateway switch 2. Set it to 20 instead of 10.
6. If multiple racks are cabled together (Oracle Big Data Appliance racks and Oracle Exalogic Elastic Cloud Machine racks), then repeat Steps 1 (page 7-17) to 5 (page 7-18) on each rack. Increment the instance numbers from the previous rack by 2 to get a unique number:
 - Set rack 2 gateway switch 1 to 12 (10+2).
 - Set rack 2 gateway switch 2 to 22 (20+2).

Continue in this way, assigning even numbers between 0 and 63.

7.4 Configuring the Power Distribution Units

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

- [Connecting the PDUs to the Network](#) (page 7-18)
- [Updating the PDU Firmware Version](#) (page 7-19)

7.4.1 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. If the PDUs do not respond, then they probably have different IP addresses:

- a. Press and hold the red RESET button.
- b. When the display starts to count down in about 10 to 15 seconds, release the button and then press it again without holding it down.

- c. When the monitor is restarting, note the current IP setting, and whether it is a dynamic host configuration protocol (DHCP) or static address. Use those addresses to connect.
2. Log in as the **admin** user on the PDU metering unit. The default password is **adm1n** or **admin**. Change this password after configuring the network.
3. To check the firmware version, select **Module Info**. The version should be 1.06. You can upgrade the firmware for a lower version. See "[Updating the PDU Firmware Version](#) (page 7-19)."
4. Click **Net Configuration** in the upper left of the page.
5. Confirm that the DHCP Enabled option is not selected.
6. Enter the following network settings for the PDU metering unit:
 - IP address
 - Subnet mask
 - Default gateway
7. Click **Submit** to set the network settings and reset the PDU metering unit.
8. Repeat Steps [6](#) (page 7-19) and [7](#) (page 7-19) for the second PDU.

7.4.2 Updating the PDU Firmware Version

To update the PDU firmware version:

1. 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 1.0.6.
 - f. Click **Search** to see the Patch Search Results page.
 - g. Click the patch name, such as **16523441**.
 - h. Download the file.
2. Unzip the file on your local system. It has a name such as **p16523441_1060_Generic.zip**. Then unzip the PDU file, which has a name like **OraclePDU_v1.0.6.zip**.
3. In the PDU web interface, select **Net Configuration** and log in as **admin** with password **adm1n**.
4. Scroll down to Firmware Update.
5. Click **Browse**, select the **MKAPP_V1.06.DL** file, and click **Submit**.
6. Click **Browse**, select the **HTML_V1.06.DL** file, and click **Submit**.
7. Click **Module Info** to verify the firmware version number.
8. Click **Net Configuration**, and then click **Logout**.

7.5 Configuring the Oracle Big Data Appliance Servers

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



Note:

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 display the server console, use this command:

```
start HOST/console
```

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. Use a laptop to connect to a server:
 - a. Open an SSH connection using PuTTY or a similar utility.
 - b. Enter the default IP address of the server.
3. Log in as the root user to the first server. The password is welcome1.
4. Set up passwordless SSH for root by entering the `setup-root-ssh` command, as described in "[Setting Up Passwordless SSH](#) (page 14-1)".

For starter racks, with or without expansion, you must use the `-j` parameter. See [setup-root-ssh](#) (page 13-32).

```
# setup-root-ssh -j "eth0_ips[range]"
```

5. Verify that SSH keys are distributed across the rack:

```
# cd /opt/oracle/bda
# dcli "hostname ; date"
```

6. 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 and using the `-j` parameter for a starter rack:

```
# setup-root-ssh -p password [-j "eth0_ips[range]"]
```

Enter the `dcli` command in Step 5 (page 7-20) again to verify the keys.

7. Compare the system serial numbers (S/N) on the front sticker to the server assignments from the bottom (1) to the top (18) of the rack, to ensure that they match:

```
# dcli "dmidecode -s "chassis-serial-number"
192.168.10.1: 1143FMM064
.
.
```

```
192.168.10.18: 1144FMM1A0
```

8. Compare the rack master serial number on the rack front sticker to the rack assignment, to ensure that they match:

```
# dcli "ipmitool sunoem cli 'show /SP system_identifier'" | grep =
192.168.10.1: system_identifier = Oracle Big Data Appliance AK00024695
```

```
192.168.10.18: system_identifier = Oracle Big Data Appliance AK00024695
```

9. If the rack master serial number is incorrect, then reset it in the ILOM on every system. Enter the following command on one line, with no breaks, replacing *R-MSN* with the rack master serial number (such as AK00024695):

```
# dcli -l root \"ipmitool sunoem cli' set /SP system_identifier=\"\"Oracle Big
Data Appliance R-MSN\"\"\" \"> /tmp/set-rack-csn.out
```

In the previous syntax, the `system_identifier` has a 40-character limit.

10. Verify that the InfiniBand ports are up, two on each server (36 total for a full rack of 18 servers).

```
# 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
```

11. Verify that the InfiniBand ports are running at 40 Gbps (4X QDR):

```
# dcli ibstatus | grep rate | more
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)
```

12. 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
```

13. Power up each server and verify that `BDA_IMAGING_SUCCEEDED` and `BDA_REBOOT_SUCCEEDED` are in the `/root` directory of every server:

```
# dcli 'ls /root/BDA*'
```

14. 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
```

15. 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 -vi SUCCESS ~/all-bdahwcheck.out
```

This command returns all lines that do not show SUCCESS, so no returned lines means no failures.

- To verify 32 cores:

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

- To verify 63 GB of memory:

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

- To verify four 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 | more
```

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

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

17. 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
```

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

```
dccli bdchecksw > filename
```

19. 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
```

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

```
# dccli "ubiosconfig export all -x /tmp/bios.xml"
[# dccli "grep -m1 -A20 boot_order /tmp/bios.xml" | more
[root@bda01 bda]# dccli "ubiosconfig export all -x /tmp/bios.xml"
[root@bda01 bda]# dccli "grep -m1 -A20 boot_order /tmp/bios.xml" | more
<boot_order>
  <boot_device>
    <description>USB:USBIN:ORACLE UNIGEN-UFD PMAP</description>
    <instance>1</instance>
  </boot_device>
  <boot_device>
```

```

        <description>RAID:PCIE6:(Bus 50 Dev 00)PCI RAID Adapter</description>
        <instance>1</instance>
    </boot_device>
    <boot_device>
        <description>PXE:NET0:IBA XE Slot 4000 v2196</description>
        <instance>1</instance>
    </boot_device>
    <boot_device>
        <description>PXE:NET1:IBA XE Slot 4001 v2196</description>
        <instance>1</instance>
    </boot_device>
    <boot_device>
        <description>PXE:NET2:IBA XE Slot 8800 v2196</description>
        <instance>1</instance>
    </boot_device>

```

21. If the boot order is incorrect:

- a.** Log in to the ILOM for that server and force the next restart into BIOS setup:

```
# set /HOST boot_device=bios
```

- b.** Restart the server and correct the BIOS.

7.6 Configuring the Network

Oracle Big Data Appliance Configuration Generation Utility generates the cluster-network.json and rack-network.json files, which are used to configure all three network interfaces and all Oracle ILOMs. See "[Generating the Configuration Files](#) (page 4-3)" if you do not have this file.

The network configuration consists of these procedures:

- [Copying the Configuration Files to Oracle Big Data Appliance](#) (page 7-23)
- [Starting the Network Configuration](#) (page 7-24)
- [Connecting to the Administrative Network](#) (page 7-26)
- [Completing the Network Configuration](#) (page 7-27)
- [Checking the Virtual Interface Cards \(VNICs\)](#) (page 7-29)
- [Verifying Access to the Network Services](#) (page 7-30)

7.6.1 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 a laptop or the KVM switch to open a console session to the first server. The first server is the lowest server in the rack. See [Figure E-1](#) (page E-2).
- 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 (page 7-24). 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
network.json
bin
boot
.
.
.
```

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

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

 Note:

If `cluster_name-config.json` is also on the drive, you can copy it to the server for use in [Installing the Oracle Big Data Appliance Software](#) (page 10-1).

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

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

11. Remove the USB flash drive from the server.

7.6.2 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. The procedure varies slightly depending on whether you are configuring a new rack—either a full rack or a starter rack—or configuring additional servers after deploying the starter rack.

To start the network configuration on a new rack:

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

```
# ssh root@192.168.10.1
```

2. For a new rack, verify that /opt/oracle/bda contains the new network.json file.

For the additional servers over the starter rack, verify that /opt/oracle/bda contains both network.json from an existing node and the new networkexpansion.json file.

3. Verify that the InfiniBand network is healthy:

```
# cd /opt/oracle/bda
# bdacheckib -s
```

```
LINK bdasw-ib1.0B ... bdasw-ib3.8B UP
LINK bdasw-ib1.1B ... bdasw-ib2.8B UP
LINK bdasw-ib3.15A ... bda2.HCA-1.2 UP
LINK bdasw-ib3.15B ... bda1.HCA-1.2 UP
LINK bdasw-ib3.14A ... bda4.HCA-1.2 UP
LINK bdasw-ib3.14B ... bda3.HCA-1.2 UP
LINK bdasw-ib3.13A ... bda6.HCA-1.2 UP
.
.
.
```

The InfiniBand network is used extensively during software installation. To perform more in-depth checks of the fabric, run the iblinkinfo utility.

4. Remove passwordless SSH. See ["remove-root-ssh" \(page 13-30\)](#).
5. Reconnect to the server using its ILOM instead of SSH:

```
# ssh root@192.168.1.101
Password: welcome1
-> start -f /HOST/console
bda1 login:
User: root
Password: welcome1
```

The networksetup-one script restarts the network services, so you do not see all of the output using an ssh connection.

6. Begin the network configuration:

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

[Example 7-1](#) (page 7-25) shows sample output from the networksetup-one script.

Example 7-1 Sample Output from networksetup-one

```
# ./networksetup-one
networksetup-one: check syntax and static semantics of /opt/oracle/bda/network.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/network.json to servers
network.json  0% 0  0.0KB/s  --:--  ETAnetwork.json 100% 4304  4.2KB/s  00:00
network.json  0% 0  0.0KB/s  --:--  ETAnetwork.json 100% 4304  4.2KB/s  00:00
.
.
.
networksetup-one: passed
networksetup-one: executing network settings on all servers
networksetup-one: wait a few seconds for the network to restart on 192.168.1.2
.
.
.
bda1node02.example.com BdaUserConfigNetwork: reset network
bda1node03.example.com BdaUserConfigNetwork: reset network
bda1node04.example.com BdaUserConfigNetwork: reset network
.
.
.
networksetup-one: deploying this server
networksetup-one: network will restart momentarily, pardon our dust
bda1node01.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

```

7.6.3 Connecting to the Administrative Network

Before completing the network configuration, you must connect Oracle Big Data Appliance to the 1 GbE administrative (management) network.

To connect Oracle Big Data Appliance to the administrative network:

1. Connect port 48 of the Cisco Ethernet switch to the administrative network.
2. To use the small form-factor pluggable (SFP) ports for a fiber uplink in port 48 (optional), reconfigure the port:

```

bda1sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

bda1sw-ip(config)# interface gigabitEthernet 1/48
bda1sw-ip(config-if)# media-type sfp
bda1sw-ip(config-if)# end
bda1sw-ip#
*Sep 15 14:12:06.309: %SYS-5-CONFIG_I: Configured from console by console
bda1sw-ip# write memory
bda1sw-ip# copy running-config startup-config

```

**See Also:**

For SFP port configuration details, the *Cisco Catalyst 4948 Install Guide* at

<http://www.cisco.com/en/US/docs/switches/lan/catalyst4900/4948/05modcfg.html>

7.6.4 Connecting to the Client Network

You must also connect Oracle Big Data Appliance to the 10 GbE client network.

To connect Oracle Big Data Appliance to the client network:

1. Connect the two Sun Network QDR InfiniBand Gateway leaf switches to the data center.
2. After the cables are routed, the network administrator may need to alter the network switch end configuration to recognize the links.
3. Use SSH to connect to a Sun Network QDR InfiniBand Gateway switch, and log in as the `ilom-admin` user.

4. Enter the Fabric Management shell:

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

5. Ensure that the Bridge entries have active links from the leaf switches to the customer network switch:

```
FabMan@bda1sw-iba-> listlinkup
```

This output shows four ports connected to each switch:

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

6. Repeat Steps 3 (page 7-27) to 5 (page 7-27) on the second gateway switch (-ib3).

7.6.5 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 Administrative Network](#) (page 7-26)".

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

```
./networksetup-two | tee /tmp/network-conf.out
```

The previous command captures the output in a file named `/tmp/network-conf.out`.

[Example 7-2](#) (page 7-28) shows sample output from the script.

Example 7-2 Sample Output from networksetup-two

```
# ./networksetup-two
networksetup-two: check syntax and static semantics of /opt/oracle/bda/network.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 bda1node02.example.com
hello from bda1node03.example.com
.
```

```

.
.
networksetup-two: passed
networksetup-two: check existence of default vlan for port 0A-ETH-1 on bda1sw-ib2
networksetup-two: no default vlan for port, create it
spawn ssh root@10.133.43.36 createvlan 0A-ETH-1 -vlan -1 -pkey default
networksetup-two: verify default vlan for port 0A-ETH-1 for bda1sw-ib2
.
.
.
networksetup-two: passed
networksetup-two: apply eoib on each server
networksetup-two: wait a few seconds for the network to restart on 10.133.42.253
networksetup-two: wait a few seconds for the network to restart on 10.133.42.254
.
.
.
check and delete vNIC for bda1node02 eth9 on switch bda1sw-ib2
check and delete vNIC for bda1node02 eth9 on switch bda1sw-ib3
create vNIC eth9 bda1node02 using switch bda1sw-ib3
vNIC created
check and delete vNIC for bda1node02 eth8 on switch bda1sw-ib2
.
.
.
networksetup-two: ping server ips on client network
networksetup-two: passed
networksetup-two: test ssh server ips on client network
hello from bda1node02.example.com
hello from bda1node03.example.com
.
.
.
networksetup-two: passed
networksetup-two: end setup client network

```

7.6.6 Checking the Virtual Interface Cards (VNICs)

Verify that the VNICs were created correctly on both gateway switches.

To check the VNICs:

1. Use SSH to connect to a Sun Network QDR InfiniBand Gateway switch, and log in as the ilom-admin user.
2. Enter the Fabric Management shell:
-> show /SYS/Fabric_Mgmt
3. Verify that the active ports are assigned to the default VLAN (0):

```

FabMan@bda1sw-ib2-> showvlan
Connector/LAG VLN PKEY
-----
0A-ETH-1      0 ffff
0A-ETH-3      0 ffff
1A-ETH-1      0 ffff
1A-ETH-3      0 ffff

```

4. Verify that the VNICs were created round-robin on each server and 10 GbE interface:

```
FabMan@hostname-> showvnics
ID STATE FLG IOA_GUID NODE
IID MAC VLN PKEY GW
-----
561 UP N 0021280001CF4C23 bda1node12 BDA 192.168.41.31 0000
CE:4C:23:85:2B:0A NO ffff 0A-ETH-1
.
.
.
```

5. Ensure that you can ping in to and out from the 10 GB interfaces.

7.6.7 Verifying Access to the Network Services

Ensure that all network services are running correctly.

To verify access to the network services:

1. Verify that the dcli command is working:

```
# cd /opt/oracle/bda
# dcli "hostname ; date"
```

2. If you are prompted for a password, then enter Ctrl+c several times. Otherwise, generate new SSH keys with the setup-root-ssh utility. See "[setup-root-ssh](#) (page 13-32)."
3. Use the output from the date command to check that the servers are synchronized within a few seconds. If they are not, then restart each device. Time differences can cause the software installation to fail.

You can also check the NTP clients:

```
# ntpq -p 127.0.0.1
```

4. Verify network connectivity:

```
# bdchecknet | tee -a /tmp/bdchecknet_exp.out
bdchecknet: check syntax and static semantics of /opt/oracle/bda/network.json
bdchecknet: passed
bdchecknet: ping test private infiniband ips (bondib0 40gbs)
bdchecknet: passed
bdchecknet: ping test admin ips (eth0 1gbs)
bdchecknet: passed
bdchecknet: ping test client access ips (bondeth0 10gbs Eoib)
bdchecknet: passed
bdchecknet: test admin network resolve and reverse resolve
bdchecknet: passed
bdchecknet: test admin name array matches ip array
bdchecknet: passed
bdchecknet: test client network (eoib) resolve and reverse resolve
bdchecknet: passed
bdchecknet: test client name array matches ip array
bdchecknet: passed
bdchecknet: test ntp servers
bdchecknet: passed
bdchecknet: test arp -a
bdchecknet: passed
```

5. If the network is not able to resolve both DNS names and IP addresses, then correct the /etc/resolv.conf files. Otherwise, the software installation will fail.

6. Check the accessibility of the PDUs by entering the IP address or host name as the address in a browser.

If the connection is successful, you see the Current Measurement page.

7.7 Configuring Multiple Oracle Big Data Appliance Racks

The general process when installing multiple racks is to configure the individual racks, cable the racks together, and validate the multirack network.

7.7.1 Configuring and Cabling Multiple Racks

Complete one of the following procedures. They differ depending on whether you are setting up a new installation or adding racks to an operational installation, and whether there are other engineered systems, such as Oracle Exadata Database Machine.

To install multiple racks:

1. Complete the configuration steps for each rack individually:
 - [Configuring the KVM Switch](#) (page 7-1)
 - [Configuring the Cisco Ethernet Switch](#) (page 7-3)
 - [Configuring the InfiniBand Leaf and Spine Switches](#) (page 7-12)
 - [Configuring the Power Distribution Units](#) (page 7-18)
 - [Configuring the Oracle Big Data Appliance Servers](#) (page 7-20)
 - [Configuring the Network](#) (page 7-23)
2. Cable the Oracle Big Data Appliance racks together. See [Connecting Multiple Oracle Big Data Appliance Racks](#) (page 9-1).
3. Revalidate the InfiniBand network. See "[Validating a Multirack InfiniBand Network](#) (page 7-33)."

To add one or more racks to an operational Oracle Big Data Appliance rack:

1. Complete the configuration steps for each new Oracle Big Data Appliance rack individually:
 - [Configuring the KVM Switch](#) (page 7-1)
 - [Configuring the Cisco Ethernet Switch](#) (page 7-3)
 - [Configuring the InfiniBand Leaf and Spine Switches](#) (page 7-12)
 - [Configuring the Power Distribution Units](#) (page 7-18)
 - [Configuring the Oracle Big Data Appliance Servers](#) (page 7-20)
 - [Configuring the Network](#) (page 7-23)
2. Unplug the power to the InfiniBand switches on the new racks.
3. Cable the Oracle Big Data Appliance racks together. See [Connecting Multiple Oracle Big Data Appliance Racks](#) (page 9-1).
4. Plug in power to the InfiniBand switches. Wait five minutes between switches while the subnet manager fabric is updated.

5. Revalidate the InfiniBand network. See "[Validating a Multirack InfiniBand Network](#) (page 7-33)."

To connect Oracle Big Data Appliance racks to a new Oracle Exadata Database Machine or Oracle Exalogic Elastic Cloud Machine rack:

1. Complete the configuration steps for each Oracle Big Data Appliance rack individually:
 - [Configuring the KVM Switch](#) (page 7-1)
 - [Configuring the Cisco Ethernet Switch](#) (page 7-3)
 - [Configuring the InfiniBand Leaf and Spine Switches](#) (page 7-12)
 - [Configuring the Power Distribution Units](#) (page 7-18)
 - [Configuring the Oracle Big Data Appliance Servers](#) (page 7-20)
 - [Configuring the Network](#) (page 7-23)
2. Cable the Oracle Big Data Appliance racks together. See [Connecting Multiple Oracle Big Data Appliance Racks](#) (page 9-1).
3. Revalidate the InfiniBand network. See "[Validating a Multirack InfiniBand Network](#) (page 7-33)."

To add one or more racks to an operational Oracle Exadata Database Machine or Oracle Exalogic Elastic Cloud Machine rack:

1. Complete the configuration steps for each new Oracle Big Data Appliance rack individually:
 - [Configuring the KVM Switch](#) (page 7-1)
 - [Configuring the Cisco Ethernet Switch](#) (page 7-3)
 - [Configuring the InfiniBand Leaf and Spine Switches](#) (page 7-12)
 - [Configuring the Power Distribution Units](#) (page 7-18)
 - [Configuring the Oracle Big Data Appliance Servers](#) (page 7-20)
 - [Configuring the Network](#) (page 7-23)
2. Unplug the power to the InfiniBand switches on the new racks.
3. Cable the Oracle Big Data Appliance, Oracle Exadata Database Machine, or Oracle Exalogic Elastic Cloud Machine racks together. See [Connecting Multiple Oracle Big Data Appliance Racks](#) (page 9-1).
4. Plug in power to the InfiniBand switches. Wait five minutes between switches while the subnet manager fabric is updated.
5. Ensure that the subnet manager only runs on the switches with the highest firmware versions:

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. On the racks that have lower firmware versions than the Oracle Big Data Appliance InfiniBand switches (both spine and gateway), disable the subnet manager.

- a. Log into the Exadata master switch as root.
- b. Verify that the Oracle Big Data Appliance switches are visible:

```
# ibswitches
```

- c. Disable the subnet manager on Oracle Exadata Database Machine:
disablesm
- d. Log in to each additional lower-version InfiniBand switch and verify that the master switch changes to the Oracle Big Data Appliance spine switch:
getmaster
- e. Disable subnet manager on each lower-version InfiniBand switch:
disablesm
6. Revalidate the InfiniBand network. See "[Validating a Multirack InfiniBand Network](#) (page 7-33)."

7.7.2 Validating a Multirack InfiniBand Network

After cabling the racks together, verify that the InfiniBand network is operating correctly.

Note:

The bdacheckib utility validates connections among Oracle Big Data Appliance and Oracle Exadata Database Machine racks. It does not check other types of racks, such as Oracle Exalogic Elastic Cloud Machine.

To validate the InfiniBand network:

1. Generate a topology file for the network:
bdacheckib -g

The name of the generated file is sample-multi-rack.json
2. Open sample-multi-rack.json in a text editor.
3. Reorder the elements to start with the rack at the left and end with the rack on the right. Save the file.
4. Validate the configuration:

```
# bdacheckib sample-multi-rack.json
Verifying rack #1
leaf: bda1sw-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: bda1sw-ib3
LINK ... to rack2 UP
LINK ... to rack2 UP
LINK ... to rack1 UP
.
.
.
```

[Example 7-3](#) (page 7-34) shows a sample configuration file.

Example 7-3 Multirack Configuration File sample-multi-rack.json

```
# cat sample-multi-rack.json
# This json multirack spec is generated. The array elements are sorted
# alphabetically. A proper arranged json spec representing racks from left to right
# can be used as input mr.json to bdacheckib (bdacheckib mr.json)
# Commas separating rack elements are optional.
[{"SPINE_NAME": "bda1sw-ib1", "LEAF1_NAME": "bda1sw-ib2", "LEAF2_NAME": "bda1sw-ib3"}
{"SPINE_NAME": "bda2sw-ib1", "LEAF1_NAME": "bda2sw-ib2", "LEAF2_NAME": "bda2sw-ib3"}
{"SPINE_NAME": "exa1sw-ib1", "LEAF1_NAME": "exa1sw-ib2", "LEAF2_NAME": "exa01sw-
ib3"}]
```

8

Expanding an Oracle Big Data Appliance Starter Rack

You can install and configure one or more Big Data Appliance X5-2 High Capacity Nodes plus InfiniBand Infrastructure to add up to 12 new servers on a starter rack. This chapter explains the configuration steps for adding new servers on a starter rack.

This chapter consists of the following sections:

- [About Expanding an Oracle Big Data Appliance Starter Rack](#) (page 8-1)
- [Installing New Servers to an Oracle Big Data Appliance Starter Rack](#) (page 8-3)
- [Configuring the New Oracle Big Data Appliance Servers](#) (page 8-4)
- [Setting Up the Network](#) (page 8-5)
- [Installing Oracle Big Data Appliance Software on the New Servers](#) (page 8-7)

Note:

Oracle recommends that an Oracle field engineer performs the procedures described in this chapter.

8.1 About Expanding an Oracle Big Data Appliance Starter Rack

An Oracle Big Data Appliance starter rack has the same hardware configuration as a full rack, except that it comes with 6 servers instead of 18. All switches and power supplies are included in the starter rack, and do not need to be upgraded or supplemented to support additional servers.

You can add up to 12 new servers on a starter rack by using a High Capacity node plus InfiniBand Infrastructure. You can either add all the new servers together or add some of the new servers at a later time on a starter rack.

All Oracle Big Data Appliance expansion nodes come with the in-rack cables required for power as well as connections to the administrative and InfiniBand networks. Adding servers to a rack requires no changes to multi-rack InfiniBand connections, since that is a separate cabling task from switch to switch.

8.1.1 Checking the Contents of a High Capacity Node plus Infiniband Infrastructure

Before starting the installation, ensure that you received all components, cables, and accessories for the High Capacity node plus InfiniBand Infrastructure as shown in the following tables.

[Table 8-1](#) (page 8-2) lists the components and accessories.

Table 8-1 A High Capacity Node plus Infiniband Infrastructure Components and Accessories

Quantity	Description
1	Oracle Big Data Appliance High Capacity Node plus InfiniBand Infrastructure
2	2U snap-in slide rail kit
1	1U-2U slim-rail CMA kit
1	Velcro 25-foot roll
1	Cable labeling kit
1	Server documentation and license packages

[Table 8-2](#) (page 8-2) lists the cables.

Table 8-2 A High Capacity Node plus Infiniband Infrastructure Cables

Length	Quantity	Description
10 feet	1	Red Ethernet Cat 5 or Cat 5e
10 feet	1	Ethernet Black Cat 5 or Cat 5e
3 meters	2	InfiniBand QSFP
2 meters	2	Power

8.1.2 Checking the Status of a Starter Rack

Ensure that the starter rack is in a good working condition before you install the High Capacity node plus InfiniBand Infrastructure:

- Check the LEDs on all components to verify a normal operating state. Resolve any issues before proceeding.
- Verify that the Sun Rack II PDU firmware is at or above the minimum version of 1.06. See "[Updating the PDU Firmware Version](#) (page 7-19)."

8.1.3 Expanding an Oracle Big Data Appliance X4-2 Starter Rack

You can also add one or more Oracle Big Data Appliance servers to an Oracle Big Data Appliance X4-2 starter rack by using a High Capacity node plus InfiniBand Infrastructure.

To prepare the starter rack for expanding the hardware:

- Upgrade the starter rack to Oracle Big Data Appliance 4.1 or later.



Note:

Starter rack hardware for the Oracle Big Data Appliance versions earlier to 4.1 cannot be expanded using a High Capacity node plus InfiniBand Infrastructure.

8.2 Installing New Servers to an Oracle Big Data Appliance Starter Rack

The following instructions describe how to physically install new servers on a starter rack.



Caution:

Two people are needed to install new servers on a starter rack.

To install new servers on a starter rack:

1. Unpack all power cables, Ethernet cables, and InfiniBand cables, and mark them with the labels provided in the kit:
 - Mark each end of a cable with the appropriate label. Take care to put the "from" and "to" labels on the correct end.
 - Select the appropriate PDU power cord labels for your installation, either single-phase or three-phase.
2. Plug in the cables at both ends. Refer to the cabling tables in the appendixes.
3. Install the server rack mount kits in the rack. Start at the first empty unit at the bottom and work up toward the top of the rack.
4. Install the servers in the rack rails. After installing each server, ensure that the server and the unit below it can slide in and out normally.
5. Install the cable arms and cables into each server, starting with the first new server at the bottom and working up to the top unit. Ensure that all InfiniBand cables are plugged in to the correct ports on both ends. Keep the power cables out of the way while you route the InfiniBand and Ethernet cables.



Caution:

Do not plug the power cables into the PDU strips until you are ready to turn on the PDU and configure it. Plug the power cables only into the servers now.

8.3 Configuring the New Oracle Big Data Appliance Servers

After installing the new servers in the rack, you can configure them as described in the following steps.

To configure the new servers:

1. Disconnect the Cisco switch uplink cable (port 48) from the network, in case the factory IP addresses duplicate those on the local network.
2. Plug the power cords into the appropriate outlets. Route the cables to PDU-A last. When you are done, Oracle ILOM starts, the service processor LED is solid green, and the system is in Standby mode with a flashing OK LED.
3. Press the power button on the servers to turn them on. The system automatically detects the factory IP addresses. See "[Factory Network Settings](#) (page 3-9)."
4. Using a laptop connected to the first new server, log in as root using SSH. The default password is welcome1.
5. Set up passwordless SSH for the root user. For *range*, for example, enter [7:10] when you install four new servers.

```
# setup-root-ssh -j "eth0_ips[range]" -p welcome1
```

See "[setup-root-ssh](#) (page 13-32)."

6. Confirm that passwordless SSH works properly across the range of new servers. Use the same value for *range* as you did in step 5 (page 8-4).
7. Compare the output of the following command with the system serial number (S/N) sticker on the front of each system.

```
# dcli -j "eth0_ips[range]" "dmidecode -s chassis-serial-number"
192.168.1.7: 1143FMM064
.
.
.
192.168.1.10: 1144FMM1A0
```

If you find inconsistencies, then verify that the InfiniBand cables are plugged into the correct ports on the gateway switches for each node.

8. Set the rack master serial number in the ILOM to the S/N on the front left rack sticker (also displayed on the rear of the rack near PDU-A):

```
# dcli -j "eth0_ips[range]" "ipmitool sunoem cli 'set /SP
\system_identifier=\""Oracle Big Data Appliance rack-master-serial-number\""'
```

This example sets the range to [7:10] for the four new servers, and sets the serial number to AK00086253:

```
# dcli -j "eth0_ips[7:10]" "ipmitool sunoem cli 'set /SP
\system_identifier=\""Oracle Big Data Appliance AK00086253\""'
```

The system identifier has a 40-character limit.

9. Verify that the rack master serial number is set correctly. This example uses the range for the four new servers:

```
# dcli -j "eth0_ips[7:10]" "ipmitool sunoem cli 'show /SP system_identifier'" |
grep =
192.168.1.7: system_identifier = Oracle Big Data Appliance AK00086253
.
.
.
192.168.1.10: system_identifier = Oracle Big Data Appliance AK00086253
```

10. Complete the physical hardware checks beginning with Step 10 (page 7-21) in ["Configuring the Oracle Big Data Appliance Servers"](#) (page 7-20). In all `dcli` commands, be sure to enter the range, using the `eth0_ips[range]` syntax shown previously.

8.4 Setting Up the Network

The following topics explain how to add the new servers to the Ethernet administrative network:

- [Copying the Configuration Files](#) (page 8-5)
- [Configuring the Network](#) (page 8-6)
- [Verifying the Network Configuration](#) (page 8-7)



See Also:

["Configuring the Network"](#) (page 7-23)

8.4.1 Copying the Configuration Files

These instructions use a serial connection to Oracle ILOM. If you use an SSH connection instead, then you lose the console connection when the IP addresses change from the factory settings to the custom settings in ["Configuring the Network"](#) (page 8-6).

Follow these steps to copy the configuration files to the first new server in the rack.

To copy the configuration files:

1. Plug your laptop into the Oracle ILOM "SER MGT" port of the first new server.
2. Log in as root with the default `welcome1` password:

```
login as: root
Using keyboard-interactive authentication.
Password: welcome1
Oracle(R) Integrated Lights Out Manager
Version 3.1.2.12 r74388
Copyright (c) 2012, Oracle and/or its affiliates. All rights reserved.
->
```

3. Open the console and log in with the same credentials:

```
-> start /SP/console
Are you sure you want to start /SP/console (y/n)? y
Serial console started. To stop, type ESC (
bda7 login:
```

```
User: root
Password: welcome1
#
```

4. If you do not see a prompt, press **Enter**.
5. Copy the original Oracle Big Data Appliance configuration file from the first server in the rack (server 1) to the new server where you are logged in. Replace *IP_address* with the real IP address of server 1:

```
# scp root@IP_address:/opt/oracle/bda/network.json /opt/oracle/bda
```

Alternatively, use a USB flash drive to copy the file between servers.

6. Copy the new networkexpansion.json configuration file from a USB flash drive to the same /opt/oracle/dba directory:

```
# mkdir /mnt/usb
# /opt/oracle/bda/bin/showusb /dev/sdm1
# mount /dev/sdm1 /mnt/usb
# cp /mnt/usb/networkexpansion.json /opt/oracle/bda
```

See "[Copying the Configuration Files to Oracle Big Data Appliance](#) (page 7-23)" for general instructions on using a flash drive.

8.4.2 Configuring the Network

Complete the following procedure to configure the network.

1. Run the networksetup-one script:

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

After this step completes, the servers have their custom IP addresses.

2. Reconnect the Cisco switch uplink cable to the local network.
3. If you lost your SSH connection when the IP addresses changed, then reconnect:
 - a. Use SSH to connect to Oracle ILOM of the first server, using the custom IP address.
 - b. Log in as the root user with the default welcome1 password.

If you use a serial connection, you may still be connected when networksetup-one completes. If not, type Esc+(to get the Oracle ILOM prompt before reconnecting.

4. Verify the InfiniBand topology by running bdacheckib -m on the rack configuration file. For example, if the rack name is rack01:

```
# cd /opt/oracle/bda
# bdacheckib -m rack01-network.json
```

5. Verify that passwordless SSH still works:

```
# cd /opt/oracle/bda
# dcli -j "eth0_ips[range]" "hostname ; date"
```

If you are prompted for a password, type Ctrl+c several times. If the dcli command fails, then repeat Steps 5 (page 8-4) and 6 (page 8-4) of the first procedure to configure the servers.

6. Use the output from the date utility in Step 5 (page 8-6) to verify that the clocks for all servers and switches are synchronized. The times should be within a few seconds.

If they are not synchronized, then you can manually change the time, restart the device, or check the routing through the Cisco switch to the NTP server. Issue this command to check the status of an NTP client:

```
# ntpq -p 127.0.0.1
```

7. Run networksetup-two:

```
# ./networksetup-two | tee /tmp/bdanet_exp.out
```

The tee utility captures the output in a file while still allowing it to appear on your screen.

8. Log out of the server, unless it is the first server in the cluster.

8.4.3 Verifying the Network Configuration

Complete the following procedure to ensure that the network is working properly

1. Log in to the first server in the cluster.
2. Verify that the network and all network services are working properly:

```
# bdachecknet | tee -a /tmp/bdanet_exp.out
bdachecknet: check syntax and static semantics of opt/oracle/bda/network.json
bdachecknet: passed
bdachecknet: ping test private infiniband ips (bondib0 40gbs)
bdachecknet: passed
bdachecknet: ping test admin ips (eth0 1gbs)
bdachecknet: passed
.
.
.
```

3. If the systems are not able to resolve and use the DNS names, then manually correct the /etc/resolv.conf files and verify the reverse lookup.
4. Submit an install base update service request in My Oracle Support to update your assets.

8.5 Installing Oracle Big Data Appliance Software on the New Servers

You can add the new servers to an existing cluster, or configure them as a new cluster.

To add the servers to an existing cluster:

- Complete the instructions in ["Adding Servers to a Cluster"](#) (page 10-8).

To configure the servers as a separate cluster:

1. Run Oracle Big Data Appliance Configuration Generation Utility, and generate the configuration files. See ["Generating the Configuration Files"](#) (page 4-3).

2. Run the Mammoth utility to install the software. See "[Installing the Software on a New Rack](#) (page 10-6)."

9

Connecting Multiple Oracle Big Data Appliance Racks

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

- [Extending a Rack by Adding Another Rack](#) (page 9-1)
- [Prerequisites for Multirack Cabling](#) (page 9-2)
- [Cabling Two Racks Together](#) (page 9-2)
- [Cabling Several Racks Together](#) (page 9-3)



See Also:

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

9.1 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 before 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:

- [Multirack Cabling Tables](#) (page F-1) for the physical connections among the racks
- *Sun Datacenter InfiniBand Switch 36 Administration Guide* for monitoring the Subnet Manager at
http://docs.oracle.com/cd/E26698_01/html/E26434/z40001039112.html#scrolltoc

9.2 Prerequisites for Multirack Cabling

Verify that all Sun Network QDR InfiniBand Gateway switches on all connected Oracle Big Data Appliance and Oracle Exadata Database Machine racks are running firmware version 2.1.3-4 NM2 or later. Update the firmware as necessary.

When updating the firmware of Sun Network QDR InfiniBand Gateway switches from v2.0.x to v2.1.x or higher, be careful to update them in descending order of the gateway instance numbers. Start with the switch with the highest gateway instance number.

From the Oracle ILOM interface, use this hardware command to obtain the gateway instance number:

```
setgwinstance --list
```

See My Oracle Support Doc ID 1534255.1 about upgrading the firmware for Oracle Big Data Appliance.

9.3 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 1.b (page 9-2) 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 F-2](#) (page F-6).
 5. Shut down 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 F-1](#) (page F-5).
 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. Shut down 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 F-1](#) (page F-5).
 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 `disableism` command to stop the Subnet Manager.
 - b. Use the `setsmpriority 8` command to set the priority to 8.
 - c. Use the `enableism` command to restart the Subnet Manager.
 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.

9.4 Cabling Several Racks Together

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

You can cable up to 18 racks together without additional switches.

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 1.b (page 9-2) to ensure that the Subnet Manager Master is running on the spine switch.
2. Ensure that the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
3. Completely shut down the new rack (R_{n+1}).
4. Cable the leaf switch in the new rack according to the appropriate table in [Multirack Cabling Tables](#) (page F-1). For example, if rack R_{n+1} was R4, then use [Table F-9](#) (page F-10).
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. Shut down 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 [Multirack Cabling Tables](#) (page F-1).
 - 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 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. Shut down 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 [Multirack Cabling Tables](#) (page F-1).
 - 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 [Multirack Cabling Tables](#) (page F-1).
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,... Rn +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 Rn+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.

Part III

Software Installation and Configuration

An Oracle service representative typically configures the network configuration, and installs and configures the software.

Part III contains the following chapter:

- [Installing the Oracle Big Data Appliance Software](#) (page 10-1)

10

Installing the Oracle Big Data Appliance Software

This chapter explains how to install, reinstall, and reconfigure the software on Oracle Big Data Appliance. It contains these sections:

- [About the Mammoth Utility](#) (page 10-1)
- [Installation Prerequisites](#) (page 10-3)
- [Downloading the Mammoth Software Deployment Bundle](#) (page 10-4)
- [Installing the Software on a New Rack](#) (page 10-6)
- [Adding Servers to a Cluster](#) (page 10-8)
- [What If an Error Occurs During the Installation?](#) (page 10-8)
- [Upgrading the Software on Oracle Big Data Appliance](#) (page 10-10)
- [Changing the Configuration of Optional Software](#) (page 10-13)
- [Reinstalling the Base Image](#) (page 10-19)
- [Installing a One-Off Patch](#) (page 10-26)
- [Mammoth Software Installation and Configuration Utility](#) (page 10-27)
- [Oracle Big Data Appliance Base Imaging Utilities](#) (page 10-32)

Note:

If you did not enter the required passwords in the Oracle Big Data Appliance Configuration Generation Utility, then you are prompted to enter them during the software installation. Ensure that you know the current passwords for the operating system root and oracle users, the Cloudera Manager admin user, and the MySQL administrator. If you are installing or reinstalling Oracle Big Data Connectors, then you also need the MySQL password for Oracle Data Integrator.

10.1 About the Mammoth Utility

Mammoth is a command-line utility for installing and configuring the Oracle Big Data Appliance software. Using Mammoth, you can:

- Set up a cluster for either CDH or Oracle NoSQL Database.
- Create a cluster on one or more racks.
- Create multiple clusters on an Oracle Big Data Appliance rack.
- Extend a cluster to new servers on the same rack or a new rack.

- Update a cluster with new software.

10.1.1 About Rolling Upgrades

A rolling upgrade reboots the nodes of a cluster in a sequence that always leaves one node out of three replicas active. This allows the cluster to stay on-line throughout the entire series of reboots. As a result, some (though not all) services of the cluster can remain accessible without interruption during the upgrade.

Rolling upgrades are only available for installations running CDH 4.1 or later.

Because sequential reboots of the nodes in pairs incrementally add time to the process, rolling upgrades take longer than a conventional Mammoth upgrade, in which all servers reboots are concurrent. Estimate an average of 8-10 minutes per pair of non-critical nodes, in addition to 4 minutes each for the Mammoth node and critical nodes (which are rebooted one at a time).

A rolling upgrade can be unattended. The reboot event will “roll” from one server to the next (or next pair) automatically if each server's post-upgrade reboot is successful. If the update detects error or warning conditions during a server reboot, then the process pauses and prompts you to continue.

```
Reboot of node mnode22bda08 had errors/warnings, do you wish to proceed anyway?  
[y/n]:
```

- If you enter n:

The upgrade stops so that you can fix the problems.

When you are ready, resume the installation using `mammoth -p`.

- if you enter y:

The upgrade continues.

This prompt is displayed once per node upgrade. If the node has already been rebooted (and the prompt has already been answered) then the upgrade will continue without displaying the prompt again.

When Rolling Upgrades are Possible

A rolling upgrade is an option for the following upgrade tasks described in this document:

- [Upgrading the Software on Oracle Big Data Appliance](#) (page 10-10)
- [Installing a One-Off Patch](#) (page 10-26)

For [Adding Servers to a Cluster](#) (page 10-8), there is normally no cluster downtime for the process. The only exception is that the first time a single-rack cluster is extended onto a second rack, there will be downtime for the cluster extension (as master roles are moved to the additional rack), but subsequent cluster extensions will not require any downtime.

In cases where rolling upgrade is an option, Mammoth displays the following message on the console and then prompts you to chose whether or not to proceed with the installation as a rolling upgrade:

You can choose a rolling upgrade or continue with a conventional upgrade. Rolling upgrades need a minimum replication factor of three (this is the default in Mammoth)

to have full availability during the upgrade. Do you wish to perform a Rolling Upgrade? [y/n]:

 Note:

If you enter y at the prompt to accept a rolling upgrade, the cluster is committed to the process and the upgrade of all nodes in the cluster must continue to completion uninterrupted.

Service Availability During a Rolling Upgrade

In a rolling upgrade, services that are configured to use HDFS High Availability (HDFS HA) can remain continuously available while each node in the cluster is upgraded and rebooted. These include HDFS, YARN, and Zookeeper. Keeping these services available during the upgrade requires a Hadoop replication factor of three or greater. This is because the rolling upgrade restarts nodes in pairs. The upgrade cannot predetermine if the replication factor meets this minimum, since global configuration may have been overwritten for individual files. An insufficient replication factor does not stop the upgrade, but in that case HDFS HA-enabled services may be subject to downtime.

In a rolling upgrade, Cloudera services are restarted in batches of two. This is what sustains access to HDFS and other services during an upgrade. Services that are not or cannot be configured to use HDFS HA may be unavailable during a rolling upgrade, regardless of the replication factor.

10.2 Installation Prerequisites

The Oracle Audit Vault, Auto Service Request, and Oracle Enterprise Manager options require software on a remote server on the same network as Oracle Big Data Appliance. Those installations must be complete before you run Mammoth, or it will fail.

Similarly, you must complete several steps for Kerberos if the key distribution center (KDC) is installed on a remote server. There are no preliminary steps if you install the KDC on Oracle Big Data Appliance.

The following list describes the prerequisites for all installation options.

Audit Vault Requirements:

- Oracle Audit Vault and Database Firewall Server Release 12.1.1 must be up and running on a separate server on the same network as Oracle Big Data Appliance. No other version of Audit Value and Database Firewall is supported by Oracle Big Data Appliance.

Auto Service Request Requirements:

- Your My Oracle Support account is set up.
- ASR Manager is up and running.

Enterprise Manager Requirements:

- Oracle Management System (OMS) version 12.1.0.4.0 or higher is up and running.

- The OMS agent pull URL is working.
- The OMS emcli download URL is working.
- Both the HTTPS upload port and the console port are open.



Note:

Double-check the OMS credentials and ensure that you enter them correctly when running Mammoth. Invalid credentials are the primary cause of failure in the Enterprise Manager discovery process.

MIT Kerberos Requirements for a Remote KDC:

1. Add cloudera-scm/admin as a user to the KDC database by running the following command from kadmin:

```
addprinc -randkey cloudera-scm/admin@<REALM NAME>
```
2. Grant cloudera-scm/admin all permissions to the Kerberos database. It must be able to add, modify, remove, and list principals from the database.
3. Create the cmf.keytab file by running the following command from kadmin:

```
xst -k cmf.keytab cloudera-scm/admin@<REALM NAME>
```
4. Move cmf.keytab to /opt/oracle/BDAMammoth.
5. To support Oozie and Hue, ensure that the remote KDC supports renewable tickets. If it does not, then follow these steps:
 - a. Open kdc.conf and set values for max_life and max_renewable_life. The max_life parameter defines the time period when the ticket is valid, and the max_renewable_life parameter defines the time period when users can renew a ticket.
 - b. Set maxrenewlife for the krbtgt principal. Use the following kadmin command, replacing *duration* with the time period and *REALM NAME* with the name of the realm:

```
modprinc -maxrenewlife duration krbtgt/REALM NAME
```

If the KDC does not support renewable tickets when Kerberos is configured, then Oozie and Hue might not work correctly.

Active Directory Kerberos Requirements for a Remote KDC

Active Directory Kerberos requirements are documented in MOS (My Oracle Support).

See MOS documents 2013585.1 and 2029378.1.

10.3 Downloading the Mammoth Software Deployment Bundle

The Mammoth bundle contains the installation files and the base image. Before you install the software, you must use Oracle Big Data Appliance Configuration Generation

Utility to generate the configuration files, as described in "[Generating the Configuration Files](#) (page 4-3)."

You use the same bundle for all procedures described in this chapter, regardless of rack size, and whether you are creating CDH or Oracle NoSQL Database clusters, or upgrading existing clusters.

To download the Mammoth bundle:

1. Locate the download site in either My Oracle Support or Automated Release Updates (ARU):

My Oracle Support

- a. Log on to [My Oracle Support](#) and search for *Oracle Big Data Appliance Patch Set Master Note (Doc ID 1485745.1)*.
- b. In the list provided, click the link to the Mammoth Bundle software version for this release.
- c. Find the link to the appropriate Mammoth patch bundle in the table.

ARU

- a. Connect to ARU.
 - b. On the Patches page, set Product to Big Data Appliance Integrated Software and Release to the appropriate release number.
 - c. Click **Search**.
2. Download the BDA Mammoth ZIP files to any directory (such as /tmp) in the first node of the cluster. Depending on the configuration of the rack, this node can be the first, seventh, or thirteenth server from the bottom of the rack. For multirack clusters, this server is in the primary rack.

The patch consists of a set of files labeled as follows: `p<patch_number>_Linux-x86-64_<file number>of<total number of files>.zip`

3. Log in to the first node of the cluster as root.
4. Extract all files from the all of the downloaded zip files. For example:

```
$ unzip p<patch_number>_Linux-x86-64_1of4.zip
Archive:  p<patch_number>_<version>_Linux-x86-64_1of4.zip
  inflating: README.txt
   creating: BDAMammoth-ol6-<version>/
...
```

5. Change to the BDAMammoth-<version> directory:

```
# cd BDAMammoth-ol6-<version>
```

6. Extract all files from BDAMammoth-<version>.run:

```
# ./BDAMammoth-ol6-<version>.run
```

```
Big Data Appliance Mammoth v<version> Self-extraction
```

```
Checking for and testing BDA Base Image in /tmp
```

```
BDABaseImage-<version>_RELEASE.iso: OK
```

```
Removing existing temporary files
```

```
Generating /tmp/BDAMammoth.tar
```

```
Verifying MD5 sum of /tmp/BDAMammoth.tar
/tmp/BDAMammoth.tar MD5 checksum matches

Extracting /tmp/BDAMammoth.tar to /opt/oracle/BDAMammoth

Extracting Base Image RPMs to bdarepo
Moving BDABaseImage into /opt/oracle/

Extracting BDA Extras

Removing temporary files
.
.
.
Please "cd /opt/oracle/BDAMammoth" before running "./mammoth -i <rack_name>"
#
```

The new version of the Mammoth software is installed in /opt/oracle/BDAMammoth, and the previous version (if you are upgrading) is saved in /opt/oracle/BDAMammoth/previous-BDAMammoth.

7. Follow the specific instructions for the type of installation you are performing.

10.4 Installing the Software on a New Rack

Mammoth installs and configures the software on Oracle Big Data Appliance using the files generated by Oracle Big Data Appliance Configuration Generation Utility. A cluster can be dedicated to either CDH (Hadoop) or Oracle NoSQL Database.

For a CDH cluster, 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. If you have a license, Mammoth optionally installs and configures all components of Oracle Big Data Connectors.

For a NoSQL cluster, Mammoth installs Oracle NoSQL Database. CDH and Oracle NoSQL Database do not share a cluster, beginning with Oracle Big Data Appliance 2.2.

In addition to installing the software across all servers in the rack, Mammoth 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.

10.4.1 Installing the Software

Follow this procedure to install and configure the software on one or more Oracle Big Data Appliance racks. You can configure one cluster on multiple racks in a single installation.

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/network.json. If the rack is still configured to the factory default IP addresses, first perform the network configuration steps described in ["Configuring the Network" \(page 7-23\)](#).

2. Verify that the software is not installed on the rack already. If the software is installed and you want to upgrade it, then use the `mammoth -p` option in Step 6 (page 10-7).
3. Download and unzip the Mammoth bundle, as described in "Downloading the Mammoth Software Deployment Bundle (page 10-4)." You must be logged in as root to the first server in the cluster.

4. Change to the BDAMammoth directory.

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

5. Copy `cluster_name-config.json` to the current directory. See "About the Configuration Files (page 4-3)."
6. Run the `mammoth` command with the appropriate options. See "Mammoth Software Installation and Configuration Utility (page 10-27)." This sample command runs all steps:

```
./mammoth -i rack_name
```

After Mammoth completes Step 3 of the installation, it prompts you to restart, if it upgraded the base image.

7. If you installed support for Auto Service Request, then complete the steps in "Verifying the Auto Service Request Installation (page 10-7)."
8. To configure another CDH cluster on the server:
 - a. Copy the BDAMammoth ZIP file to any directory on the first server of the cluster, which is either server 7 or server 13.
 - b. Repeat Steps 3 (page 10-7) to 7 (page 10-7). Each cluster has its own `cluster_name-config.json` file. Oracle Big Data Appliance Configuration Generation Utility creates the files in separate directories named for the clusters.

 Note:

Mammoth stores the current configuration in the `/opt/oracle/bda/install/` state directory. *Do not delete the files in this directory.* Mammoth fails without this information if you must use it again, such as adding a rack to the cluster.

Verifying the Auto Service Request Installation

1. Log in to My Oracle Support at <http://support.oracle.com>.
2. Search for document ID 2103715.1, *Engineered Systems ASR Configuration Check tool (asrexacheck version 4.x)*, and download the `asrexacheck` script.

Although this check was originally intended for Oracle Exadata Database Machine, it is now supported on Oracle Big Data Appliance.
3. Copy `asrexacheck` to a server in the Oracle Big Data Appliance cluster.
4. Log in to the server as root.
5. Copy `asrexacheck` to all servers in the cluster:

```
# dcli -C -f asrexacheck -d /opt/oracle.SupportTools
```

See [Executing Commands Across a Cluster Using the dcli Utility](#) (page 14-1) for information about dcli.

6. Change the file access permissions:

```
# dcli -C chmod 755 /opt/oracle.SupportTools/asrexacheck
```
7. Run the script:

```
# dcli -C /opt/oracle.SupportTools/asrexacheck
```
8. File an Oracle Support Request (SR) to validate the ASR installation. Note the following choices:
 - a. Under "What is the Problem?" click the Hardware tab.
 - b. For Products Grouped By, select Hardware License.
 - c. For Problem Type, select My - Auto Service Request (ASR) Installation and Configuration Issues.
 - d. Include the output of asrexacheck.
9. Continue with Step 8 (page 10-7) of the software installation procedure.

10.5 What If an Error Occurs During the Installation?

If the Mammoth utility fails, take these steps to resolve the problem:

1. Read the error message to see if it suggests a cause or a resolution.
2. Make the recommended changes and rerun the step.
3. If the error message does not recommend a resolution, or it does not resolve the problem:
 - File a service request (SR) with My Oracle Support.
 - Upload the diagnostic zip file, which Mammoth generates when an error occurs.



See Also:

Oracle Big Data Appliance Software User's Guide

10.6 Adding Servers to a Cluster

To add servers to an existing cluster, use the Mammoth Utility to extend the cluster onto the new servers. Do not use the Oracle Big Data Appliance Configuration Generation Utility for this purpose.

For each rack in a cluster, the minimum number of servers that the rack can contribute to the cluster is three. Therefore, initially you must add a group of at least three. After the minimum number is fulfilled, there are no further constraints. You can then add more servers from the rack to the cluster individually or in groups of any number. You can also extend the cluster to all of the servers on the rack at the same time.

If you are adding servers to same rack (and to the same cluster) then there is no cluster downtime for the process. The first time a single-rack cluster is extended onto a

second rack, there will be downtime for the cluster extension (as master roles are moved to the additional rack) but subsequent cluster extensions will not require any downtime. (See [About Rolling Upgrades](#) (page 10-2).)

It is not necessary to decommission edge nodes when adding new nodes to an existing cluster.

There is no significant difference between how new server nodes and old nodes are configured. However, when a new node has a larger disk capacity, the data partitions created are correspondingly larger in order to use the entire disk. HDFS handles the fact that some data nodes have more space available than other nodes. Oracle Big Data Appliance performs no additional special configuration.

To install the software on additional servers in a cluster:

1. Ensure that all servers are running the same software version. The additional servers must not have an Oracle Big Data Appliance base image that is newer than the existing cluster. See "[About Software Version Differences](#) (page 10-9)".
2. Ensure that all racks that form a single Hadoop cluster are cabled together. See [Connecting Multiple Oracle Big Data Appliance Racks](#) (page 9-1).

3. 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.

4. Generate a parameter file for the server group. The following example adds six servers beginning with node13:

```
./mammoth -e node13 node14 node15 node16 node17 node18
```

The servers can be on the same rack or multiple racks. See the `-e` option in "[Mammoth Software Installation and Configuration Utility](#) (page 10-27)."

After Mammoth completes Step 3 of the installation, it prompts you to restart, if it upgraded the base image.

5. If you are using Oracle Enterprise Manager Cloud Control to monitor Oracle Big Data Appliance, then run rediscovery to identify the hardware and software changes.

If you have a license for Oracle Big Data Connectors, then they are installed on all nodes of the non-primary racks, although the services do not run on them. Oracle Data Integrator agent still runs on node03 of the primary rack.

Mammoth 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 Mammoth fails.

About Software Version Differences

All servers configured as one Hadoop cluster must have the same image. A new Oracle Big Data Appliance rack might 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. When all servers of a single Hadoop cluster have the same image version, you can install the software.

To synchronize the new servers 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 servers.

To upgrade the image version:

- Run Mammoth with the `-p` option to upgrade the software on the cluster to the latest version. See "[Upgrading the Software on Oracle Big Data Appliance](#) (page 10-10)."

To downgrade the image version:

- Reimage the new rack to the older version installed on the cluster. See My Oracle Support Information Center ID 1445745.2.
- Use the old version of the Mammoth utility, which is on the first server of the existing cluster, to extend the cluster onto the new rack.

If you add a newer server model, then you can downgrade only to the first software version available for that model. For example, if you add Sun Server X3-2L servers, then you must install Oracle Big Data Appliance software version 2.3 or higher.

10.7 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, Oracle Linux Unbreakable Enterprise Kernel (UEK), CDH, JDK, and Oracle Big Data Connectors (if previously installed). Two versions of the software bundle are available, one for Oracle Linux 5 and the other for Oracle Linux 6.

To upgrade only Oracle Big Data Connectors, and no other components of the software stack, contact Oracle Support for assistance.

Software downgrades are not supported.



Note:

The upgrade will prompt you to either choose a rolling upgrade or to continue with a conventional upgrade. A conventional upgrade process automatically stops and starts services as needed. Therefore the cluster is unavailable while the `mammoth` command is executing. The advantage of a rolling upgrade is that it may allow uninterrupted access to some services throughout the upgrade process. However, a rolling upgrade adds some extra time per server because the reboots are not concurrent.

Currently, Oracle Big Data Appliance software upgrades do not require corresponding upgrades of the InfiniBand or Cisco switches.

10.7.1 About the Operating System Versions

Oracle Big Data Appliance 2.3 and later software runs on either Oracle Linux version 5 or version 6. When upgrading the software, you can choose whether to upgrade the operating system also:

- **Retaining Oracle Linux 5:** Download the Mammoth bundle for Oracle Linux 5. It retains Oracle Linux 5, but upgrades Oracle Unbreakable Enterprise Kernel (UEK) to version 2. It installs all of the Oracle Big Data Appliance software for this release.
- **Upgrading to Oracle Linux 6:** Download the Mammoth bundle for Oracle Linux 6. To upgrade the operating system, you must first reimage the servers. Reimaging erases all files and data, therefore Oracle does not recommend this type of upgrade for a production system, unless you have a second cluster functioning as a backup with the same data. After reimaging, you can install the software with Oracle Linux 6 and UEK version 2.

10.7.2 Upgrading the Software

Follow these procedures to upgrade the software on an Oracle Big Data Appliance cluster to the current version.

10.7.2.1 Prerequisites

You must know the passwords currently in effect for the cluster, which the Mammoth utility will prompt you for:

- oracle
- root
- Cloudera Manager admin
- MySQL Database admin
- MySQL Database for Oracle Data Integrator (if Oracle Data Integrator agent is installed)

In Oracle Big Data Appliance 4.1, Apache Sentry changes from a file-based provider to the Sentry service. If Sentry is enabled when you upgrade the cluster, then the file-based provider is retained. To switch to the Sentry service, disable Sentry before you upgrade the software, and then re-enable Sentry afterward.

10.7.2.2 Upgrading to the Current Software Version

Making sure cluster services are healthy before upgrade, and especially after reboot is very important. This is even more important in the case of rolling upgrade. In the case of rolling upgrade, if previous errors exist, the parcel upgrade will stop (even if the parcel is already upgraded). This will prevent services from starting. Manual steps will be needed to resume.

Upgrade the Oracle Big Data Appliance software to the current software version as follows. This is a summary. Refer to MOS Doc ID 2101906.1 for details, including prerequisites, further information on the steps below, and known issues.



Note:

All Oozie jobs should be stopped before the upgrade. Failure to do this may cause the upgrade to fail.

1. Download and unzip the Mammoth bundle, as described in "[Downloading the Mammoth Software Deployment Bundle](#) (page 10-4)." You must be logged in as root to node 1 of the primary rack.

2. Change to the BDAMammoth directory.

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

3. Run the mammoth command with the -p option:

```
# ./mammoth -p
```

Mammoth automatically upgrades the base image if necessary.

4. After all nodes reboot, perform the following checks.

- a. Check uptime.

```
# dcli -C uptime
```

- b. Check /root/BDA_REBOOT_SUCCEEDED.

```
# dcli -C ls -ltr /root/BDA_REBOOT_SUCCEEDED
```

Note: If there is no BDA_REBOOT_SUCCEEDED, check for /root/BDA_REBOOT_* and /root/bda_reboot_status.

- c. Verify that the kernel and JDK are upgraded.

```
# dcli -C uname -a
# dcli -C java -version
```

- d. Check that all Cloudera Configuration Manager services are healthy. You may need to manually restart some services.



Note:

During any upgrade, but especially a rolling upgrade, it is **crucial that all services in CM are healthy after the reboot before continuing**. Failure to do so will result in upgrade failures.

5. After the reboot and the post reboot checks, log on to node 1 of the primary rack and rerun mammoth -p in order to resume the upgrade.

```
# cd /opt/oracle/BDAMammoth
# ./mammoth -p
```

6. When the upgrade is complete, perform the post-upgrade validation steps described in the MOS document (2101906.1).

10.7.2.3 Performing Parcel Upgrades Outside of the One-Off Patch Mechanism

Oracle one-off patches for Oracle Big Data Appliance software are by definition supported. Parcel upgrades that are not part of a one-off patch are unsupported unless affirmed by Oracle Support.

The CDH 5.5.2 Parcel Upgrade

Mammoth bundles include a recent version of the parcel for Cloudera's Distribution including Apache Hadoop (CDH). However, CDH is released on Cloudera's schedule and sometimes a newer CDH release than the one in the current Mammoth bundle may be available. A new CDH version may be supported for integration with the current Oracle Big Data Appliance release. Consult with Oracle Support if you have questions.

Oracle Big Data Appliance Release 4.4 includes CDH 5.5.1. A parcel upgrade to CDH 5.5.2 in this release is supported. See MOS Doc ID 2105676.1 for instructions. This upgrade is optional. Before you chose to do the upgrade, consider the issues described in the *Known Limitations* sections of the MOS note.

10.8 Changing the Configuration of Optional Software

During the initial configuration of Oracle Big Data Appliance, the optional software components may or may not be installed. Using the `bdacli` utility, you can reverse some of those decisions. You must provide the relevant server names, ports, user names, and passwords.

This section provides examples of some reconfiguration options. It has the following topics:

- [Changing Support for Oracle Big Data Connectors](#) (page 10-13)
- [Adding Support for Auto Service Request](#) (page 10-16)
- [Adding Support for Oracle Enterprise Manager Cloud Control](#) (page 10-16)
- [Adding Support for Oracle Audit Vault and Database Firewall](#) (page 10-17)
- [Adding Kerberos Authentication](#) (page 10-18)

 Note:

See "[bdacli](#) (page 13-7)" for the complete syntax.

10.8.1 Changing Support for Oracle Big Data Connectors

You can add or remove support for Oracle Big Data Connectors:

10.8.1.1 Adding Oracle Big Data Connectors

When adding support for Oracle Big Data Connectors, you can choose whether to install Oracle Data Integrator agent. If you do, then you must provide passwords for the MySQL Database root user and the Oracle Data Integrator user of MySQL

Database (BDA_ODI_REPO). You must also know these passwords if they are not saved in `cluster_name-config.json`:

- Cloudera Manager admin user
- Oracle Audit Vault and Database Firewall, if enabled

The following procedure uses the `bdaccli` utility.

To add Oracle Big Data Connectors to a cluster:

1. Log in to the first NameNode (node01) of the primary rack as root.
2. Enable Oracle Big Data Connectors:

```
# bdaccli enable bdc
INFO: Logging all actions in /opt/oracle/BDAMammoth/bdaconfig/tmp/
bda1node03-20140805110007.log and traces in /opt/oracle/BDAMammoth/bdaconfig/tmp/
bda1node03-20140805110007.trc
INFO: This is the install of the primary rack
INFO: Checking if password-less ssh is set up
.
.
.
Do you wish to enable ODI? [y/n]: y
Enter password for the BDA_ODI_REPO mysql user
Enter password: odi_password
Enter password again: odi_password
Enter password for the mysql root user
Enter password: root_password
Enter password again: root_password
WARNING: The password for the Cloudera Manager admin user is missing from the
parameters file and is required for the installation.
Enter password: admin_password
Enter password again: admin_password
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
   0      2    0     2    0     0    181     0  --:--:--  --:--:--  --:--:--   250
INFO: The password for audit vault server is not needed since feature is not
enabled
INFO: Creating environment.pp file ...
INFO: Making sure all puppet agents can be accessed.
INFO: Pinging puppet agents
INFO: Adding BDC to the cluster. This will take some time ...
.
.
.
SUCCESS: Successfully reconfigured service
```

10.8.1.2 Removing Oracle Big Data Connectors

When removing support for Oracle Big Data Connectors, you must provide passwords for the following users, if the passwords are not saved in `cluster_name-config.json`:

- Cloudera Manager admin user
- MySQL Database root user, if Oracle Data Integrator agent is enabled
- BDA_ODI_REPO user of MySQL Database, if Oracle Data Integrator agent is enabled
- Oracle Audit Vault and Database Firewall, if enabled

The following procedure uses the `bdacli` utility.

To remove Oracle Big Data Connectors from a cluster:

1. Log in to the first NameNode (node01) of the primary rack as root.
2. Remove Oracle Big Data Connectors:

```
# bdacli disable bdc
INFO: Logging all actions in /opt/oracle/BDAMammoth/bdaconfig/tmp/
bda1node03-20140805104603.log and traces in /opt/oracle/BDAMammoth/bdaconfig/tmp/
bda1node03-20140805104603.trc
INFO: This is the install of the primary rack
INFO: Checking if password-less ssh is set up
INFO: Executing checkRoot.sh on nodes

/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
SUCCESS: Executed checkRoot.sh on nodes

/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
INFO: Executing checkSSHAllNodes.sh on nodes

/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
SUCCESS: Executed checkSSHAllNodes.sh on nodes

/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
INFO: Reading component versions from

/opt/oracle/BDAMammoth/bdaconfig/COMPONENTS
INFO: Creating nodelist files...
WARNING: The password for the Cloudera Manager admin user is missing from the
parameters file and is required for the installation.
Enter password: admin_password
Enter password again: admin_password
  % Total    % Received % Xferd  Average Speed   Time    Time       Time  Current
                                 Dload  Upload    Total   Spent    Left   Speed
    0     2    0     2    0     0    184      0  --:--:--  --:--:--  --:--:--   250
WARNING: The password for the MySQL root user is missing from the parameters
file and is required for the installation.
Enter password: root_password
Enter password again: root_password
INFO: Executing verifyMySQLPasswd.sh on nodes
/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
SUCCESS: Executed verifyMySQLPasswd.sh on nodes

/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
WARNING: The password for the MySQL BDA_ODI_REPO user is missing from the
parameters file and is required for the installation.
Enter password: odi_password
Enter password again: odi_password
INFO: Executing verifyMySQLPasswd.sh on nodes
/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
INFO: The password for audit vault server is not needed since feature is not
enabled
INFO: Creating environment.pp file ...
INFO: Making sure all puppet agents can be accessed.
INFO: Pinging puppet agents
INFO: Removing big data connectors. This will take some time ...
.
.
.
SUCCESS: Successfully reconfigured service
```

10.8.2 Adding Support for Auto Service Request

The following procedure shows how to add support for Auto Service Request.

To support Auto Service Request:

1. Set up your My Oracle Support account and install ASR Manager. You must do this before activating Auto Service Request on Oracle Big Data Appliance. See [Setting Up Auto Service Request](#) (page 5-1).
2. Turn on Auto Service Request monitoring and activate the assets:

```
# bdaccli enable asr
INFO: Logging all actions in /opt/oracle/BDAMammoth/bdaconfig/tmp/
bda1node01-20130205075303.log and traces in /opt/oracle/BDAMammoth/bdaconfig/tmp/
bda1node01-20130205075303.trc
.
.
.
Enter the value for ASR_HOST [Default: ]: asr-host.example.com
Enter the value for ASR_PORT [Default: 162]:
Enter the value for ASR_SERVER_USER: jdoe

Please Check the values you entered for the ASR parameters

ASR_HOST = asr-host.example.com
ASR_PORT = 162
ASR_SERVER_USER = jdoe

Are these values correct (y/n): y
Enter password for user jdoe on machine asr-host.example.com
Enter password: password
Enter password again: password
INFO: Creating environment.pp file ...
INFO: Making sure all puppet agents can be accessed.
INFO: Pinging puppet agents
INFO: Setting up ASR on all nodes. This will take some time ...
.
.
.
```

3. Complete the steps in "[Verifying the Auto Service Request Installation](#) (page 10-7)."

10.8.3 Adding Support for Oracle Enterprise Manager Cloud Control

Oracle Big Data Appliance supports these versions of the Enterprise Manager System Monitoring plug-in: 13.2 PG, 13.2, 13.1, and 12.c. You can install the Enterprise Manager System Monitoring plug-in for Oracle Big Data Appliance in an Oracle Enterprise Manager Cloud Control installation on the same network.

Installing the Oracle Enterprise Manager Cloud Control Plug-In

1. Install the preliminary patches (if required for the plug-in version you choose).

To enable system monitoring over HTTPS for either the 13.1 or 12.c plug-in, you must first install the patches listed in the table below. You can download these patches from the Oracle Automated Release Updates (ARU) site. the 13.2 and 13.2PG plug-ins do not require additional patches.

Enterprise Manager System Monitoring Plug-In	Patch Prerequisites for Support of the Plug-In on Oracle Big Data Appliance
13.2 and 13.2 PG Enterprise Manager System Monitoring Plug-Ins. For Oracle Big Data Appliance, the specific 13.2 PG version is Enterprise Manager for Big Data Appliance (13.2.2.0.0).	No patches needed.
13.1 Enterprise Manager System Monitoring Plug-In	• OMS side patch: 23294904 • Agent plugin : 23294778 • Discovery plugin : 23294785
12.c Enterprise Manager System Monitoring Plug-In	• OMS side patch: 23218275 • BDA Plugin Discovery OH patches: 23218115 and 23218109.

2. Install the Enterprise Manager System Monitoring plug-in.

For instructions on installing the version 13.1 plug-in, see *Oracle Enterprise Manager System Monitoring Plug-in Installation Guide for Oracle Big Data Appliance* . You can use these same instructions to install version 13.2 and 13.2 PG.

For instructions on installing the 12.c plug-in, see this My Oracle Support Document: [Download/Deployment of the 12c EM May Bundle Patch to Support Discovery/Monitoring of Cloudera Manager Admin Console HTTPS URLs \(Doc ID 2148611.1\)](#).

What's New in 13.2 PG for Oracle Big Data Appliance?

The 13.2 PG plug-in for Oracle Big Data Appliance adds support for the following targets:

- Kafka
- Sentry
- Sqoop2
- Key Management Service
- Key Trustee Service
- Cluster Security

10.8.4 Adding Support for Oracle Audit Vault and Database Firewall

Only Release 12.1.1 of Oracle Audit Vault and Database Firewall Server is supported for use with Oracle Big Data Appliance at this time.

Before installing support on Oracle Big Data Appliance, ensure that Oracle Audit Vault and Database Firewall Server is up and running. The software must be installed on a separate server on the same network as Oracle Big Data Appliance.

You must also have the following information about the Audit Vault Server installation:

- Audit Vault Server administration user name and password
- Database service name
- IP address

- Port number

To add support for Oracle Audit Vault and Database Firewall:

1. Log in to the first NameNode (node01) of the primary rack.
2. Add support for Oracle Audit Vault and Database Firewall:

```
# bdaccli enable auditvault
INFO: Logging all actions in /opt/oracle/BDAMammoth/bdaconfig/tmp/
bda1node01-20140805072714.log and traces in /opt/oracle/BDAMammoth/bdaconfig/
tmp/ bda1node01-20140805072714.trc
INFO: This is the install of the primary rack
INFO: Checking if password-less ssh is set up
INFO: Executing checkRoot.sh on nodes

/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
SUCCESS: Executed checkRoot.sh on nodes

/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
INFO: Executing checkSSHAllNodes.sh on nodes

/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
SUCCESS: Executed checkSSHAllNodes.sh on nodes

/opt/oracle/BDAMammoth/bdaconfig/tmp/all_nodes #Step -1#
INFO: Reading component versions from

/opt/oracle/BDAMammoth/bdaconfig/COMPONENTS
INFO: Creating nodelist files...
Please enter the Audit Vault Server Admin Username: admin_username
Please enter the Audit Vault Server Admin Password: admin_password
Enter password again: admin_password
Please enter the Audit Vault Server Database Service Name: service_name
Please enter the Audit Vault Server IP Address: IP address
Please enter the Audit Vault Server Port: port_number
INFO: Creating environment.pp file ...
INFO: Making sure all puppet agents can be accessed.
INFO: Pinging puppet agents
INFO: Adding audit Vault Service. This will take some time ...
.
.
.
```

10.8.5 Adding Kerberos Authentication

The following procedure configures Kerberos authentication.



Note:

If you choose to enable Active Directory Kerberos, first read MOS (My Oracle Support) documents 2029378.1 and 2013585.1. These documents explain required preliminary steps and provide important information on known issues.

To support Kerberos authentication:

1. Ensure that you complete the Kerberos prerequisites listed in "[Installation Prerequisites](#) (page 10-3)."
2. Log into the first NameNode (node01) of the primary rack.
3. Configure Kerberos:

```
# bdacli enable kerberos
INFO: Logging all actions in /opt/oracle/BDAMammoth/bdaconfig/tmp/
bda1node01-20131104072502.log and traces in /opt/oracle/BDAMammoth/bdaconfig/tmp/
bda1node01-20131104072502.trc
.
.
.
INFO: Executing checkRoot.sh on nodes /opt/oracle/BDAMammoth/bdaconfig/tmp/
all_nodes #Step -1#
SUCCESS: Executed checkRoot.sh on nodes /opt/oracle/BDAMammoth/bdaconfig/tmp/
all_nodes #Step -1#
SUCCESS: Password-less root SSH is setup.

Do you want to setup KDC on a BDA node (yes/no): yes

Please enter the realm name: EXAMPLE.COM

Enter password for Kerberos database: password
Enter password again: password
INFO: Executing changekrb5_kdc.sh on nodes bda1node01 #Step -1#
SUCCESS: Executed changekrb5_kdc.sh on nodes bda1node01 #Step -1#
SUCCESS: Successfully set the Kerberos configuration on the KDC
INFO: Setting up Master KDC
INFO: Executing setupKDC.sh on nodes bda1node01 #Step -1#
.
.
.
```

10.9 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](#) (page 1-6)". 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 replaced a failed server. Mammoth automatically upgrades the base image as necessary before upgrading the Oracle Big Data Appliance software to a newer version.

You can reimage all or part of a rack. However, all the servers in a cluster must have the same image. Follow the appropriate instructions:

- [Reimaging a Single Oracle Big Data Appliance Server](#) (page 10-19)
- [Reimaging an Oracle Big Data Appliance Rack](#) (page 10-21)
- [Reimaging an Oracle Big Data Appliance Cluster](#) (page 10-23)

10.9.1 Reimaging a Single Oracle Big Data Appliance Server

Follow this procedure to reimage one server, for example, following replacement of a failed server.

 Caution:

If you reimage a server, then all files and data are erased.

To reinstall the base image on one server:

1. Download the base image patch from My Oracle Support or Oracle Automated Release Updates (ARU), and copy it to the server being reimaged. You can use the base image downloaded with the Mammoth bundle, or download the separate base image patch.

 Caution:

Use the most recent 4.x version of the base image. Do not use the version included in the Mammoth bundle.

See "[Downloading the Mammoth Software Deployment Bundle](#) (page 10-4)." You can take the same basic steps to download the base image patch from My Oracle Support.

2. If you are reimaging the server to the current customer settings, verify that `/opt/oracle/bda/network.json` reflects the intended network configuration. If it does not, then generate a new file using Oracle Big Data Appliance Configuration Generation Utility. See "[Generating the Configuration Files](#) (page 4-3)."
3. Ensure that 4 GB or more disk space are free in the partition:

```
$ df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/md2        161G   23G  130G  15% /
/dev/md0        194M   40M  145M  22% /boot
tmpfs           24G     0   24G   0% /dev/shm
/dev/sda4       1.7T  197M  1.7T   1% /u01
/dev/sdb4       1.7T  197M  1.7T   1% /u02
/dev/sdc1       1.8T  199M  1.8T   1% /u03
/dev/sdd1       1.8T  199M  1.8T   1% /u04
.
.
.
```

4. Unzip the downloaded base image ZIP file. The following example shows a partial listing of the output to the command line:

```
# unzip p<number>_<version>_Linux-x86-64.zip
Archive: p<number>_Linux-x86-64.zip
creating: BDABaseImage-ol6-<version>_RELEASE/
creating: BDABaseImage-ol6-<version>_RELEASE/Extras/
creating: BDABaseImage-ol6-<version>_RELEASE/Extras/BALANCER/
extracting: BDABaseImage-ol6-<version>_RELEASE/Extras/BALANCER/orabalancer-
<version>-h2.zip
inflating: BDABaseImage-ol6-<version>_RELEASE/Extras/BALANCER/orabalancer-
<version>-h2.noarch.rpm
extracting: BDABaseImage-ol6-<version>_RELEASE/Extras/BALANCER/orabalancerdemo-
<version>-h2.zip
creating: BDABaseImage-ol6-<version>_RELEASE/Extras/CELL/
```

```
inflating: BDABaseImage-ol6-<version>_RELEASE/Extras/CELL/bd_cell-  
<version>_LINUX.X64_160506.11  
...
```

5. Change to the subdirectory created in the previous step:

```
# cd BDABaseImage-ol6-<version>_RELEASE
```

6. Reimage the server using the `makebdimage` command. The following example reimages server 4, including the internal USB, from the base image to the custom settings in `network.json`. You must be logged in to the server being reimaged, which is server 4 in this example.

```
./makebdimage -usbint BDABaseImage-<version>_RELEASE.iso /opt/oracle/bda/  
network.json 4
```

7. If the `makebdimage` command succeeds without errors, then restart the server.

10.9.2 Reimaging an Oracle Big Data Appliance Rack

Follow this procedure to reimage an entire rack.

▲ Caution:

If you reimage an entire rack, then all clusters, files, and data on the rack are erased. Reimaging is *not* required for a software upgrade.

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/cluster_name-config.json` file to a safe place outside Oracle Big Data Appliance.
2. Download the most recent base image patch from My Oracle Support or Oracle Automated Release Updates (ARU), and copy it to the first (bottom) server of the rack being reimaged. You can use the base image downloaded with the Mammoth bundle, or download the separate base image patch.

See "[Downloading the Mammoth Software Deployment Bundle](#) (page 10-4)." You can take the same basic steps to download the base image patch from My Oracle Support.

▲ Caution:

Use the most recent 4.x version of the base image. Do not use the version included in the Mammoth bundle.

3. Establish an SSH connection to the first server and log in as `root`.
4. If you are reimaging to existing customer network settings, then verify that `/opt/oracle/bda/network.json` reflects the intended network configuration. If it does not, then generate a new file using Oracle Big Data Appliance Configuration Generation Utility. See "[Generating the Configuration Files](#) (page 4-3)."

5. Ensure that passwordless SSH is set up:

```
# dcli hostname
192.168.41.37: bda1node01.example.com
192.168.41.38: bda1node02.example.com
192.168.41.39: bda1node03.example.com
.
.
.
```

This command must run without errors and return the host names of all Oracle Big Data Appliance servers. If not, then follow the steps in "[Setting Up Passwordless SSH](#) (page 14-1)". Do not continue until the `dcli hostname` command runs successfully on all servers.

6. Check all Oracle Big Data Appliance servers for hardware issues:

```
# dcli bdacheckhw | grep -v SUCCESS
```

7. Resolve any hardware errors and warnings before reimaging the rack.
8. Verify that at least 4 GB are available in the root (/) partition of all servers:

9.

```
# dcli df -h /
```

192.168.41.37: Filesystem	Size	Used	Avail	Use%	Mounted on
192.168.41.37: /dev/md2	161G	21G	132G	14%	/
192.168.41.38: Filesystem	Size	Used	Avail	Use%	Mounted on
192.168.41.38: /dev/md2	161G	19G	135G	12%	/
192.168.41.39: Filesystem	Size	Used	Avail	Use%	Mounted on
192.168.41.39: /dev/md2	161G	23G	131G	15%	/
.					
.					
.					

10. Unzip the downloaded base image ZIP file. For example:

```
# unzip p<number>_<version>_Linux-x86-64.zip
Archive:  p<number>_<version>_Linux-x86-64.zip
creating: BDABaseImage-ol6-<version>_RELEASE/
inflating: BDABaseImage-ol6-<version>_RELEASE/biosconfig
inflating: BDABaseImage-ol6-<version>_RELEASE/makebdaimage
extracting: BDABaseImage-ol6-<version>_RELEASE/BDABaseImage-ol6-
<version>_RELEASE.md5sum
inflating: BDABaseImage-ol6-<version>_RELEASE/reimagerack
inflating: BDABaseImage-ol6-<version>_RELEASE/reimagecluster
inflating: BDABaseImage-ol6-<version>_RELEASE/BDABaseImage-ol6-
<version>_RELEASE.iso
.
.
.
```

11. Change to the subdirectory created in the previous step:

```
# cd BDABaseImage-ol6-<version>_RELEASE
```

12. Complete one of the following procedures:

- To reimage an Oracle Big Data Appliance that was configured for a customer network to the same customer network settings, execute the `./reimagerack` command.
- To reimage an appliance that still has the factory settings:
 - a. Ensure that `/opt/oracle/bda/network.json` does not exist.

- b.** Execute the `./reimagerack` command.
- To restore the factory network settings on a rack configured with custom network settings:
- a.** Copy `/opt/oracle/bda/network.json` to a safe location outside Oracle Big Data Appliance.
 - b.** Ensure that `/opt/oracle/bda/BdaShip.json` exists.
 - c.** Reimage the rack:
`./reimagerack deploy ship`

See [reimagerack](#) (page 10-34) for the complete syntax of the command.

- 13.** Run Mammoth. See "[Installing the Software](#) (page 10-6)."

10.9.3 Reimaging an Oracle Big Data Appliance Cluster

Follow this procedure to reimage a group of servers that the Mammoth utility has deployed as a single cluster. This can only be used on a group of servers that has been deployed as a single cluster using Mammoth. You can either reapply the network settings that are in currently in effect or revert to factory settings.

 Caution:

If you reimage a cluster, then all files and data on the cluster are deleted. Reimaging is *not* required for a software upgrade.

Mandatory Prerequisites

You must rename the file `<rack_name>-network.json` generated by the Oracle Big Data Appliance configuration utility to `network.json`. **Before reimaging, `/opt/oracle/bda/network.json` must exist on all nodes of the cluster.**

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

 Tip:

The `reimagecluster` utility creates an installer on the internal USB drive of each server in the rack, initializes the installation, and reboots each server. While the servers are down for reboot, you can log on to Integrated Lights Out Manager (iLOM) on each server to monitor progress:

```
<server name>-ilom.<domain>.com
```

Click on **Remote Control -> Redirection -> Launch Remote Console** to launch a Java console where you can monitor the re-imaging.

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 (or `/opt/oracle/BDAMammoth/<cluster_name>-config.json`) to a safe place outside Oracle Big

Data Appliance. If the cluster has been upgraded, the file may exist in `/opt/oracle/BDAMammoth/previous-BDAMammoth` directory.

2. Download the latest Oracle Big Data Appliance Base Image tarball from My Oracle Support to a temporary directory. Copy it onto the first (bottom) server of the cluster to be re-imaged.

For Oracle Big Data Appliance 4.8, use the 4.5 base image, which is [Patch 25395169: ORACLE BIG DATA APPLIANCE BASE IMAGE VERSION 4.5.0-4 FOR ORACLE LINUX 6](#). See this MOS note for additional details on reimaging: [Oracle Big Data Appliance Base Image Version 4.5.0 for New Installations or Reprovisioning Existing Installations on Oracle Linux 6 \(Doc ID 2173434.1\)](#)

3. If re-imaging to your current settings, verify that `/opt/oracle/bda/network.json` reflects the intended network configuration. If it does not, then generate a new file using Oracle Big Data Appliance Configuration Generation Utility.
4. Establish an SSH connection to the first server and log in as root.
5. Ensure that password-less SSH has been setup on the cluster (`dcli -C hostname` should run without errors). If it has not been setup, run `setup-root-ssh -C`.

```
# setup-root-ssh -C
Enter root password:
...
...
setup-root-ssh succeeded
```

`dcli -C hostname` must run without errors and return the host names of all Oracle Big Data Appliance servers in the cluster. If not, then follow the steps in "[Setting Up Passwordless SSH](#) (page 14-1)". Do not continue until the `dcli -C hostname` command runs successfully on all servers.

6. Ensure that there is at least 4 GB free in the `/` partition of all servers:

```
# dcli -C df -h /
10.xxx.xxx.xx: Filesystem      Size  Used Avail Use% Mounted on
10.xxx.xxx.xx: /dev/md2          459G  49G  387G  12% /
10.xxx.xxx.xx: Filesystem      Size  Used Avail Use% Mounted on
...
```

7. Unzip the downloaded base image ZIP file. For example:

```
# unzip p<number>_<version>_Linux-x86-64.zip
Archive: p<number>_<version>_Linux-x86-64.zip
creating: BDABaseImage-ol6-<version>_RELEASE/
creating: BDABaseImage-ol6-<version>_RELEASE/Extras/
creating: BDABaseImage-ol6-<version>_RELEASE/Extras/BALANCER/
extracting: BDABaseImage-ol6-<version>_RELEASE/Extras/BALANCER/orabalancer-
<version>-h2.zip
inflating: BDABaseImage-ol6-<version>_RELEASE/Extras/BALANCER/orabalancer-
<version>-h2.noarch.rpm
extracting: BDABaseImage-ol6-<version>_RELEASE/Extras/BALANCER/orabalancerdemo-
<version>-h2.zip
creating: BDABaseImage-ol6-<version>_RELEASE/Extras/CELL/
inflating: BDABaseImage-ol6-<version>_RELEASE/Extras/CELL/bd_cell-
<version>_LINUX.X64_160506.11
...
```

8. Change to the subdirectory created in the previous step:

```
# cd BDABaseImage-ol6-<version>_RELEASE
```

9. Use one of the following procedures to re-image the entire Oracle Big Data Appliance. The first alternative is to retain your existing configuration . The second is to restore to factory settings (prior to Oracle Big Data Appliance Configuration Utility setup) .

Reimage and Retain Existing Network Settings

- a. At /opt/oracle/bda/, check that both network.json and BdaShip.json are listed.

```
# ls *.json
network.json BdaShip.json
```

If network.json is not listed, locate the customer-specific network.json file (the one that describes the configuration you want to keep). First check to ensure that this file defines correct network settings. If it does not, generate a new file using Oracle Big Data Appliance Configuration Generation Utility. Then copy the file to /opt/oracle/bda and rename it to network.json.

- b. Use dcli to copy network.json to all nodes of the cluster.

```
# dcli -C -f /opt/oracle/bda/network.json -d /opt/oracle/bda/
```

- c. Execute the ./reimagecluster command. Partial output of the command is shown below.

```
# ./reimagerack
Passwordless SSH set up to all hosts : xxx.xxx.xx.xx xxx.xxx.xx.xx
xxx.xxx.xx.xx xxx.xxx.xx.xx xxx.xxx.xx.xx
xxx.xxx.xx.xx
Using /tmp/BDABaseImage-ol6-<version>_RELEASE/BDABaseImage-ol6-
<version>_RELEASE.iso to re-image
BDABaseImage-ol6-<version>_RELEASE.iso: OK
Re-imaging to existing customer network settings

FINAL CONFIRMATION : All of the following hosts will be completely erased,
re-imaged and initialized including resetting of the iloms: xxx.xxx.xx.xx
xxx.xxx.xx.xx xxx.xxx.xx.xx xxx.xxx.xx.xx xxx.xxx.xx.xx xxx.xxx.xx.xx
Proceed [y/n]?
y
Running hardware checks on all nodes...
...
Rebooting all servers
Broadcast message from root@bda1node01.example.com (unknown) at 15:58 ...
The system is going down for reboot NOW!
```

To Reimage to Factory Network Settings

- a. As described in Step 1, check that /opt/oracle/bda/network.json has been copied to a safe location outside Oracle Big Data Appliance and that /opt/oracle/bda/BdaShip.json exists.
- b. Reimage the cluster to factory-shipped settings with the following command:

```
# ./reimagecluster deploy ship
```

10. Whether you re-imaged with existing network settings or restored to factory settings, the final steps are the same. After reboot from reimagecluster completes, run setup-root-ssh to re-establish a connection to all the servers in the rack. This is necessary so that the dcli command can be run to reboot all the servers again.

```
# setup-root-ssh
Enter root password:
```

```
...
setup-root-ssh succeeded
```

11. Reboot all servers once again using the `dcli reboot` command.

```
# dcli reboot
```

12. Verify that the files `/root/BDA_IMAGING_SUCCEEDED` and `/root/BDA_REBOOT_SUCCEEDED` exist on each server in the rack.

```
# dcli ls -lrt /root/BDA_*
```

Check that the date and timestamp of `/root/BDA_REBOOT_SUCCEEDED` correspond to the time of the reboot in the reimaging process. The file may be zero-byte or may contain some INFO messages.

13. Verify the installed base image version:

```
# imageinfo
Big Data Appliance Image Info
IMAGE_VERSION : <version>
IMAGE_CREATION_DATE : Wed May 18 20:19:54 UTC 2016
IMAGE_LABEL : BDA-<version>_LINUX.X64_RELEASE
LINUX_VERSION : Oracle Linux Server release <version>
KERNEL_VERSION : 2.<version>.el6uek.x86_64
BDA_RPM_VERSION : bda-<version>-1.el6.x86_64
OFED_VERSION : OFED-IOV-<version>
JDK_VERSION : jdk1.<version>-fcs.x86_64
```

You can now run the Mammoth utility to bring the system up to the latest Oracle Big Data Appliance version. See [About the Mammoth Utility](#) (page 10-1).

10.10 Installing a One-Off Patch

One-off patch bundles provide a fix to specific bugs in one or more releases. You use Mammoth to apply the patch to your cluster.

To install a one-off patch bundle:

1. Download the patch bundle from the Automated Release Update (ARU) system to a directory such as `/tmp` on the first node of the Oracle Big Data Appliance cluster.

The file is named `BDA-patch-release-patch.zip`. The examples in this procedure use the name `BDA-patch-<version>-123456.zip`.

2. Unzip the file. For example:

```
# unzip BDA-patch-<version>-123456.zip
Archive:  BDA-patch-<version>-123456.zip
  creating: BDA-patch-<version>-123456/
  inflating: BDA-patch-<version>-123456/BDA-patch-<version>-123456.run
  inflating: BDA-patch-<version>-123456/README.txt
```

3. Change to the patch directory created in Step 2. For example:

```
$ cd BDA-patch-<version>-123456
```

4. Extract the contents of the run file. For example:

```
$ ./BDA-patch-<version>-123456.run
Big Data Appliance one-off patch 123456 for v<version> Self-extraction

Removing existing temporary files
```

```
Generating /tmp/BDA-patch-<version>-123456.tar
Verifying MD5 sum of /tmp/BDA-patch-<version>-123456.tar
/tmp/BDA-patch-<version>-123456.tar MD5 checksum matches
```

```
Extracting /tmp/BDA-patch-<version>-123456.tar to /opt/oracle/BDAMammoth/patches/
123456
Removing temporary files
```

Please "cd /opt/oracle/BDAMammoth" before running "./mammoth -p 123456"

5. Change to the BDAMammoth directory:

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

6. Install the patch. For example:

```
$ ./mammoth -p 123456
```

Mammoth prompts you to choose whether or not you want to load the patch as rolling upgrade. (See "[About Rolling Upgrades](#) (page 10-2)")

Alternatively, you can use the `bdacli` command. See "[bdacli](#) (page 13-7)".

10.11 Mammoth Software Installation and Configuration Utility

You must log in as root on the first server and change to the /opt/oracle/BDAMammoth directory to use Mammoth. It has this syntax:

```
./mammoth option [cluster_name] ]
```

In this command, *cluster_name* is the name of the cluster You must enter *cluster_name* in the first command exactly as it appears in (*cluster_name*-config.json). Afterward, *cluster_name* defaults to the rack specified in a previous mammoth command.

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

Example 10-1 Mammoth Syntax Examples

This command displays Help for the Mammoth utility:

```
./mammoth -h
```

This command does a complete install on rack bda3:

```
./mammoth -i bda3
```

This command runs steps 2 through 6 on the rack being set up:

```
./mammoth -r 2-6
```

This command generates a parameter file to add six new servers over a starter rack, beginning with node07, to an existing cluster:

```
./mammoth -e node07 node08 node09 node10 node11 node12
```

10.11.1 Mammoth Options

The syntax of the `mammoth` command supports the configuration of new clusters. You can also use the Mammoth bundle to upgrade from earlier releases.

-c

Run the Oracle Big Data Appliance cluster checks.

-e *newnode1 newnode2 newnode3...*

Generates a parameter file for the list of servers to be added to a cluster. It then adds the new node or nodes to the cluster and installs required software on these nodes. The servers in the list can be in the same rack or multiple racks.

Reboots normally occur only for the new nodes and only at the initial update of the factory image.

The parameter file generated is `opt/oracle/BDAMammoth/<cluster_name>-config.json`, where `<cluster_name>` is the name of the cluster being extended. No passwords are included in the parameter file, so you must enter them when running Mammoth. Cluster extension expects the default root password on nodes in the extension, so you should retain that password until the extension has completed. Be sure to change passwords afterward.

On Oracle NoSQL Database clusters, Mammoth prompts for the kind of zone for the new nodes. You can choose from an existing zone, a new primary zone, or a new secondary zone. When adding to an existing zone, Mammoth lists the zones that you can use. When creating a new zone, Mammoth prompts for the zone name and replication factor.



Note: If Oracle Big Data SQL is Enabled:

In release 4.8.0, `mammoth -e` does not make necessary Oracle Big Data SQL changes. Correct the problem by calling `./setup-bds extend bds-config.json` after `mammoth -e` completes.

-h

Displays command Help including command usage and a list of steps.

-i *cluster_name*

Runs all mandatory steps on the cluster, equivalent to `-r 1-18` for a full rack. Use this option when configuring a new rack or adding a group of servers to a cluster.

-l

Lists the steps of the Mammoth utility.

-p

Upgrades the software on the cluster to the current version or installs a one-off patch.

-r *n-N*

Runs steps *n* through *N* of Mammoth while no errors occur.

-s *n* [*cluster_name*]

Runs step *n*. Enter *cluster_name* to identify another cluster on the same rack. See the `-e` option.

-v

Displays the version number of the Mammoth.

10.11.2 Mammoth Installation Steps

Following are descriptions of the steps that the Mammoth performs when installing the software.

Step 1 PreinstallChecks

This step performs several tasks:

- Validates the configuration files and prompts for the passwords.
- 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.
- Sets up passwordless SSH for the root user on the InfiniBand network.
- 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 an alias to identify the node where the Mammoth is run as the puppet master node. For example, if you run the Mammoth from `bda1node01` with an IP address `192.168.41.1`, then a list of aliases for that IP address includes `bda1node01-master`. The Mammoth uses Puppet for the software installation.
- 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.

This step also performs a variety of hardware and software checks. A failure in any of these checks causes the Mammoth 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 2 SetupPuppet

This step configures puppet agents on all nodes and start them, configures a puppet master on the node where the Mammoth is being run, waits for the agents to submit their certificates, and automates their signing. This step also changes the root password on all nodes (optional). 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 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?"](#) (page 10-8).

Step 3 UpdateBaselImage

Installs the most recent Oracle Big Data Appliance image and system parameter settings. The system automatically restarts if the kernel is upgraded.

Step 4 PrepareBaselImage

- Copies third-party licenses to `/opt/oss/src/OSSLICENSES.pdf` on every server, as required by the licensing agreements.
- Copies third-party software source code to `/opt/oss/src/` on every server, as required by the licensing agreements.

Mammoth does not copy the source code to Oracle NoSQL Database clusters.

- 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.

- The NameNode 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.

Step 5 SetupMySQL

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

Mammoth does not install MySQL Database on Oracle NoSQL Database clusters.

Step 6 SetupKDC

Configures the Kerberos Key Distribution Center (KDC) that is used to set up security on the cluster.

Step 7 InstallHadoop

Installs all packets in Cloudera's Distribution including Apache Hadoop (CDH) and Cloudera Manager. It then starts the Cloudera Manager server on node03 and configures the cluster.

Mammoth does not install CDH or Cloudera Manager on Oracle NoSQL Database clusters.

Step 8 StartHadoopServices

Starts the agents on all nodes and starts all CDH services. Also enables HTTPS for Cloudera Manager, Hue, and Oozie if HTTPS/ Network Encryption was enabled. After this step, you have a fully functional Hadoop installation.

Cloudera Manager runs on port 7180 of node03. You can open it in a browser, for example:

```
http://bda1node03.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. You are prompted to change the password at the end of this step.

Note that Mammoth does not install or start CDH services on Oracle NoSQL Database clusters.

Step 9 InstallBDASoftware

Installs the server-side components of Oracle Big Data Connectors, if this option was selected in Oracle Big Data Appliance Configuration Generation Utility. Oracle Big Data Connectors must be licensed separately. Optional.

Installs the server-side components of Oracle Big Data SQL, if this option was selected in Oracle Big Data Appliance Configuration Generation Utility. For Oracle Big Data SQL support of Oracle NoSQL Database, this step installs the client libraries (kvclient.jar) on the CDH nodes. Oracle Big Data SQL must be licensed separately. Optional.

Installs Oracle NoSQL Database on clusters allocated to its use. Enterprise Edition requires a separate license.

Step 10 SetupHadoopSecurity

Actions performed in this step are based upon the choices you made in the Oracle Big Data Appliance Configuration Utility prior to installation.

- Configures HTTPS/ Network Encryption if this option was selected.
- Configures HDFS Transparent Encryption if this option was selected.

Mammoth sets up the The Navigator Key Trustee services required by HDFS Transparent Encryption based on whether you chose local or external Key Trustee servers in the configuration utility prior to installation:

- If you had selected **Setup Key Trustee Servers on the BDA** in the configuration utility, Mammoth will setup up active and passive Key Trustee servers locally on the Oracle Big Data Appliance.
- If you chose to use external Key Trustee Servers, the connection to the external servers is configured in this setup. (Note that you must have already configured the external Navigator Key Trustee Servers as described in MOS Doc ID 2112644.1)
- Configures Kerberos authentication if this option was selected. There are no prerequisites if you set up the key distribution center on Oracle Big Data Appliance. Otherwise, see "[Installation Prerequisites](#) (page 10-3)."

Step 11 SetupASR

Installs and configures Auto Service Request (ASR). Optional.

This step does the following:

- Installs the required software packages
- Configures the trap destinations
- Starts the monitoring daemon

To activate the assets from ASR Manager, see "[Verifying ASR Assets](#) (page 5-4)".



Note:

For this step to run successfully, the ASR host system must be up with ASR Manager running and configured properly. See [Setting Up Auto Service Request](#) (page 5-1).

Step 12 CleanupInstall

Performs the following:

- 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 13 CleanupSSHroot (Optional)

Removes passwordless SSH for root that was set up in "PreinstallChecks".

10.12 Oracle Big Data Appliance Base Imaging Utilities

The following utilities are distributed in the base image bundle. To run the utilities, you must be logged in as root.

- [makebdaimage](#) (page 10-32)
- [reimagecluster](#) (page 10-33)
- [reimagerack](#) (page 10-34)

10.12.1 makebdaimage

Reinstalls the base image on a single server, a cluster, or an entire rack. Both [reimagecluster](#) (page 10-33) and [reimagerack](#) (page 10-34) call this utility.

The makebdaimage utility has this syntax:

```
makebdaimage [--usb | --usbint] [--noiloms] /path/BDABaseImage-version_RELEASE.iso /
path/BDdaDeploy.json target_servers
```

Options

--usb | --usbint

Identifies the USB port that will be used for reimaging. Use --usb for the external USB drive, or --usbint for the internal drive. The internal USB drive contains a full Linux installation.

To use the `--usbint` option, you must be logged in to the target server; otherwise, you reimage the server using the wrong network information.

--noiloms

The reimaging process does not alter the Oracle ILOMs on the target servers.

target_servers

One or more servers where the image will be installed, which you specify using one of these formats:

- **node_number**
Identifies one server for reimaging. Enter the position of the server in the rack, from 1 to 18, counting from bottom to top.
- **from.json to.json**
Identifies the current configuration (*from.json*) and the desired configuration (*to.json*). The JSON files can be either the factory configuration in *BdaShip.json* or a custom configuration in *network.json*. You can find them in */opt/oracle/bda* on configured servers.
- **to.json**: Concatenation of JSON files similar to *network.json* and *BdaShip.json*, but containing an extra `ETH0_MACS` array entry. The first entry in the concatenated file with a matching `eth0_mac` entry is used when reimaging. The *to.json* file is used as is.

10.12.2 reimagedcluster

Reimages all servers in the cluster in parallel using `dcli` and `makebdaimage`.

The `reimagedcluster` utility has this syntax:

```
reimagedcluster [--no-iloms] [from.json [to.json]]
```

Prerequisites

- Verify that the following command returns the list of servers in the cluster:

```
$ dcli -C hostname
```
- Ensure that exactly one `BDABaseImage-version_RELEASE*.iso` file is in the current directory.

Options**--no-iloms**

The reimaging process does not alter the Oracle ILOMs on the target servers.

from.json

The full path to the current configuration file, either *BdaShip.json* or *network.json*. This option defaults to */opt/oracle/bda/network.json*. If *network.json* is missing, then the option defaults to */opt/oracle/bda/BdaShip.json*. The servers must already be set to the values in the JSON file used by `reimagedcluster`.

to.json

The full path to the new network configuration in effect after reimaging. It can be either a *network.json* or a *BdaShip.json* file.

10.12.3 reimagerack

Reimages all servers in the rack in parallel using `dcli` and `makedbaimage`.

The `reimagerack` utility has this syntax:

```
reimagerack [--no-iloms] [--no-macs] [--hosts n1, n2...] [from.json [to.json]]
```

Prerequisites

- Verify that the following command returns the list of servers in the rack:

```
$ dcli -hostname
```
- Ensure that exactly one `BDABaseImage-version_RELEASE*.iso` file is in the current directory.

Options

--no-iloms

The reimaging process does not alter the Oracle ILOMs on the target servers.

--no-macs

The utility does not retrieve the server MAC addresses. Instead, it uses the InfiniBand cabling. Use this option only when restoring the factory settings; you must include both JSON files (*from.json* and *to.json*) in the command line.

--hosts *n1, n2...*

Restricts reimaging to the servers identified by a comma-delimited list of host names or IP addresses, whichever one `dcli` accepts. All servers must be in the list of target servers identified by the `dcli -t` command.

This option enforces the `--no-macs` option, so its restrictions also apply.

from.json

The full path to the current configuration file, either `BdaShip.json` or `network.json`.

This option defaults to `/opt/oracle/bda/network.json`. If `network.json` is missing, then the option defaults to `/opt/oracle/bda/BdaShip.json`. The servers must already be set to the values in the JSON file used by `reimagerack`.

to.json

The full path to the new network configuration in effect after reimaging. It can be either a `network.json` or a `BdaShip.json` file.

Part IV

Maintenance

Part IV describes procedures and utilities that you use to maintain Oracle Big Data Appliance. It contains the following chapters:

- [Maintaining Oracle Big Data Appliance](#) (page 11-1)
- [Replacing a Server Disk](#) (page 12-1)
- [Oracle Big Data Appliance Utilities](#) (page 13-1)
- [Executing Commands Across a Cluster Using the dcli Utility](#) (page 14-1)
- [Using Oracle Integrated Lights Out Manager](#) (page 15-1)

11

Maintaining Oracle Big Data Appliance

This chapter describes how to monitor and maintain Oracle Big Data Appliance. Some 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](#) (page 11-1)
- [Powering On and Off Oracle Big Data Appliance](#) (page 11-2)
- [Adding Memory to the Servers](#) (page 11-4)
- [Maintaining the Physical Disks of Servers](#) (page 11-6)
- [Maintaining the InfiniBand Network](#) (page 11-7)
- [Changing the Number of Connections to a Gateway Switch](#) (page 11-12)
- [Changing the NTP Servers](#) (page 11-12)
- [Monitoring the PDU Current](#) (page 11-13)
- [Node Migration](#) (page 11-15)



See Also:

[Executing Commands Across a Cluster Using the `dcli` Utility](#) (page 14-1)

11.1 Monitoring the Ambient Temperature of Servers

Maintaining environmental temperature conditions within design specification for a 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](#) (page 14-1)".
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 2-13](#) (page 2-12).

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
bda1node04-adm.example.com: value = 23.000 degree C
.
.
.
```

11.2 Powering On and Off Oracle Big Data Appliance

This section includes the following topics:

- [Nonemergency Power Procedures](#) (page 11-2)
- [Emergency Power-Off Considerations](#) (page 11-3)
- [Cautions and Warnings](#) (page 11-4)

11.2.1 Nonemergency Power Procedures

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



See Also:

Oracle Big Data Appliance Software User's Guide for powering on and off gracefully when the software is installed and running.

11.2.1.1 Powering On Oracle Big Data Appliance

To turn on Oracle Big Data Appliance:

- Turn on all 12 breakers on both PDUs.

Oracle ILOM and the Linux operating system start automatically.

11.2.1.2 Powering Off Oracle Big Data Appliance

To turn off Oracle Big Data Appliance:

1. Turn off the servers.
2. Turn off all 12 breakers on both PDUs.

11.2.1.2.1 Powering Off the Servers

Use the Linux shutdown command to turn 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

11.2.1.2.2 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](#) (page 14-1)".

The following command shows the syntax of the command:

```
# 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:**

[Executing Commands Across a Cluster Using the dcli Utility](#) (page 14-1)

11.2.1.3 Powering On and Off Network Switches

The network switches do not have power switches. They turn off when power is removed by turning off a PDU or a breaker in the data center.

11.2.2 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

11.2.2.1 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.

11.2.2.2 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.

11.2.3 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 turn off Oracle Big Data Appliance unless there is an emergency. In that case, follow the "[Emergency Power-Off Procedure](#) (page 11-3)".
- 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.

11.3 Adding Memory to the Servers

You can add memory to all servers in the cluster or to specific servers.

See Also:

"[Parts for Oracle Big Data Appliance Servers](#) (page B-5)"

11.3.1 Adding Memory to an Oracle Server X6-2L

Oracle Big Data Appliance X6-2L ships from the factory with 256 GB of memory in each server. Eight of the 24 slots are populated with 32 GB DIMMs. Memory is expandable up to 768 GB (in the case of 32 GB DIMMs in all 24 slots).

See the [Oracle® Server X6-2L Service Manual](#) for instruction on DIMM population scenarios and rules, installation, and other information.

11.3.2 Adding Memory to an Oracle Server X5-2L, Sun Server X4-2L, or Sun Server X3-2L

Oracle Big Data Appliance X5-2 ships from the factory with 128 GB of memory in each server. Eight of the 24 slots are populated with 16 GB DIMMs. Memory is expandable up to 768 GB (in the case of 32 GB DIMMs in all 24 slots).

Oracle Big Data Appliance X4-2 and X3-2 servers are shipped with 64 GB of memory. The eight slots are populated with 8 GB DIMMs. These appliances support 8 GB, 16 GB, and 32 GB DIMMs. You can expand the amount of memory for a maximum of 512 GB (16 x 32 GB) in a server. You can use the 8 * 32 GB Memory Kit.

You can mix DIMM sizes, but they must be installed in order from largest to smallest. You can achieve the best performance by preserving symmetry. For example, add four of the same size DIMMs, one for each memory channel, to each processor, and ensure that both processors have the same size DIMMs installed in the same order.

To add memory to an Oracle Server X5-2L, Sun Server X4-2L, or Sun Server X3-2L:

1. If you are mixing DIMM sizes, then review the DIMM population rules in the *Oracle Server X5-2L Service Manual* at
http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.gnvje.html#scrolltoc
2. Power down the server.
3. Install the new DIMMs. If you are installing 16 or 32 GB DIMMs, then replace the existing 8 GB DIMMs first, and then replace the plastic fillers. You must install the largest DIMMs first, then the next largest, and so forth. You can reinstall the original DIMMs last.

See the *Oracle Server X5-2L Service Manual* at

http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.ceiebfdg.html#scrolltoc

4. Power on the server.

11.3.3 Adding Memory to Sun Fire X4270 M2 Servers

Oracle Big Data Appliance ships from the factory with 48 GB of memory in each server. 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.

To add memory to a Sun Fire X4270 M2 server:

1. Power down the server.
2. Replace the plastic fillers with the DIMMs. See the *Sun Fire X4270 M2 Server Service Manual* at
http://docs.oracle.com/cd/E19245-01/E21671/motherboard.html#50503715_71311
3. Power on the server.

11.4 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 Oracle Big Data Appliance Servers](#) (page B-5)" for the repair procedures

11.4.1 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.

11.4.1.1 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.

11.4.1.2 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
```

11.4.1.3 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  
.  
.  
.
```

11.5 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](#) (page 15-7)
- [Replacing a Failed InfiniBand Switch](#) (page 11-7)
- [Verifying InfiniBand Network Operation](#) (page 11-9)
- [Understanding the Network Subnet Manager Master](#) (page 11-11)

11.5.1 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:

- "In-Rack InfiniBand Switch-to-Server Cable Connections (page E-8)" for information about cabling
- *Sun Network QDR InfiniBand Gateway Switch Installation Guide* at http://docs.oracle.com/cd/E26699_01/html/E26706/gentextid-125.html#scrolltoc
- *Sun Datacenter InfiniBand Switch 36 User's Guide* at http://docs.oracle.com/cd/E26698_01/html/E26434/gentextid-119.html#scrolltoc

To replace a failed InfiniBand switch:

1. Turn off both power supplies on the switch by removing the power plugs.
2. 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.
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](#) (page 15-7)".
6. Connect to the switch as `ilom_admin` and open the Fabric Management shell:

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

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

7. Disable the Subnet Manager:

```
FabMan@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 reenabling the Subnet Manager on all the other switches.

11.5.2 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/E26699_01/html/E26706/gentextid-28027.html#scrolltoc

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:

```
# bdacheksw

[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 : 5.6
[SUCCESS: All required programs are accessible in $PATH
[SUCCESS: All required RPMs are installed and valid
[SUCCESS: Big Data Appliance software validation checks succeeded
```

11.5.3 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.
- For multirack configurations joining different types of racks, such as Oracle Big Data Appliance and Exalogic, see My Oracle Support note [1682501.1](#).
- 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/E26699_01/index.html
- Sun Datacenter InfiniBand Switch 36 library at http://docs.oracle.com/cd/E26698_01/index.html

11.6 Changing the Number of Connections to a Gateway Switch

If you change the number of 10 GbE connections to a Sun Network QDR InfiniBand Gateway switch, then you must run the `bdaredoclientnet` utility. See "[bdaredoclientnet](#) (page 13-23)."

To re-create the VNICs in a rack:

1. Verify that `/opt/oracle/bda/network.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/network.json
```
2. Connect to node01 (bottom of rack) using the administrative network. 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](#) (page 14-1)" 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 7-2](#) (page 7-28).

6. Restore passwordless SSH (optional):

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

11.7 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 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](#) (page 7-14)".

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](#) (page 7-3)"

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

11.8 Monitoring the PDU Current

The PDU current can be monitored directly. Configure threshold settings as a way 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/html/E23956/index.html>

Table 11-1 (page 11-14) lists the threshold values for the Oracle Big Data Appliance rack using a single-phase, low-voltage PDU.

Table 11-1 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 11-2 (page 11-14) lists the threshold values for the Oracle Big Data Appliance rack using a three-phase, low-voltage PDU.

Table 11-2 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 11-3 (page 11-14) lists the threshold values for the Oracle Big Data Appliance rack using a single-phase, high-voltage PDU.

Table 11-3 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 11-4 (page 11-15) lists the threshold values for the Oracle Big Data Appliance rack using a three-phase, high-voltage PDU.

Table 11-4 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

11.9 Node Migration

There is no automated failover for the cluster nodes. You must perform the recovery process manually and the process differs for critical and non-critical nodes.

Nodes 1 through 4 host services that are essential to Oracle Big Data Appliance. Therefore, these are critical nodes. Recover from a critical node failure by migrating all of its services to another node. Nodes 5 and above are generally DataNodes. Since data is replicated across three nodes, the failure of a single DataNode is non-critical. To recover from a non-critical node failure, attempt to restart the node and if that fails, replace it.

Use the `bdacli admin_cluster` (page 13-7) command to decommission, commission, and migrate nodes.

See Also:

See Managing a Hardware Failure in the *Oracle Big Data Appliance Software User's Guide* for instructions how to replace a failing CDH or NoSQL cluster node.

12

Replacing a Server Disk

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

This chapter contains the following sections:

- [Overview of the Disk Replacement Process](#) (page 12-1)
- [About Disk Drive Identifiers](#) (page 12-2)
- [Prerequisites for Replacing a Failing Disk](#) (page 12-4)
- [What If a Server Fails to Restart?](#) (page 12-6)
- [Replacing a Disk Drive](#) (page 12-7)
- [Correcting a Mounted Partitions Error](#) (page 12-9)
- [Identifying the Function of a Disk Drive](#) (page 12-10)
- [Configuring an Operating System Disk](#) (page 12-11)
- [Configuring an HDFS or Oracle NoSQL Database Disk](#) (page 12-19)
- [Verifying the Disk Configuration](#) (page 12-21)



See Also:

My Oracle Support Doc ID 1581331.1 [My Oracle Support Doc ID 1581331.1](#)

12.1 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, as an operating system disk, or as either an HDFS disk or an Oracle NoSQL Database disk.
4. Configure the disk for its dedicated function.
5. Verify that the configuration is correct.
6. Install the Oracle Big Data Appliance software.

**See Also:**

"Servicing Storage Drives and Rear Drives" in the *Oracle Server X6-2L Service Manual* at

http://docs.oracle.com/cd/E62172_01/html/E62184/z400001c165586.html#scrolltoc

"Servicing Storage Drives and Rear Drives" in the *Oracle Server X5-2L Service Manual* at

http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40000091011460.html#scrolltoc

"Servicing Storage Drives and Boot Drives" in the *Sun Fire X4270M2 Server Service Manual* at

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

12.2 About Disk Drive Identifiers

The Oracle Big Data Appliance servers contain a disk enclosure cage that is controlled by the host bus adapter (HBA). The enclosure holds 12 disk drives that are identified by slot numbers 0 to 11. The drives can be dedicated to specific functions, as shown in [Table 12-1](#) (page 12-3).

Oracle Big Data Appliance uses symbolic links, which are defined in `/dev/disk/by_hba_slot`, to identify the slot number of a disk. The links have the form `snpm`, where `n` is the slot number and `m` is the partition number. For example, `/dev/disk/by_hba_slot/s0p1` initially corresponds to `/dev/sda1`.

When a disk is hot swapped, the operating system cannot reuse the kernel device name. Instead, it allocates a new device name. For example, if you hot swap `/dev/sda`, then the disk corresponding `/dev/disk/by_hba_slot/s0` might link to `/dev/sdn` instead of `/dev/sda`. Therefore, the links in `/dev/disk/by_hba_slot/` are automatically updated when devices are added or removed.

The command output lists device names as kernel device names instead of symbolic link names. Thus, `/dev/disk/by_hba_slot/s0` might be identified as `/dev/sda` in the output of a command.

12.2.1 Standard Disk Drive Mappings

[Table 12-1](#) (page 12-3) shows typical initial mappings between the RAID logical drives and the operating system identifiers. Nonetheless, you must use the mappings that exist for your system, which might be different from the ones listed here. The table also identifies the dedicated function of each drive in an Oracle Big Data Appliance server. The server with the failed drive is part of either a CDH cluster (HDFS) or an Oracle NoSQL Database cluster.

Table 12-1 Disk Drive Identifiers

Symbolic Link to Physical Slot	Typical Initial Kernel Device Name	Dedicated Function
/dev/disk/by-hba-slot/s0	/dev/sda	Operating system
/dev/disk/by-hba-slot/s1	/dev/sdb	Operating system
/dev/disk/by-hba-slot/s2	/dev/sdc	HDFS or Oracle NoSQL Database
/dev/disk/by-hba-slot/s3	/dev/sdd	HDFS or Oracle NoSQL Database
/dev/disk/by-hba-slot/s4	/dev/sde	HDFS or Oracle NoSQL Database
/dev/disk/by-hba-slot/s5	/dev/sdf	HDFS or Oracle NoSQL Database
/dev/disk/by-hba-slot/s6	/dev/sdg	HDFS or Oracle NoSQL Database
/dev/disk/by-hba-slot/s7	/dev/sdh	HDFS or Oracle NoSQL Database
/dev/disk/by-hba-slot/s8	/dev/sdi	HDFS or Oracle NoSQL Database
/dev/disk/by-hba-slot/s9	/dev/sdj	HDFS or Oracle NoSQL Database
/dev/disk/by-hba-slot/s10	/dev/sdk	HDFS or Oracle NoSQL Database
/dev/disk/by-hba-slot/s11	/dev/sdl	HDFS or Oracle NoSQL Database

12.2.2 Standard Mount Points

[Table 12-2](#) (page 12-3) show the mappings between HDFS partitions and mount points.

Table 12-2 Mount Points

Symbolic Link to Physical Slot and Partition	HDFS Partition	Mount Point
/dev/disk/by-hba-slot/s0p4	/dev/sda4	/u01
/dev/disk/by-hba-slot/s1p4	/dev/sdb4	/u02
/dev/disk/by-hba-slot/s2p1	/dev/sdc1	/u03
/dev/disk/by-hba-slot/s3p1	/dev/sdd1	/u04
/dev/disk/by-hba-slot/s4p1	/dev/sde1	/u05
/dev/disk/by-hba-slot/s5p1	/dev/sdf1	/u06
/dev/disk/by-hba-slot/s6p1	/dev/sdg1	/u07
/dev/disk/by-hba-slot/s7p1	/dev/sdh1	/u08
/dev/disk/by-hba-slot/s8p1	/dev/sdi1	/u09
/dev/disk/by-hba-slot/s9p1	/dev/sdj1	/u10
/dev/disk/by-hba-slot/s10p1	/dev/sdk1	/u11
/dev/disk/by-hba-slot/s11p1	/dev/sdl1	/u12

12.2.3 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](#) (page 12-7)."

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

12.3 Prerequisites for Replacing a Failing Disk

To replace an HDFS disk or an operating system disk that is in a state of predictive failure, you must first dismount the HDFS partitions. You must also turn off swapping before replacing an operating system disk.

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 HDFS partitions:

1. Log in to the server with the failing drive.
2. If the failing drive supported the operating system, then turn off swapping:

```
# bdaswapoff
```

Removing a disk with active swapping crashes the kernel.

3. List the mounted HDFS partitions:

```
# mount -l

/dev/md2 on / type ext4 (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 ext4 (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 failing disk. If the disk has no partitions listed, then proceed to ["Replacing a Disk Drive"](#) (page 12-7). 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 /u11` removes the mount point for partition `/dev/sdk1`.

If the `umount` commands succeed, then proceed to "Replacing a Disk Drive (page 12-7)." 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://bda1node03.example.com:7180`

7. Complete these steps in Cloudera Manager:



Note:

If you remove mount points in Cloudera Manager as described in the following steps, then you must restore these mount points in Cloudera Manager after finishing all other configuration procedures.

- a. Log in as `admin`.
 - b. On the Services page, click **hdfs**.
 - c. Click the **Instances** subtab.
 - d. In the Host column, locate the server with the failing disk. Then click the service in the Name column, such as **datanode**, 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 this DataNode**.
8. In Cloudera Manager, remove the mount point from NodeManager Local Directories:
 - a. On the Services page, click **Yarn**.
 - b. In the Status Summary, click **NodeManager**.
 - c. From the list, click to select the NodeManager that is on the host with the failed disk.
 - d. Click the **Configuration** subtab.
 - e. Remove the mount point from the NodeManager.
 - f. Click **Save Changes**.
 - g. Restart the NodeManager.
 9. If you have added any other roles that store data on the same HDFS mount point (such as HBase Region Server), then remove and restore the mount points for these roles in the same way.
 10. Return to your session on the server with the failed drive.
 11. Reissue the `umount` command:

```
# umount mountpoint
```

If the `umount` still fails, run `lsdf` to list open files under the HDFS mount point and the processes that opened them. This may help you to identify the process that is preventing the unmount. For example:

```
# lsof | grep /u11
```

12. Bring the disk offline:

```
# MegaCli64 PDoffline "physdrv[enclosure:slot]" a0
```

For example, "physdrv[20:10]" identifies disk s11, which is located in slot 10 of enclosure 20.

13. Delete the disk from the controller configuration table:

```
MegaCli64 CfgLDDel Lslot a0
```

For example, L10 identifies slot 10.

14. Complete the steps in "[Replacing a Disk Drive](#) (page 12-7)."

12.4 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/>

12.5 Replacing a Disk Drive

Complete this procedure to replace a failed or failing disk drives.

1. Before replacing a failing disk, see "[Prerequisites for Replacing a Failing Disk](#) (page 12-4)."
2. Replace the failed disk drive.
See "[Parts for Oracle Big Data Appliance Servers](#) (page B-5)".
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 (page 12-7) in a text editor and search for the following:

- Disks that have a Foreign State of Foreign
- 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.

- c. Check the current status of Foreign State again:

```
# MegaCli64 pdlist a0 | grep foreign
```

- d. If the Foreign State is still Foreign, then repeat the clear command:

```
# MegaCli64 CfgForeign clear a0
```

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. If the CfgLdAdd command in Step 9 (page 12-8) fails because of cached data, then clear the cache:

```
# MegaCli64 discardpreservedcache l1 a0
```

11. Verify that the disk is recognized by the operating system:

```
# ls SCSI
```

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.

The `lsscsi` output might not show the correct order, but you can continue with the configuration. While the same physical to logical disk mapping is required, the same disk to device mapping for the kernel is not required. The disk configuration is based on `/dev/disks/by-hba-slot` device names.

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 ORACLE  CONCORD14      0960  -
[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   ORACLE  UNIGEN-UFD      PMAP  /dev/sdm
[
```

12. Check the hardware profile of the server, and correct any errors:

```
# bdacheckhw
```

13. Check the software profile of the server, and correct any errors:

```
# bdachecksw
```

If you see a "Wrong mounted partitions" error and the device is missing from the list, then you can ignore the error and continue. However, if you see a "Duplicate mount points" error or the slot numbers are switched, then see ["Correcting a Mounted Partitions Error"](#) (page 12-9).

14. Identify the function of the drive, so you configure it properly. See ["Identifying the Function of a Disk Drive"](#) (page 12-10).

12.6 Correcting a Mounted Partitions Error

When the `bdachecksw` utility finds a problem, it typically concerns the mounted partitions.

An old mount point might appear in the `mount` command output, so that the same mount point, such as `/u03`, appears twice.

To fix duplicate mount points:

1. Dismount both mount points by using the `umount` command twice. This example dismounts two instances of `/u03`:

```
# umount /u03
# umount /u03
```

2. Remount the mount point. This example remounts `/u03`:

```
# mount /u03
```

If a disk is in the wrong slot (that is, the virtual drive number), then you can switch two drives.

To switch slots:

1. Remove the mappings for both drives. This example removes the drives from slots 4 and 10:

```
# MegaCli64 cfgLdDel L4 a0
# MegaCli64 cfgLdDel L10 a0
```

2. Add the drives in the order you want them to appear; the first command obtains the first available slot number:

```
# MegaCli64 cfgLdAdd [20:4] a0
# MegaCli64 cfgLdAdd [20:5] a0
```

3. If mount errors persist even when the slot numbers are correct, then you can restart the server.

12.7 Identifying the Function of a Disk Drive

The server with the failed disk is configured to support either HDFS or Oracle NoSQL Database, and most disks are dedicated to that purpose. However, two disks are dedicated to the operating system. Before configuring the new disk, find out how the failed disk was configured.

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:

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

```
# ls SCSI
```

See the output from Step 11 (page 12-8) of "[Replacing a Disk Drive](#) (page 12-7)".

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

```
# mdadm -Q --detail /dev/md2
/dev/md2:
  Version : 0.90
  Creation Time : Mon Jul 22 22:56:19 2013
  Raid Level : raid1
  .
  .
  .
  Number   Major   Minor   RaidDevice State
    0         8        2         0   active sync  /dev/sda2
    1         8       18         1   active sync  /dev/sdb2
```

3. If the previous steps indicate that the failed disk is an operating system disk, then proceed to "[Configuring an Operating System Disk](#) (page 12-11)".

12.8 Configuring an Operating System Disk

The first two disks 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 the disk in either slot 0 or slot 1.

- [Partitioning the Operating System Disk](#) (page 12-11)
- [Repairing the RAID Arrays](#) (page 12-14)
- [Formatting the HDFS Partition of an Operating System Disk](#) (page 12-16)
- [Restoring the Swap Partition](#) (page 12-16)
- [Restoring the GRUB Master Boot Records and HBA Boot Order](#) (page 12-17)

12.8.1 Partitioning the Operating System Disk

The partitioning procedure differs slightly between Oracle Linux versions 5 and 6. Follow the appropriate procedure for your system:

- [Partitioning for Oracle Linux 6](#) (page 12-11)
- [Partitioning for Oracle Linux 5](#) (page 12-13)

 Note:

Replace `/dev/disk/by-hba-slot/sn` in the following commands with the appropriate symbolic link, either `/dev/disk/by-hba-slot/s0` or `/dev/disk/by-hba-slot/s1`.

12.8.1.1 Partitioning for Oracle Linux 6

To partition a logical drive for Oracle Linux 6:

1. Complete the steps in "[Replacing a Disk Drive](#) (page 12-7)".
2. Confirm that the new disk does not have a partition table:

```
# parted /dev/disk/by-hba-slot/sn -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/disk/by-hba-slot/sn 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/disk/by-hba-slot/sn -s mklabel gpt print
```

5. List the Cylinder, Head, Sector (CHS) partition information of the *surviving* disk. Thus, if you are partitioning /dev/disk/by-hba-slot/s0, then enter /dev/disk/by-hba-slot/s1 for /dev/disk/by-hba-slot/sm in the following command:

```
# parted /dev/disk/by-hba-slot/sm -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	ext4	primary	boot
2	25,127,8	21697,116,20		primary	raid
3	21697,116,21	23227,61,35	linux-swap(v1)	primary	
4	23227,61,36	243031,29,36	ext4	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/disk/by-hba-slot/sn -s mkpart primary file_system start end
```

Use the start and end addresses that you obtained in Step 5 (page 12-12) instead of the addresses shown in the following example:

```
# parted /dev/disk/by-hba-slot/sn -s mkpart primary ext4 0,0,34 25,127,7
# parted /dev/disk/by-hba-slot/sn -s mkpart primary ext4 25,127,8 21697,116,20
# parted /dev/disk/by-hba-slot/sn -s mkpart primary linux-swap 21697,116,21 23227,61,35
```

7. Create primary partition 4 using the start address obtained in Step 5 (page 12-12) and an end address of 100%:

```
# parted /dev/disk/by-hba-slot/sn -s mkpart primary ext4 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/disk/by-hba-slot/sn set 1 raid
# parted -s /dev/disk/by-hba-slot/sn set 2 raid
```

9. Set the boot flag:

```
# parted -s /dev/disk/by-hba-slot/sn set 1 boot
```

10. Verify that the boot flag is set:

```
# parted /dev/disk/by-hba-slot/sn -s unit chs print
```

An example of the output is as follows:

```

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	0,32,32	24,89,0	ext4	primary	boot
2	24,89,1	60812,135,58		primary	raid
3	60812,135,59	65618,90,7	linux-swap(v1)	primary	
4	65618,90,8	243030,252,37	ext4	primary	

11. Complete the steps in "[Repairing the RAID Arrays](#) (page 12-14)".

12.8.1.2 Partitioning for Oracle Linux 5

To partition a logical drive for Oracle Linux 5:

1. Complete the steps in "[Replacing a Disk Drive](#) (page 12-7)".
2. Confirm that the new disk does not have a partition table:

```
# parted /dev/disk/by-hba-slot/sn -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/disk/by-hba-slot/sn 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/disk/by-hba-slot/sn -s mklabel gpt print
```

5. List the Cylinder, Head, Sector (CHS) partition information of the *surviving* disk. Thus, if you are partitioning /dev/disk/by-hba-slot/s0, then enter /dev/disk/by-hba-slot/s1 for /dev/disk/by-hba-slot/sm in the following command:

```
# parted /dev/disk/by-hba-slot/sm -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/disk/by-hba-slot/sn -s mkpart file_system start end
```

Use the start and end addresses that you obtained in Step 5 (page 12-13) instead of the addresses shown in the following example:

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

7. Create primary partition 4 using the start address obtained in Step 5 (page 12-13) and an end address of 100%:

```
# parted /dev/disk/by-hba-slot/sn -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/disk/by-hba-slot/sn set 1 raid
# parted -s /dev/disk/by-hba-slot/sn set 2 raid
```

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

```
# parted /dev/disk/by-hba-slot/sn

GNU Parted 1.8.1
Using /dev/sdn
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](#) (page 12-14)".

12.8.2 Repairing the RAID Arrays

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

- /dev/md0 contains /dev/disk/by-hba-slot/s0p1 and /dev/disk/by-hba-slot/s1p1. It is mounted as /boot.
- /dev/md2 contains /dev/disk/by-hba-slot/s0p2 and /dev/disk/by-hba-slot/s1p2. 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. Remove the partitions from the RAID arrays:

```
# mdadm /dev/md0 -r detached
# mdadm /dev/md2 -r detached
```

2. Verify that the RAID arrays are degraded:

```
# mdadm -Q --detail /dev/md0
# mdadm -Q --detail /dev/md2
```

3. 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
```

4. Restore the partitions to the RAID arrays:

```
# mdadm --add /dev/md0 /dev/disk/by-hba-slot/snp1
# mdadm --add /dev/md2 /dev/disk/by-hba-slot/snp2
```

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

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

6. 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>
```

7. Display the content of /etc/mdadm.conf:

```
# 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...
```

8. Compare the output of the following command with the content of /etc/mdadm.conf from Step 7 (page 12-15):

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

9. If the UUIDs in the file are different from UUIDs in the output of the mdadm command:
 - a. Open /etc/mdadm.conf in a text editor.
 - b. Select from ARRAY to the end of the file, and delete the selected lines.
 - c. Copy the output of the command into the file where you deleted the old lines.

- d. Save the modified file and exit.
10. Complete the steps in ["Formatting the HDFS Partition of an Operating System Disk" \(page 12-16\)](#).

12.8.3 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:

```
# mkfs -t ext4 /dev/disk/by-hba-slot/snp4
```

Note:

If this command fails because the device is mounted, then dismount the drive now and skip step 3 (page 12-16). See ["Prerequisites for Replacing a Failing Disk" \(page 12-4\)](#) for dismounting instructions.

2. Verify that the partition label (such as /u01 for s0p4) 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:

```
# tune2fs -c -1 -i 0 -m 0.2 -L /u0n /dev/disk/by-hba-slot/snp4
```
5. Mount the HDFS partition:

```
# mount /u0n
```
6. Complete the steps in ["Restoring the Swap Partition" \(page 12-16\)](#).

12.8.4 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-sdn3 /dev/disk/by-hba-slot/snp3
Setting up swapspace version 1, size = 12582907 kB
LABEL=SWAP-sdn3, no uuid
```

2. Verify that the swap partition is restored:

```
# bdaswapon; bdaswapoff
```

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

3. Verify that the replaced disk is recognized by the operating system:

```
$ ls -l /dev/disk/by-label
total 0
lrwxrwxrwx 1 root root 10 Aug 3 01:22 BDAUSB -> ../../sdn1
lrwxrwxrwx 1 root root 10 Aug 3 01:22 BDAUSBB00T -> ../../sdm1
lrwxrwxrwx 1 root root 10 Aug 3 01:22 SWAP-sda3 -> ../../sda3
lrwxrwxrwx 1 root root 10 Aug 3 01:22 SWAP-sdb3 -> ../../sdb3
lrwxrwxrwx 1 root root 10 Aug 3 01:22 u01 -> ../../sda4
lrwxrwxrwx 1 root root 10 Aug 3 01:22 u02 -> ../../sdb4
lrwxrwxrwx 1 root root 10 Aug 3 01:22 u03 -> ../../sdc1
lrwxrwxrwx 1 root root 10 Aug 3 01:22 u04 -> ../../sdd1
.
```

4. If the output does not list the replaced disk:

- On Linux 5, run `udevtrigger`.
- On Linux 6, run `udevadm trigger`.

Then repeat step 3 (page 12-17). The `lsscsi` command should also report the correct order of the disks.

5. Complete the steps in "[Restoring the GRUB Master Boot Records and HBA Boot Order](#) (page 12-17)".

12.8.5 Restoring the GRUB Master Boot Records and HBA Boot Order

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

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
```

However, the GRUB device map does not support symbolic links, and the mappings in the device map might not correspond to those used by `/dev/disk/by-hba-slot`. The following procedure explains how you can correct the device map if necessary.

To restore the GRUB boot record:

1. Check which kernel device the drive is using in slot1

```
# ls -ld /dev/disk/by-hba-slot/s1
lrwxrwxrwx 1 root root 9 Apr 22 12:54 /dev/disk/by-hba-slot/s1 -> ../../sdb
```

2. If the output displays `/dev/sdb` as shown in step 1 (page 12-17), then proceed to the next step (open GRUB).

If another device is displayed, such as `/dev/sdn`, then you must first set `hd1` to point to the correct device:

- a. Make a copy of the `device.map` file:

```
# cd /boot/grub
# cp device.map mydevice.map
# ls -l *device*
-rw-r--r-- 1 root root 85 Apr 22 14:50 device.map
```

```
-rw-r--r-- 1 root root 85 Apr 24 09:24 mydevice.map
```

- b.** Edit mydevice.map to point hd1 to the new device. In this example, s1 pointed to /dev/sdn in step 1 (page 12-17).

```
# more /boot/grub/mydevice.map
# this device map was generated by bda install
(hd0)      /dev/sda
(hd1)      /dev/sdn
```

- c.** Use the edited device map (mydevice.map) in the remaining steps.

- 3.** Open GRUB, using either device.map as shown, or the edited mydevice.map:

```
# 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.
]
```

- 4.** 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
```

- 5.** 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.
```

- 6.** Close the GRUB command-line interface:

```
grub> quit
```

- 7.** Ensure that the boot drive in the HBA is set correctly:

```
# MegaCli64 /c0 show bootdrive
```

If BootDrive VD:0 is set, the command output is as follows:

```
Controller = 0
Status = Success
Description = None
Controller Properties :
=====
-----
Ctrl_Prop Value
-----
```

```
BootDrive VD:0
-----
```

If BootDrive VD:0 is not set, the command output shows No Boot Drive:

```
Controller = 0
Status = Success
Description = None
Controller Properties :
=====
-----
Ctrl_Prop Value
-----
BootDrive No Boot Drive
-----
```

8. If MegaCli64 /c0 show bootdrive reports that the boot drive is not set, then set it as follows:

```
# MegaCli64 /c0/v0 set bootdrive=on
Controller = 0
Status = Success
Description = None
Detailed Status :
=====
-----
VD  Property  Value Status  ErrCd ErrMsg
-----
0   Boot Drive On    Success  0      -
-----
```

9. Verify that the boot drive is now set:

```
# MegaCli64 /c0 show bootdrive
Controller = 0
Status = Success
Description = None
Controller Properties :
=====
-----
Ctrl_Prop Value
-----
BootDrive VD:0
-----
```

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

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

11. Check the configuration. See "[Verifying the Disk Configuration](#) (page 12-21)".

12.9 Configuring an HDFS or Oracle NoSQL Database Disk

Complete the following instructions for any disk that is not used by the operating system. See "[Identifying the Function of a Disk Drive](#) (page 12-10)".

To configure a disk, you must partition and format it.



Note:

Replace *snp1* in the following commands with the appropriate symbolic name, such as *s4p1*.

To format a disk for use by HDFS or Oracle NoSQL Database:

1. Complete the steps in "[Replacing a Disk Drive](#) (page 12-7)", if you have not done so already.

2. Partition the drive:

```
# parted /dev/disk/by-hba-slot/sn -s mklabel gpt mkpart primary ext4 0% 100%
```

3. Format the partition for an ext4 file system:

```
# mkfs -t ext4 /dev/disk/by-hba-slot/snp1
```

4. Reset the appropriate partition label to the missing device. See [Table 12-2](#) (page 12-3).

```
# tune2fs -c -1 -i 0 -m 0.2 -L /unn /dev/disk/by-hba-slot/snp1
```

For example, this command resets the label for `/dev/disk/by-hba-slot/s2p1` to `/u03`:

```
# tune2fs -c -1 -i 0 -m 0.2 -L /u03 /dev/disk/by-hba-slot/s2p1
Setting maximal mount count to -1
Setting interval between checks to 0 seconds
Setting reserved blocks percentage to 0.2% (976073 blocks)
```

5. Verify that the replaced disk is recognized by the operating system:

```
$ ls -l /dev/disk/by-label
total 0
lrwxrwxrwx 1 root root 10 Aug 3 01:22 BDAUSB -> ../../sdn1
lrwxrwxrwx 1 root root 10 Aug 3 01:22 BDAUSBB00T -> ../../sdm1
lrwxrwxrwx 1 root root 10 Aug 3 01:22 SWAP-sda3 -> ../../sda3
lrwxrwxrwx 1 root root 10 Aug 3 01:22 SWAP-sdb3 -> ../../sdb3
lrwxrwxrwx 1 root root 10 Aug 3 01:22 u01 -> ../../sda4
lrwxrwxrwx 1 root root 10 Aug 3 01:22 u02 -> ../../sdb4
lrwxrwxrwx 1 root root 10 Aug 3 01:22 u03 -> ../../sdc1
lrwxrwxrwx 1 root root 10 Aug 3 01:22 u04 -> ../../sdd1
.
```

6. If the output does not list the replaced disk:

- On Linux 5, run `udevtrigger`.
- On Linux 6, run `udevadm trigger`.

Then repeat step 5 (page 12-20). The `ls SCSI` command should also report the correct order of the disks.

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

```
# mount /unn
```

For example, mount `/u03`.

8. If you are configuring multiple drives, then repeat the previous steps.
9. If you previously removed a mount point in Cloudera Manager for an HDFS drive, then restore it to the list.
 - a. Open a browser window to Cloudera Manager. For example:
`http://bda1node03.example.com:7180`
 - b. Open Cloudera Manager and log in as admin.
 - c. On the Services page, click **hdfs**.
 - d. Click the **Instances** subtab.
 - e. In the Host column, locate the server with the replaced disk. Then click the service in the Name column, such as **datanode**, to open its page.
 - f. Click the **Configuration** subtab.
 - g. If the mount point is missing from the Directory field, then add it to the list.
 - h. Click **Save Changes**.
 - i. From the Actions list, choose **Restart**.
10. If you previously removed a mount point from NodeManager Local Directories, then also restore it to the list using Cloudera Manager.
 - a. On the Services page, click **Yarn**.
 - b. In the Status Summary, click **NodeManager**.
 - c. From the list, click to select the NodeManager that is on the host with the failed disk.
 - d. Click the Configuration sub-tab.
 - e. If the mount point is missing from the NodeManager Local Directories field, then add it to the list.
 - f. Click **Save Changes**.
 - g. From the Actions list, choose **Restart**.
11. Check the configuration. See "[Verifying the Disk Configuration](#) (page 12-21)" .

12.10 Verifying the Disk Configuration

Before you can reinstall the Oracle Big Data Appliance software on the server, you must verify that the configuration is correct on the new disk drive.

To verify the disk configuration:

1. Check the software configuration:

```
# bdachecksw
```
2. If there are errors, then redo the configuration steps as necessary to correct the problem.
3. Check the `/root` directory for a file named `BDA_REBOOT_SUCCEEDED`.
4. If you find a file named `BDA_REBOOT_FAILED`, then read the file to identify and fix any additional problems.
5. Use this script to generate a `BDA_REBOOT_SUCCEEDED` file:

```
# /opt/oracle/bda/lib/bdastartup.sh
```

6. Verify that BDA_REBOOT_SUCCEEDED exists. If you still find a BDA_REBOOT_FAILED file, then redo the previous steps.

13

Oracle Big Data Appliance Utilities

This chapter describes the utilities available on Oracle Big Data Appliance. Most of the utilities are for monitoring the health of the hardware and the network.

- [bdcheckcluster](#) (page 13-2)
- [bdcheckhw](#) (page 13-3)
- [bdcheckib](#) (page 13-3)
- [bdchecknet](#) (page 13-5)
- [bdchecksw](#) (page 13-5)
- [bdclustersynctime](#) (page 13-6)
- [bdcli](#) (page 13-7)
- [bdclustersynctime](#) (page 13-6)
- [bdadiag](#) (page 13-17)
- [bdadiagcluster](#) (page 13-19)
- [bdagetclientconfig](#) (page 13-20)
- [bdaibstatus](#) (page 13-21)
- [bdaid](#) (page 13-21)
- [bdaimagevalidate](#) (page 13-22)
- [bdanodedesc](#) (page 13-23)
- [bdaredoclientnet](#) (page 13-23)
- [bdaserials](#) (page 13-25)
- [bdaswapoff](#) (page 13-26)
- [bdaswapon](#) (page 13-26)
- [bdaupdatefw](#) (page 13-27)
- [iblinkinfo](#) (page 13-28)
- [imagehistory](#) (page 13-29)
- [imageinfo](#) (page 13-29)
- [listlinkup](#) (page 13-30)
- [mount_hadoop_dirs](#) (page 13-30)
- [remove-root-ssh](#) (page 13-30)
- [reset-boot-order](#) (page 13-31)
- [setup-root-ssh](#) (page 13-32)
- [showusb](#) (page 13-33)
- [showvlan](#) (page 13-34)

- [showvnics](#) (page 13-34)

13.1.1 bdacheckcluster

Checks the health of a CDH cluster, including the software, hardware, and network, and logs the results in a file in the /tmp directory.

To also run the CDH cluster health checks, enter the Cloudera Manager admin password in response to the prompt. If you do not enter the password, then these checks are skipped.

Syntax

```
bdacheckcluster
```

Notes

If the cluster is protected by Kerberos authentication, then you must obtain a ticket for the hdfs user before running bdacheckcluster.

Example

This example shows the output from the utility:

```
# bdacheckcluster
INFO: Logging results to /tmp/bdacheckcluster_1373393815/
Enter CM admin password to enable check for CM services and hosts
Press ENTER twice to skip CM services and hosts checks
Enter password: password
Enter password again: password
SUCCESS: Mammoth configuration file is valid.
SUCCESS: hdfs is in good health
SUCCESS: mapreduce is in good health
SUCCESS: oozie is in good health
SUCCESS: zookeeper is in good health
SUCCESS: hive is in good health
SUCCESS: hue is in good health
SUCCESS: Cluster passed checks on all hadoop services health check
SUCCESS: bda1node01.example.com is in good health
SUCCESS: bda1node02.example.com is in good health
SUCCESS: bda1node03.example.com is in good health
.
.
.
SUCCESS: Cluster passed checks on all hosts health check
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
SUCCESS: All cluster hosts resolve public hostnames to private IPs
.
.
.
```

```

INFO: Checking local reverse DNS resolve of private IPs on all cluster hosts
SUCCESS: All cluster hosts resolve private IPs to public hostnames
SUCCESS: 2 virtual NICs available on all cluster hosts
SUCCESS: NTP service running on all cluster hosts
SUCCESS: At least one valid NTP server accessible from all cluster servers.
SUCCESS: Max clock drift of 0 seconds is within limits
SUCCESS: Big Data Appliance cluster health checks succeeded

```

13.1.2 bdacheckhw

Checks the hardware profile of the server.

Syntax

```
bdacheckhw
```

Usage Notes

See "[Configuring the Oracle Big Data Appliance Servers](#) (page 7-20)" for tips about using this utility.

Example

This example shows the output from the utility:

```

# bdacheckhw
SUCCESS: Found BDA v2 server : SUN FIRE X4270 M3
SUCCESS: Correct processor info : Intel(R) Xeon(R) CPU E5-2660 0 @ 2.20GHz
SUCCESS: Correct number of types of CPU : 1
SUCCESS: Correct number of CPU cores : 32
SUCCESS: Sufficient GB of memory (>=63): 63
SUCCESS: Correct BIOS vendor : American Megatrends Inc.
SUCCESS: Sufficient BIOS version (>=08080102): 18021300
SUCCESS: Recent enough BIOS release date (>=05/23/2011):06/19/2012
SUCCESS: Correct ILOM major version : 3.1.2.12
SUCCESS: Sufficient ILOM minor version (>=74388): 74388
SUCCESS: Correct number of fans : 4
SUCCESS: Correct fan 0 status : ok
SUCCESS: Correct fan 1 status : ok
.
.
.
SUCCESS: Big Data Appliance hardware validation checks succeeded

```

13.1.3 bdacheckib

Checks the InfiniBand cabling between the servers and switches of a single rack, when entered with no options.

Run this command after connecting as root to any server.

Syntax

```
bdacheckib [parameter]
```

Parameters

-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 racks 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.

Usage Notes

The network must be configured with custom settings as described by `/opt/oracle/bda/network.json`.

Examples

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

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

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

```
[root@bda1node01 bda]# bdacheckib -g
[root@bda1node01 bda]# cat sample-multi-rack.json
# This json multirack spec is generated. The array elements are sorted
# alphabetically. A 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": "bda1sw-ib1", "LEAF1_NAME": "bda1sw-ib2", "LEAF2_NAME": "bda1sw-ib3"},
{"SPINE_NAME": "bda2sw-ib1", "LEAF1_NAME": "bda2sw-ib2", "LEAF2_NAME": "bda2sw-ib3"},
{"SPINE_NAME": "dm01sw-ib1", "LEAF1_NAME": "dm01sw-ib2", "LEAF2_NAME": "dm01sw-ib3"}
]
```

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
```

```
Rack #1 leaf to spines topology check
leaf: bda1sw-ib2
    expected 2 links to rack 1, found 4 OK
    expected 2 links to rack 2, found 4 OK
    expected 2 links to rack 3, found 3 OK
    expected 2 links to rack 4, found 3 OK
leaf: bda1sw-ib3
    expected 2 links to rack 1, found 3 OK
    expected 2 links to rack 2, found 4 OK
```

```

        expected 2 links to rack 3, found 3 OK
        expected 2 links to rack 4, found 3 OK
        .
        .
        .

Rack #1 cabling details
leaf: bda1sw-ib2
LINK ... to rack2 ..... UP
LINK ... to rack2 ..... UP
LINK ... to rack1 ..... UP
LINK ... to rack1 ..... UP
LINK ... to rack3 ..... UP
LINK ... to rack3 ..... UP
LINK ... to rack4 ..... UP
LINK ... to rack4 ..... UP
.
.
.

```

13.1.4 bdachecknet

Checks whether the network configuration is working properly.

Syntax

```
bdachecknet
```

Usage Notes

Run this command after connecting as root to any server.

Example

This example shows the output from the utility:

```

[root@node01 network]# bdachecknet
bdachecknet: do basic sanity checks on /opt/oracle/bda/network.json
bdachecknet: warning: bda deploy /opt/oracle/bda/network.json check odd nr of
elements 7
bdachecknet: passed
bdachecknet: checking for networkexpansion.json
bdachecknet: ping test private infiniband ips (bondib0 40gbs)
.
.
.

```

13.1.5 bdachecksw

Checks the software profile of a server.

Syntax

```
bdachecksw
```

Usage Notes

See "[Configuring the Oracle Big Data Appliance Servers](#) (page 7-20)" for tips about using this utility.

Example

This example shows the output from the utility:

```
# bdachecksw
SUCCESS: Correct OS disk s0 partition info : 1 ext3 raid 2 ext3 raid 3 linux-swap 4
ext3 primary
SUCCESS: Correct OS disk s1 partition info : 1 ext3 raid 2 ext3 raid 3 linux-swap 4
ext3 primary
SUCCESS: Correct data disk s2 partition info : 1 ext3 primary
SUCCESS: Correct data disk s3 partition info : 1 ext3 primary
SUCCESS: Correct data disk s4 partition info : 1 ext3 primary
SUCCESS: Correct data disk s5 partition info : 1 ext3 primary
SUCCESS: Correct data disk s6 partition info : 1 ext3 primary
SUCCESS: Correct data disk s7 partition info : 1 ext3 primary
SUCCESS: Correct data disk s8 partition info : 1 ext3 primary
SUCCESS: Correct data disk s9 partition info : 1 ext3 primary
SUCCESS: Correct data disk s10 partition info : 1 primary
SUCCESS: Correct data disk s11 partition info : 1 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/mapper/lvg1-lv1 /lv1 ext4 /dev/md0 /boot
ext3 /dev/md2 / ext3 /dev/sd4 /u01 ext4 /dev/sd4 /u02 ext4 /dev/sd1 /u03 ext4 /dev/
sd1 /u04 ext4 /dev/sd1 /u05 ext4 /dev/sd1 /u06 ext4 /dev/sd1 /u07 ext4 /dev/sd1 /u08
ext4 /dev/sd1 /u09 ext4 /dev/sd1 /u10 ext4
SUCCESS: Correct matching label and slot : symbolic link to `../../sda4'
SUCCESS: Correct matching label and slot : symbolic link to `../../sdb4'
.
.
.
SUCCESS: Correct Linux kernel version 2.6-39-400.200 series : Linux
2.6.39-400.2 14.3.el6uek.x86_64
SUCCESS: Correct Java Virtual Machine version : HotSpot(TM) 64-Bit Server
1.7.0_25
SUCCESS: Correct puppet version : 2.6.11
SUCCESS: Correct MySQL version : 5.5.35
SUCCESS: All required programs are accessible in $PATH
SUCCESS: All required RPMs are installed and valid
SUCCESS: Correct bda-monitor status : bda monitor is running
SUCCESS: Big Data Appliance software validation checks succeeded
```

13.1.6 bdaclustersyncntime

Synchronizes the time of all servers in a cluster.

Syntax

bdaclustersyncntime

Usage Notes

To use this utility, you must log in as root to the first server in the node. Passwordless ssh must also be set up for the cluster. See the -c parameter for "[setup-root-ssh](#) (page 13-32)."

This utility creates a log file named bdaclustersyncntime.log in the directory identified in the output.

Example

The following example successfully runs `bdacclustersynctime`:

```
# bdacclustersynctime
INFO: Logging results to /tmp/bdaccluster_1373485952/
SUCCESS: Mammoth configuration file is valid.
SUCCESS: All cluster host names are pingable
SUCCESS: NTP service running on all cluster hosts
SUCCESS: At least one valid NTP server found
SUCCESS: No errors found syncing date and time on all nodes
SUCCESS: Max clock drift of 0 seconds is within limits
SUCCESS: Sync date and time of cluster succeeded
```

13.1.7 bdaccli

The Oracle Big Data Appliance Command-Line Interface (`bdaccli`) queries various configuration files to return information about the rack, cluster, server, InfiniBand network, and software patches.

The `bdaccli` utility also adds and removes patches and optional services. It can migrate critical services between critical nodes, and add and remove servers from a cluster.

The `bdaccli` utility displays usage information if no parameters are included on the command line or the values are undefined.

Syntax

`bdaccli action [parameters]`

Actions

help

Displays general usage information for `bdaccli`, a list of actions, and a list of supported parameters for the `getinfo` action.

{add | remove} patch *patch_number*

Adds or removes a software patch on Oracle Big Data Appliance that matches *patch_number*. You must log in as root to use add or remove.

admin_cluster *parameter node_name*

Enables you to administer the nodes in a cluster in response to a failing server. [Table 13-1](#) (page 13-7) describes the parameters.

Parameter	Description
decommissi on	Removes the specified node from the cluster and decommissions the server in Cloudera Manager. It also updates the Mammoth files. You can decommission a failing, noncritical node. Note that critical services on the node must be moved first. This parameter is for temporarily removing a node from a cluster for maintenance and is not a method to delete a node from a cluster.

Parameter	Description
recommissi on	Removes the node from the list of decommissioned nodes, and recommissions the server in Cloudera Manager. Use this command after decommissioning and repairing a failing server.
migrate	Moves the services from a critical node to a noncritical node, and decommissions the failing server in Cloudera Manager. You specify the name of the failing critical node, and the utility selects the noncritical node for the migration. When migration is complete, the new node has all of the functionality of the original critical node. You can only migrate a critical node, and should do so only when it is failing.
reprovisio n	Restores a server to the cluster as a noncritical node, and recommissions the server in Cloudera Manager. Use this command after migrating the services of a critical node and repairing the failing server.

 Note:

In Oracle Big Data Appliance 4.8, migration does not support Oracle Big Data SQL. If Oracle Big Data SQL is installed, disable it first, run the migration, and then re-enable Oracle Big Data SQL.
Use `bdaccli disable big_data_sql` and `bdaccli enable big_data_sql`.

{enable | disable} service

Enables or disables an optional software component on Oracle Big Data Appliance. You must log in as root to use enable or disable.

[Table 13-2](#) (page 13-8) describes the component parameters.

Parameter	Description
asr	Auto Service Request
auditvault	Oracle Audit Vault and Database Firewall plugin
bdc	Oracle Big Data Connectors
big_data_sql	Oracle Big Data SQL
em	Oracle Enterprise Manager Cloud Control agent
disk_encrypt ion	Disable existing eCryptfs encryption. disable only. There is no enable action for this parameter. You have the option to use the <code>hdfs_transparent_encryption</code> parameter to encrypt data at rest.

Parameter	Description
hdfs_transparent_encryption	HDFS Transparent Encryption on the cluster. Respond Yes or No to a prompt to set up Key Trustee Servers on the BDA. If you respond No, you are further prompted for the addresses of the Active and Passive Key Trustee Servers, as well as the Key Trustee organization and authorization code. If you respond Yes, the Key Trustee Servers are set up on Oracle Big Data Appliance in a predetermined configuration. Kerberos is required. See Document 2111343.1 in My Oracle Support for important prerequisites and procedure details. After the initial configuration, starting/stopping the KMS/KTS servers can be done only through Cloudera Manager. Mammoth does not have direct access to edge nodes and therefore cannot set up or detect KMS services on these nodes. Customers can set up KMS services on edge nodes themselves. These must exist same cluster as the services that require access to the Key Trustee Servers. See Cloudera's description of this architecture.
ad_kerberos	Active Directory Kerberos. Also prompts the user to choose whether or not to enable the features <i>Encrypt Hadoop Services</i> and <i>Encrypt HDFS Data Transport</i>. The disable option disables these services as well as ad_kerberos. Note that there are important preliminary steps. See MOS (My Oracle Support) documents 2013585.1 and 2029378.1.
kerberos	MIT Kerberos. Also prompts the user to choose whether or not to enable the features <i>Encrypt Hadoop Services</i> and <i>Encrypt HDFS Data Transport</i>. The disable option disables these services as well as MIT Kerberos. This parameter does not control Active Directory Kerberos. See ad_kerberos in this table.
https_cm_hue_oozie	HTTPS for Cloudera Manager, Hue, and Oozie. There is no disable option. Kerberos is required.
hdfs_encrypted_data_transport	Encryption of data transfer between DataNodes and clients, and among DataNodes. Kerberos is required.
hadoop_network_encryption	HTTPS encryption and Kerberos authentication for HDFS and YARN, also enables encrypted web shuffling for YARN. Data spills to non-HDFS disk storage from Spark shuffle, MapReduce intermediate files, map and reduce operations, and Impala SQL are also encrypted. Kerberos is required.
sentry	Apache Sentry authorization

**Note:**

If HDFS Transparent Encryption is disabled, data stored in the HDFS Transparent Encryption zones in the cluster will remain encrypted and therefore inaccessible. To restore access to the data, re-enable HDFS Transparent Encryption using the same key provider.

{start | stop | restart | status} service

Starts, stops, restarts, or returns the current status of a service on a cluster or a specific node.

Table 13-3 (page 13-10) describes the service parameters.

Parameter	Description
<code>big_data_sql_cluster</code>	Oracle Big Data SQL on all nodes of the cluster
<code>big_data_sql_server node_name</code>	Oracle Big Data SQL on a specified node of the cluster. Use <code>bdaccli</code> with this parameter only from the first node of the cluster, where the current <code>config.json</code> file is stored.

reset scope_parameter network_parameter

Reconfigures the network (or networks) within the specified scope. The network configuration will set the IP addresses, subnets, domain and DNS. In the case of a virtual adapter such as client network or private network, the VNICs are deleted and recreated with the new IP addresses.

Requires root privileges.

Examples:

```
bdaccli reset server admin_network
bdaccli reset cluster client_network
bdaccli reset cluster private_network
bdaccli reset server all_networks
```

Scope Parameter	Description
<code>server</code>	The action is performed for the node where the command is being executed.
<code>cluster</code>	The action is performed for all the nodes in the cluster.
<code>rack</code>	The action is performed for all the nodes in the rack.

Network Parameter	Description
<code>admin_network</code>	Reconfigures the admin network interface <code>eth0</code> to what is specified in <code><rack-name>-network.json</code> .
<code>client_network</code>	Reconfigures the client <code>eth8</code> , <code>eth9</code> , <code>bondeth0</code> and all the extra client networks to what is specified in <code><cluster-name>-network.json</code> .

Network Parameter	Description
private_network	Reconfigures client ib0, ib1, bondib0 and all the extra private networks to what is specified in <cluster-name>-network.json.
all_networks	Reconfigures all networks (admin, client and private) to what is specified in the corresponding network files.

getinfo [parameter]

Returns a list of getinfo parameters. If you include a parameter name in the command, then getinfo returns information about that system component:

- **InfiniBand parameters:** The bdacli command queries the InfiniBand fabric. See ["InfiniBand Parameters" \(page 13-11\)](#).
- **Rack parameters:** Describes a physical Oracle Big Data Appliance rack. The bdacli command queries the current network.json configuration file for the rack where the command executes. See ["Rack Parameters" \(page 13-11\)](#).
- **Cluster parameters:** Describes a logical Oracle Big Data Appliance cluster. The bdacli command queries the current config.json file for the Hadoop cluster where the command executes. See ["Cluster Parameters" \(page 13-13\)](#).
- **Server parameters:** Describes a server. The bdacli command queries the operating system of the server where the bdacli command executes. See ["Server Parameters" \(page 13-15\)](#).
- **One-off patch parameters:** Provides information about one-off patches. See ["One-Off Patch Parameters" \(page 13-15\)](#).

Parameter names that end with an "s" return lists. Boolean parameters return a string of either true or false.

InfiniBand Parameters

[Table 13-6](#) (page 13-11) describes the InfiniBand parameters for bdacli getinfo.

Table 13-6 InfiniBand Parameters

Parameter	Returns
ib_bda_servers	A list of all Oracle Big Data Appliance servers on the InfiniBand fabric. The list contains the unqualified client host names, which are sorted alphabetically. The list can include servers in multiple racks that are cabled together.
ib_switches	Lists the switches on the InfiniBand fabric, including the switch name, Globally Unique Identifier (GUID), and type (GTW for a gateway switch, or 36P for a spine switch). The list contains the unqualified administration names, which are sorted alphabetically by switch name. See "Examples" (page 13-16) .

Rack Parameters

[Table 13-7](#) (page 13-12) describes the rack parameters for bdacli getinfo.

Table 13-7 Rack Parameters

Parameter	Returns
rack_admin_domain_name	The administrative network domain name for the rack, such as <code>example.com</code> . Because all cluster nodes must have the same administrative domain name, the return value is also the administrative network domain for all nodes in the cluster.
rack_domain_name	The client network domain name for the rack, such as <code>example.com</code> . Because all cluster nodes must have the same client domain name, the return value is also the client network domain for all nodes in the cluster.
rack_ib_switch_ips	The IP addresses of the three InfiniBand switches in the rack in this order: spine switch, first leaf switch (upper), second leaf switch (lower)
rack_ib_switch_names	The unqualified names of the three InfiniBand switches in the rack in this order: spine switch, first leaf switch (upper), second leaf switch (lower). For example, <code>bda1sw-ib1</code> , <code>bda1sw-ib2</code> , and <code>bda1sw-ib3</code> .
rack_ilom_ips	An ordered list of IP addresses of all Oracle ILOMs in the rack, starting with the lowest server
rack_ilom_names	An ordered list of host names of all Oracle ILOMs in the rack, such as <code>bda1node01-ilom</code> and <code>bda1node02-ilom</code> , starting with the lowest server
rack_ip_switch_ip	The IP address of the Cisco Ethernet switch
rack_ip_switch_name	The host name of the Cisco Ethernet switch, such as <code>bda1sw-ip</code>
rack_name	The name of this rack, such as <code>bda1</code>
rack_pdu_ips	An ordered list of IP addresses of PDUs in the rack, starting with PDU-A
rack_pdu_names	An ordered list of unqualified names of PDUs in the rack, starting with PDU-A (such as <code>bda1-pdua</code> and <code>bda1-pdub</code>)
rack_serial_number	The rack serial number
rack_server_admin_ips	An ordered list of IP addresses on the administrative network of all servers in the rack, starting with the lowest server
rack_server_admin_names	An ordered list of host names on the administrative network of all servers in the rack (such as <code>bda1node01-adm</code> , <code>bda1node02-adm</code> , and so forth) starting with the lowest server
rack_server_ips	An ordered list of IP addresses on the client network of all servers in the rack, starting with the lowest server
rack_server_names	An ordered list of host names on the client network of all servers in the rack (such as <code>bda1node01</code> , <code>bda1node02</code> , and so forth) starting with the lowest server
rack_server_priv_ips	A list of IP addresses for ILOM devices in the rack
rack_server_priv_names	The unqualified names of the ILOM devices in the rack (such as <code>bda1node01-priv</code> , <code>bda1node02-priv</code> , and so forth)

Cluster Parameters

The following tables describe the cluster parameters:

- [Table 13-8](#) (page 13-13)
- [Table 13-9](#) (page 13-13)
- [Table 13-10](#) (page 13-14)
- [Table 13-11](#) (page 13-14)

[Table 13-8](#) (page 13-13) describes the general cluster parameters for `bdaccli getinfo`.

Table 13-8 General Cluster Parameters

Parameter	Returns
<code>cluster_asr_installed</code>	true if Auto Service Request is configured for this cluster; false otherwise
<code>cluster_big_data_sql_enabled</code>	true if Oracle Big Data SQL is enabled for this cluster; false otherwise.
<code>cluster_cdh_version</code>	The version of Cloudera's Distribution including Apache Hadoop installed on this cluster, such as 4.5.0-016.
<code>cluster_cm_server</code>	The Cloudera Manager address, including the server name and port number, such as <code>bda1node03.example.com:7180</code> .
<code>cluster_cm_version</code>	The version of Cloudera Manager running on this cluster, such as 4.8.0-016.
<code>cluster_em_enabled</code>	true if Oracle Enterprise Manager is configured for this cluster; false otherwise
<code>cluster_name</code>	The name of the cluster, such as <code>cluster-c</code> .
<code>cluster_nosqlldb_type</code>	The edition of Oracle NoSQL Database installed on this cluster (community edition or enterprise edition), or none if Oracle NoSQL Database is not installed.
<code>cluster_nosqlldb_version</code>	The version of Oracle NoSQL Database installed on the cluster.
<code>cluster_primary_host</code>	The unqualified host name of the server that hosts the puppet master. The Mammoth utility was deployed from this host, and any reconfiguration of the cluster must be done while logged in to that server.
<code>cluster_type</code>	The type of cluster, either <code>hadoop</code> or <code>nosqlldb</code> .
<code>cluster_version</code>	The software version installed on this cluster by the Mammoth utility, such as 3.1.0.

[Table 13-9](#) (page 13-13) describes the cluster parameters related to Oracle Big Data Connectors for `bdaccli getinfo`.

Table 13-9 Oracle Big Data Connectors Status Parameters

Parameter	Returns
<code>cluster_bdc_installed</code>	true if Oracle Big Data Connectors is installed; false otherwise

Table 13-9 (Cont.) Oracle Big Data Connectors Status Parameters

Parameter	Returns
cluster_odi_enabled	true if the Oracle Data Integrator agent is enabled; false otherwise.
cluster_odi_version	The version of Oracle Data Integrator agent installed on the cluster.
cluster_oraah_version	The version of Oracle R Advanced Analytics for Hadoop installed on the cluster
cluster_oraloader_version	The version of Oracle Loader for Hadoop installed on the cluster
cluster_osch_version	The version of Oracle SQL Connector for HDFS installed on the cluster
cluster_oxh_version	The version of Oracle XQuery for Hadoop installed on the cluster

[Table 13-10](#) (page 13-14) describes the cluster network parameters for `bdaccli getinfo`.

Table 13-10 Cluster Network Parameters

Parameter	Returns
cluster_hosts_entries	A list of <code>/etc/hosts</code> entries that you can append to the <code>/etc/hosts</code> file on any device on the same InfiniBand fabric as the Oracle Big Data Appliance cluster, to ensure that Hadoop traffic uses the InfiniBand network. Do not add these entries to a device on a different fabric. Each entry is on a separate line with three parts: the InfiniBand IP address, the full client host name, and the short client host name.
cluster_ilom_ips	An ordered list of IP addresses for the Oracle ILOMs in the servers, starting with the first node in the cluster
cluster_ilom_names	A list of unqualified host names on the administrative network for the Oracle ILOMs in the servers, in order starting with the first server in the cluster
cluster_node_ips	The IP addresses on the client network of all nodes in this cluster
cluster_node_names	The host names on the client network of all nodes in the cluster, such as <code>bda1node01</code>

[Table 13-11](#) (page 13-14) describes the cluster security parameters for `bdaccli getinfo`.

Table 13-11 Cluster Security Parameters

Parameter	Returns
cluster_av_admin	The name of the Oracle Audit Vault and Database Firewall administration user. Returns an error if Audit Vault is not configured for this cluster.

Table 13-11 (Cont.) Cluster Security Parameters

Parameter	Returns
cluster_av_enabled	true if Oracle Audit Vault and Database Firewall auditing is enabled; false otherwise
cluster_av_port	The port number that the Audit Vault server listens on. Returns an error if Oracle Audit Vault and Database Firewall is not configured on this cluster.
cluster_av_server	The IP address of the Audit Vault server. Returns an error if Oracle Audit Vault and Database Firewall is not configured on this cluster.
cluster_av_service	The database service name for the Audit Vault server. Returns an error if Oracle Audit Vault and Database Firewall is not configured on this cluster.
cluster_hdfs_encryption_enabled	true if network encryption of Hadoop data is enabled for this cluster; false otherwise
cluster_hdfs_transparent_encryption_enabled	true if HDFS Transparent Encryption of Hadoop data at rest is enabled for this cluster; false otherwise
cluster_kerberos_enabled	true if Kerberos security is enabled; false otherwise.
cluster_kerberos_kdc_hosts	A list of key distribution center (KDC) hosts external to Oracle Big Data Appliance. Returns an error if Kerberos is not enabled.
cluster_kerberos_kdc_on_bda	true if the Kerberos KDC is on Oracle Big Data Appliance; false otherwise. Returns an error if Kerberos is not enabled.
cluster_kerberos_realm	The Kerberos realm for the cluster. Returns an error if Kerberos is not enabled.
cluster_sentry_enabled	true if Sentry is configured on the cluster; false otherwise.

Server Parameters

[Table 13-12](#) (page 13-15) describes the server parameters for `bdaccli getinfo`.

Table 13-12 Server Parameters

Parameter	Returns
server_mammoth_installed	true if the Mammoth utility has deployed the Oracle Big Data Appliance software on this server; false otherwise.
server_name	The name of this server on the client network, such as <code>bda1node01</code> .
server_os_version	The version of Oracle Linux on this server, such as 6.4.

One-Off Patch Parameters

[Table 13-13](#) (page 13-16) describes the one-off patch parameters for `bdaccli getinfo`.

Table 13-13 One-Off Patch Parameters

Parameter	Returns
available_patches	A list of valid patches available for installation. A valid patch has a directory under /opt/oracle/bda/patches or /opt/oracle/BDAMammoth/patches that contains a file named inventory.
installed_patches	A list of patches already installed. An installed patch has a directory in both /opt/oracle/bda/patches and /opt/oracle/BDAMammoth/patches.

See Also

"[Changing the Configuration of Optional Software](#) (page 10-13)" for more information about the optional services.

Examples

The following commands provide information about the optional software on the cluster:

```
# bdaccli getinfo cluster_bdc_installed
true
# bdaccli getinfo cluster_odi_version
11.1.1.7.0

# bdaccli getinfo cluster_hdfs_encryption_enabled
true
```

The following command lists all switches on the current InfiniBand fabric. In this example, three Oracle Big Data Appliance racks are on the fabric with the standard hardware configuration of one spine switch and two gateway switches each.

```
$ bdaccli getinfo ib_switches
bda1sw-iba0 00:21:28:6c:c8:af:a0:a0 36P
bda1sw-ibb0 00:21:28:46:9e:3b:a0:a0 36P
bda1sw-ibs0 00:21:28:6c:c8:ae:a0:a0 36P
bda2sw-ib1 00:21:28:46:98:d3:a0:a0 36P
bda2sw-ib2 00:21:28:de:ae:4a:c0:a0 GTW
bda2sw-ib3 00:21:28:c3:70:9a:c0:a0 GTW
bda3sw-ib1 00:21:28:46:90:ee:a0:a0 36P
bda3sw-ib2 00:21:28:df:34:8a:c0:a0 GTW
bda3sw-ib3 00:21:28:df:0f:0a:c0:a0 GTW
bda4sw-ib1 00:21:28:e8:af:23:a0:a0 36P
bda4sw-ib2 00:10:e0:0c:48:a0:c0:a0 GTW
bda4sw-ib3 00:21:28:f4:82:ce:c0:a0 GTW
```

This example installs patch 1234:

```
$ bdaccli add patch 1234
```

13.1.8 bdaclustersynctime

Uses network time protocol (NTP) to synchronize the time of all units in the cluster.

Syntax

`bdacustersynctime`

Usage Notes

You must be logged in as the root user.

Example

The following example shows that the cluster is synchronized.

```
# bdacustersynctime
INFO: Logging results to /tmp/bdacuster_1399300113/
SUCCESS: Mammoth configuration file is valid.
SUCCESS: All cluster host names are pingable
SUCCESS: NTP service running on all cluster hosts
SUCCESS: At least one valid NTP server found
SUCCESS: No errors found syncing date and time on all nodes
SUCCESS: Max clock drift of 0 seconds is within limits
SUCCESS: Sync date and time of cluster succeeded
```

13.1.9 bdadiag

Collects diagnostic information about an individual server for Oracle Support.

Syntax

`bdadiag [parameter]`

Parameters

cm

Downloads diagnostics from Cloudera Manager. You must know the Cloudera Manager admin password to use this parameter.

hadoop

Collects the CDH cluster logs from `/var/log`.

hdfs

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

ilom

Gathers ILOM data using `ipmitool`. You cannot use `ilom` in the same command as `snapshot`.

osw

Collects Oracle OS Watcher logs, which include historical operating system performance and monitoring data. The output can consume several hundred megabytes of disk space.

snapshot

Collects ILOM snapshot data over the network, and provides most useful output than the `ilom` option. You must know the server root password to use this parameter. You cannot use `snapshot` in the same command as `ilom`.

Return Value

The name of the compressed file in the /tmp directory where bdadiag stored the data. The file name has the form bdadiag_server-name_server-serial-number_date.tar.bz2.

The logs are organized in subdirectories, including the following:

```
asr
ilom
install
messages
net
raid
sysconfig
```

Usage Notes

You must be connected to the server as root.

The optional parameters instruct bdadiag to collect additional diagnostics. You can enter the options together on the command line to collect the most information.

You run bdadiag at the request of Oracle Support and associate it with an open Service Request (SR). See the *Oracle Big Data Appliance Software User's Guide* for details about providing diagnostics to Oracle Support.

Example

This example shows the basic output from the utility:

```
# bdadiag

Big Data Appliance Diagnostics Collection Tool v4.0.0

Gathering Linux information

Skipping ILOM collection. Use the ilom or snapshot options, or login to
ILOM over the network and run Snapshot separately if necessary.

Generating diagnostics tarball and removing temp directory

=====

Done. The report files are bzip2 compressed in /tmp/
bdadiag_bda1node0101_12      16FM5497_2014_01_18_06_49.tar.bz2
=====
```

The next example shows the additional output from the cm option.

```
# bdadiag cm

Big Data Appliance Diagnostics Collection Tool v4.0.0

Getting Cloudera Manager Diagnostics
Password for the Cloudera Manager admin user is needed
Enter password: password
Enter password again: password
Passwords match
Waiting for Cloudera Manager ...
Succeeded. Output in : /opt/oracle/BDAMammoth/bdaconfig/tmp/cm_commands.out
```

```

Collecting diagnostic data ...
{ "startTime" : "2014-09-07T14:18", "endTime" : "2014-09-06T14:18" }
Succeeded. Output in : /opt/oracle/BDAMammoth/bdaconfig/tmp/
cm_commands_collectD
iagnosticData.out
Command ID is 526
.....
Command 526 finished after 30 seconds
Operation completed successfully
Diagnostic data successfully collected
Can be downloaded from URL http://bda1node03.example.com:7180/cmf/command/526/
download
Downloading diagnostic data ...
Original Cloudera Manager Diagnostics Bundle Name : 3609df48-4930-11e1-9006-
b8ac6f8061c1.cluster1.20140907-18-59.support-bundle.zip
Data successfully downloaded and placed in /tmp/
bdadiag_bda2node01_1408NM50AE_2014_09_07_14_18/3609df48-4930-11e1-9006-
b8ac6f8061c1.AK00198532.cluster1.20140907-18-59.support-bundle.zip
Gathering Linux information
.
.
.

```

The next example shows the additional output from the snapshot option:

```

# bdadiag snapshot

Big Data Appliance Diagnostics Collection Tool v4.0.0

Please enter Host OS root password required for snapshot: password

Gathering Linux information

Gathering ILOM Snapshot data - please be patient, this may take a long time

snapshot running: Tue Sep  2 13:42:28 PDT 2014
snapshot running: Tue Sep  2 13:43:21 PDT 2014
snapshot running: Tue Sep  2 13:44:32 PDT 2014
snapshot running: Tue Sep  2 13:45:35 PDT 2014
snapshot running: Tue Sep  2 13:46:31 PDT 2014
snapshot running: Tue Sep  2 13:47:43 PDT 2014
snapshot running: Tue Sep  2 13:48:47 PDT 2014
Snapshot Collection completed.

Generating diagnostics tarball and removing temp directory

=====
Done. The report files are bzip2 compressed in /tmp/
bdadiag_bda1node01_1143FMM06E_2014_09_02_14_40.tar.bz2
=====

```

13.1.10 bdadiagcluster

Creates a diagnostic bundle for the cluster.

Syntax

```
bdadiagcluster
```

Usage Notes

You must be logged in as the root user. To include the diagnostics provided by Cloudera Manager, you must provide the Cloudera Manager password twice when prompted for it.

Example

The following example creates a bundled zip file from the diagnostic zip files creates on the servers in a six-node cluster. It does not include the Cloudera Manager diagnostics.

```
# bdadiagcluster
Enter CM admin password to run dumpcluster
Press ENTER twice to skip CM services and hosts checks
Enter password: Enter
Enter password again: Enter
INFO: Starting Big Data Appliance diagnose cluster at Mon May  5 07:34:03 2014
INFO: Logging results to /tmp/bdadiagcluster_1399300440.log
SUCCESS: Created BDA diagcluster zipfile on node bda01node01
SUCCESS: Created BDA diagcluster zipfile on node bda01node02
SUCCESS: Created BDA diagcluster zipfile on node bda01node03
SUCCESS: Created BDA diagcluster zipfile on node bda01node04
SUCCESS: Created BDA diagcluster zipfile on node bda01node05
SUCCESS: Created BDA diagcluster zipfile on node bda01node06
SUCCESS: bdadiagcluster_1399300440.zip created
INFO: Big Data Appliance diagnose cluster complete at Mon May  5 07:34:48 2014
INFO: Please get the Big Data Appliance cluster diagnostic bundle at /tmp/
bdadiagcluster_1399300440.zip
```

13.1.11 bdagetclientconfig

Deploys the HDFS, MapReduce, and Hive client configuration files from Cloudera Manager.

Syntax

```
bdagetclientconfig
```

Usage Notes

You must be connected to the server as root.

To deploy a new client configuration to all nodes of the cluster, use the `dcli -c` command.

Example

This example shows the output from one node in the cluster:

```
# bdagetclientconfig

bdagetclientconfig : Download and deploy HDFS, Map-Reduce and Hive client
configuration files

Logging to /tmp/bdagetclientconfig-1368541073.out

Downloading HDFS and Map-Reduce client configuration zipfile
```

Downloading Hive client configuration zipfile

Deploying HDFS, Map-Reduce and Hive client configurations

Successfully downloaded and deployed HDFS, Map-Reduce and Hive client configurations !

13.1.12 bdaid

Returns information about an individual server.

If you need to contact Oracle Support about an issue with Cloudera's Distribution including Apache Hadoop, then run this command first.

Syntax

bdaid

Usage Notes

You must be connected to the server as root.

Example

This example shows the output from the utility:

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

13.1.13 bdaibstatus

The bdaibstatus utility checks status for the following networks by pinging the gateway for each network.

- The client network bondeth0 and its slaves eth8 and eth9 .
- The private network bondib0 and its slaves ib0 and ib1.

For releases Oracle Big Data Appliance 4.7 and greater, bdaibstatus replaces bdacheckvnics, which provided similar functionality.

Run this command after connecting as root to any server.

Syntax

bdaibstatus [-bsgpch]

Parameters

-b

Brief. No heading.

-s

Test the passive network interface also. Temporarily make each passive slave interface active, ping the gateway from the slave, and then return the slave to passive mode.

-g

Bring down the primary interface if both of the following are true:

- The primary interface is bad (fails to ping the gateway).
- The non-primary interface is good (successfully pings the gateway).

This parameter includes an implicit -s.

-p

Display results for private InfiniBand interfaces only.

-c

Display results for client network interfaces only.

-h

Display this help information.

Example

```
[root@node01 network]# bdaibstatus
host if status actv primary pkey
=====
bdanode01 ib0 up yes yes NO
bdanode01 ib1 up no no NO
host if status actv primary switch gw port ping gw vlan
=====
=====
bdanode01 eth9 up yes yes bdasw-ib3 0A-ETH-1 yes N/A
```

13.1.14 bdaimagevalidate

Validates the hardware and software on a server by running [bdacheckhw](#) (page 13-3), and then [bdachecksw](#) (page 13-5).

Syntax

```
bdaimagevalidate [startup]
```

Parameters**startup**

Regenerates the bda_reboot_status and BDA_REBOOT_* files in /root, in addition to performing the validation checks. Use this parameter if the checks fail after restarting the server, such that either BDA_REBOOT_FAILED or BDA_REBOOT_WARNINGS exist, and the issue is resolved. Do not use this parameter for the initial set of checks, that is, if /root/bda_reboot_status does not exist.

Example

```
# bdaimagevalidate
SUCCESS: Found BDA v3 server : SUN SERVER X4-2L
SUCCESS: Correct processor info : Intel(R) Xeon(R) CPU E5-2650 v2 @ 2.60GHz
```

```
SUCCESS: Correct number of types of CPU : 1
SUCCESS: Correct number of CPU cores : 32
SUCCESS: Sufficient GB of memory (>=63): 63
SUCCESS: Correct BIOS vendor : American Megatrends Inc.
SUCCESS: Sufficient BIOS version (>=08080102): 26010600
SUCCESS: Recent enough BIOS release date (>=05/23/2011):07/08/2013
SUCCESS: Correct ILOM major version : 3.1.2.32
SUCCESS: Sufficient ILOM minor version (>=82440): 82440
.
.
.
SUCCESS: Correct bda-monitor status : bda monitor is running
SUCCESS: Big Data Appliance software validation checks succeeded
SUCCESS: All Big Data Appliance validation checks succeeded
```

13.1.15 bdanodedesc

Displays the server name and IP address.

Syntax

```
bdanodedesc
```

Example

This example shows that the user is logged in to a BDA server named bda1node03 with an IP address of 192.168.40.61:

```
$ bdanodedesc
bda1node03 BDA 192.168.40.61 HCA-
```

13.1.16 bdaredoclientnet

Re-creates the virtual network interface cards (VNICs) for all servers in the rack and spreads them across the available 10 GbE ports.

Syntax

```
bdaredoclientnet
```

Usage Notes

Log in to server 1 and change to the /opt/oracle/bda/network directory to run this utility.

You must run this utility after changing the number of 10 GbE connections to a Sun Network QDR InfiniBand Gateway switch. See "[Changing the Number of Connections to a Gateway Switch](#) (page 11-12)".

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

Example

The following example shows the output from the bdaredoclientnet utility:

```
# cd /opt/oracle/bda/network
# bdaredoclientnet
bdaredoclientnet: do basic sanity checks on /opt/oracle/bda/network.json
bdaredoclientnet: warning: bda deploy /opt/oracle/bda/network.json check odd nr of
elements 7
bdaredoclientnet: passed
bdaredoclientnet: checking for networkexpansion.json
src/query_smp.c:196; umad (DR path slid 0; dlid 0; 0,1,13,34,20 Attr 0xff90:1) bad
status 110; Connection timed out
src/query_smp.c:196; umad (DR path slid 0; dlid 0; 0,1,13,31,20 Attr 0xff90:2) bad
status 110; Connection timed out
bdaredoclientnet: ping servers by name on admin network
bdaredoclientnet: ping servers by name on admin network
bdaredoclientnet: passed
bdaredoclientnet: verify infiniband topology
bdaredoclientnet: passed
bdaredoclientnet: start setup client network (10gigE over Infiniband)
bdaredoclientnet: ping both gtw leaf switches
bdaredoclientnet: passed
bdaredoclientnet: verify existence of gateway ports
bdaredoclientnet: passed
bdaredoclientnet: removing existing eoib setup for this server
Shutting down interface bondeth0:          [ OK ]
Shutting down interface bondib0:           [ OK ]
Shutting down interface eth0:              [ OK ]
Shutting down loopback interface:          [ OK ]
Disabling IPv4 packet forwarding: net.ipv4.ip_forward = 0
                                           [ OK ]
Bringing up loopback interface:             [ OK ]
Bringing up interface bondib0:              [ OK ]
Bringing up interface eth0:                 [ OK ]
bdaredoclientnet: ping server ips on admin network
bdaredoclientnet: passed
bdaredoclientnet: ping servers by name on admin network
bdaredoclientnet: passed
bdaredoclientnet: test ssh server ips on admin network
hello from bda1node02.example.com
hello from bda1node03.example.com
.
.
.
bdaredoclientnet: passed
bdaredoclientnet: check existence of default vlan for port 0A-ETH-1 on bda1sw-ib2
bdaredoclientnet: use existing default vlan for port 0A-ETH-1 on bda1sw-ib2
bdaredoclientnet: check existence of default vlan for port 0A-ETH-1 on bda1sw-ib3
bdaredoclientnet: use existing default vlan for port 0A-ETH-1 on bda1sw-ib3
bdaredoclientnet: passed
bdaredoclientnet: apply eoib on each server
bdaredoclientnet: wait a few seconds for the network to restart on 10.111.22.001
bdaredoclientnet: wait a few seconds for the network to restart on 10.111.22.002
.
.
```

```

.
check and delete vNIC for bda1node02 on switch bda1sw-ib2
vNIC ID 757 deleted
IO Adapter for vNIC deleted
check and delete vNIC for bda1node02 on switch bda1sw-ib3
check and delete vNIC for bda1node02 on switch bda1sw-ib2
check and delete vNIC for bda1node02 on switch bda1sw-ib3
vNIC ID 707 deleted
IO Adapter for vNIC deleted
create vNIC eth9 bda1node02 on switch bda1sw-ib3
vNIC created
create vNIC eth8 bda1node02 on switch bda1sw-ib2
vNIC created
.
.
.
bdaredoclientnet: ping server ips on client network
bdaredoclientnet: passed
bdaredoclientnet: test ssh server ips on client network
hello from bda1node02.example.com
hello from bda1node03.example.com
.
.
.
bdaredoclientnet: passed
bdaredoclientnet: end setup client network

```

13.1.17 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.

Syntax

```
bdaserials
```

Example

This example shows the output from the utility:

```

# bdaserials
Rack serial number : AK00023713
System serial number : 1137FMM0BY
System UUID : 080020FF-FFFF-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:80:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
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:80:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
ib1 InfiniBand : 80:00:00:4B:FE:80:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00

```

13.1.18 bdaswapoff

Turns off swapping by the operating system.

Syntax

```
bdaswapoff
```

Usage Notes

The `bdaswapoff` utility disables both swap partitions on a server, which disables all swapping by the operating system. This state persists when the server restarts; you must run `bdaswapon` to restore swapping. Swapping is turned off by default to improve performance and to allow high availability if a disk fails.

Use `bdaswapoff` instead of the Linux `swapoff` utility.

See Also

"[bdaswapon](#) (page 13-26)"

Example

This example turns off swapping:

```
# bdaswapoff
#
```

13.1.19 bdaswapon

Turns on paging and swapping by the operating system.

Syntax

```
bdaswapon
```

Usage Notes

Swapping is turned off by default to improve performance and the ability to recover from disk failure.

Use `bdaswapon` instead of the Linux `swapon` utility.

See Also

["bdaswapoff \(page 13-26\)"](#)

Example

This example turns on swapping:

```
# bdaswapon
Filename                                Type      Size    Used    Priority
/dev/sda3                              partition 12287992 0        1
/dev/sdb3                              partition 12287992 0        1
```

13.1.20 bdaupdatefw

Updates the firmware of a particular component of a server, such as a replacement disk drive.

Syntax

`bdaupdatefw parameter`

Parameters

-d *N*

Updates the LSI disk firmware for the specified disk (*N*). Each server has 12 disks, which are numbered from 0 to 11.

-f

Specifies the file path to the firmware. If the path is omitted, then `bdaupdatefw` uses the default firmware for the specified component from `/opt/oracle/bda/firmware`.

-h

Displays syntax and usage information for `bdaupdatefw`.

-i

Updates the Oracle ILOM firmware.

-l

Updates the LSI disk controller firmware.

-m

Updates the firmware for the Mellanox host channel adapter (InfiniBand card).

Usage

This utility is typically run by Oracle field engineers when installing or replacing hardware components, which may not be factory-installed with a supported firmware version. During a software installation, Mammoth copies the currently supported firmware to Oracle Big Data Appliance. The `bdaupdatefw` command uses those files when they are needed to update the firmware of a server component.

You can update one firmware package in a single command. Thus, you can specify only one of the following parameters: `-d`, `-i`, `-l`, or `-m`.

 Caution:

Only use the firmware provided in a Mammoth bundle. Do not attempt to install firmware downloaded from a third-party site. Doing so may result in the loss of warranty and support.

Examples

This example shows the output from a command to update the Oracle ILOM firmware. To perform the update, you must execute the `ipmiflash` command provided in the output.

```
# bdaupdatefw -i
[INFO:GENERAL] No firmware file specified. Using default firmware file - /opt/oracle/bda/firmware/ILOM-3_2_0_r74388-Sun_Fire_X4270_M3.pkg
[INFO:GENERAL] Updating ILOM firmware with the firmware file /opt/oracle/bda/firmware/ILOM-3_2_0_r74388-Sun_Fire_X4270_M3.pkg
[INFO:GENERAL] Original version is: 3.1.2.12 r74388
[INFO:GENERAL]
[INFO:GENERAL] Please run the following command and enter the root password
[INFO:GENERAL] for the ILOM when requested
[INFO:GENERAL]
[INFO:GENERAL] Note that this command will shutdown the server after flashing.
[INFO:GENERAL] You will need to login to the ILOM to power on the server afterwards.
[INFO:GENERAL]
[INFO:GENERAL] ipmiflash -v -I lanplus -H 10.133.46.218 -U root write /opt/oracle/bda/firmware/ILOM-3_1_2_12_r74388-Sun_Fire_X4270_M3.pkg
[INFO:GENERAL]
```

13.1.21 iblinkinfo

Lists all InfiniBand connections in the InfiniBand network.

Syntax

```
iblinkinfo
```

Usage Notes

Run this command as root from any server.

Example

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 bda1sw-ib2 10.133.43.36:
149 1[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 130 2[ ] "SUN IB QDR GW switch bda1sw-ib2 10.133...
149 2[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 127 1[ ] "SUN IB QDR GW switch bda1sw-ib2 10.133...
149 3[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 111 2[ ] "SUN IB QDR GW switch bda1sw-ib2 10.133...
149 4[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 109 1[ ] "SUN IB QDR GW switch bda1sw-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) ==> [ ] "" ( )
.
.
.

```

13.1.22 imagehistory

Displays a history of operating system upgrades.

Syntax

```
imagehistory
```

Example

This example shows that the appliance was imaged with version 3.1.0 with an upgrade to 4.0.0

```
$ imagehistory
```

```
Big Data Appliance Image History
```

```

IMAGE_VERSION           : 3.1.0
IMAGE_CREATION_DATE      : Tue Jun 24 06:44:22 UTC 2014
IMAGING_START_DATE      : Sat Aug 30 18:15:01 UTC 2014
IMAGING_END_DATE        : Sat Aug 30 14:59:39 EDT 2014
DEPLOYMENT_VERSION      : 4.0.0
DEPLOYMENT_START_DATE   : Sat Aug 30 15:18:07 EDT 2014
DEPLOYMENT_END_DATE     : Sat Aug 30 16:39:07 EDT 2014

```

13.1.23 imageinfo

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

Syntax

```
imageinfo
```

Example

This example identifies the 4.0.0 image:

```
$ imageinfo
```

```
Big Data Appliance Image Info
```

```

IMAGE_CREATION_DATE      : Sat Aug 30 13:55:26 EDT 2014
IMAGE_LABEL              : BDA_MAIN_LINUX.X64_140829
IMAGE_VERSION            : 4.0.0
LINUX_VERSION            : Oracle Linux Server release 5.8
KERNEL_VERSION           : 2.6.39-400.215.9.el5uek
BDA_RPM_VERSION          : bda-4.0.0-1.el5
OFED_VERSION             : OFED-IOV-1.5.5-2.0.0088
JDK_VERSION              : jdk-1.7.0_65-fcs
HADOOP_VERSION           : 2.3.0-cdh5.1.0

```

13.1.24 listlinkup

Shows the Ethernet bridge ports with active links.

Syntax

```
listlinkup
```

Usage Notes

Run this command after connecting as root to a Sun Network QDR InfiniBand Gateway switch.

Example

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)
```

13.1.25 mount_hadoop_dirs

Makes password-encrypted data accessible after the server is restarted. Until you run this utility and enter the correct password, the data is unreadable.

Syntax

```
mount_hadoop_dirs
```

Example

This example shows the prompt for entering the password:

```
$ mount_hadoop_dirs
Enter password to mount Hadoop directories: password
```

13.1.26 remove-root-ssh

Removes passwordless SSH previously established by the [setup-root-ssh](#) (page 13-32) command.

Syntax

```
remove-root-ssh [-C | -c | -g | -j] [-p]
```

```
remove-root-ssh -h
```

Parameters

-C

Targets all servers in the cluster, using the list of servers in `/opt/oracle/bda/cluster-hosts-infiniband`.

-c *host1, host2,...*

Targets the servers specified as *host1*, *host2*, and so forth, on the command line.

-g *groupfile*

Targets a user-defined set of servers listed in *groupfile*. You can enter either server names or IP addresses in the file, one per line.

-j "etho0_ips[range]"

Specifies the range of servers in a starter rack [1-6] or a starter rack with additional servers [1-12]. This parameter is required in the 2.2.x base image when the utility is used before network configuration.

-h

Displays Help.

-p *password*

Specifies the root password on the command line.

Oracle recommends that you omit this parameter. You will be prompted to enter the password, which the utility does not display on your screen.

Usage Notes

You must know the root password to use this command.

If you do not specify the target servers, then `remove-root-ssh` uses all servers in the rack.

See Also

["setup-root-ssh \(page 13-32\)"](#), ["Setting Up Passwordless SSH \(page 14-1\)"](#)

Example

This example shows passwordless SSH being removed:

```
# remove-root-ssh
calling /opt/oracle/bda/bin/dcli -c 192.168.42.37,192.168.42.38,192.168.42.39 --unkey
192.168.42.37: ssh key dropped
192.168.42.38: ssh key dropped
192.168.42.39: ssh key dropped
remove-root-ssh succeeded
```

13.1.27 reset-boot-order

Resets the boot order of the server to the factory defaults, as specified in the BIOS. By doing so, it clears any ILOM booting overrides.

Syntax

```
reset-boot-order
```

Example

The following example resets the boot order of the current server:

```
# reset-boot-order
Set Boot Device to none

Cleared ILOM boot override - Boot device: none

Found BDA v1 Hardware - setting boot order using biosconfig

Copyright (C) SUN Microsystems 2009.
BIOSconfig Utility Version 2.2.1
Build Date: Aug 24 2009
Build Time: 09:01:30

BIOSconfig Specification Version 2.4

Processing Input BIOS Data....

Success

Found USB device name : USB:02.82;01 Unigen PSA4000
.
.
.
BIOSconfig Specification Version 2.4

Processing Input BIOS Data....

Success
New BIOS boot order :
    USB:02.82;01 Unigen PSA4000
    RAID:Slot0.F0:(Bus 13 Dev 00)PCI RAID Adapter
    PXE:IBA GE Slot 0100 v1331
    PXE:IBA GE Slot 0101 v1331
    PXE:IBA GE Slot 0700 v1331
    PXE:IBA GE Slot 0701 v1331
```

13.1.28 setup-root-ssh

Establishes passwordless SSH for the root user.

Syntax

```
setup-root-ssh [-C | -c | -g | -j] [-p]
```

```
setup-root-ssh -h
```

Parameters

-C

Targets all servers in the cluster, using the list of servers in /opt/oracle/bda/cluster-hosts-infiniband.

-c *host1*, *host2*,...

Targets the servers specified as *host1*, *host2*, and so forth, on the command line.

-g *groupfile*

Targets a user-defined set of servers listed in *groupfile*. You can enter either server names or IP addresses in the file, one per line.

-j "etho0_ips[range]"

Specifies the range of servers in a starter rack [1-6] or a starter rack with additional servers [1-12]. This parameter is required in the 2.2.x base image when the utility is used before network configuration.

-h

Displays Help.

-p *password*

Specifies the root password on the command line.

Oracle recommends that you omit this parameter. You will be prompted to enter the password, which the utility does not display on your screen.

Usage Notes

You must know the root password to use this command.

If you do not specify the target servers, then `setup-root-ssh` uses all servers in the rack.

See Also

["remove-root-ssh \(page 13-30\)"](#), ["Setting Up Passwordless SSH \(page 14-1\)"](#)

Example

This example shows passwordless SSH being set up for root:

```
# setup-root-ssh
Enter root password: password
spawn /opt/oracle/bda/bin/dcli -c 192.168.42.37,192.168.42.38... -k
root@192.168.42.37's password:
root@192.168.42.38's password:
.
.
.
192.168.42.37: ssh key added
192.168.42.38: ssh key added
.
.
.
setup-root-ssh succeeded
```

13.1.29 showusb

Shows the device location of an inserted USB drive as it is known to the operating system, such as `/dev/sdn`.

Syntax

```
showusb
```

13.1.30 showvlan

Lists the VLANs configured on the switch.

Syntax

```
showvlan
```

Usage Notes

Run this command after connecting as root to a Sun Network QDR InfiniBand Gateway switch.

Example

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

```
# showvlan
Connector/LAG  VLN  PKEY
-----
0A-ETH-1      0    ffff
0A-ETH-2      0    ffff
0A-ETH-3      0    ffff
0A-ETH-4      0    ffff
1A-ETH-1      0    ffff
1A-ETH-3      0    ffff
1A-ETH-4      0    ffff
```

13.1.31 showvnics

Lists the virtual network interface cards (VNICs) created for the switch.

Syntax

```
showvnics
```

Usage Notes

Run this command after connecting as root to a Sun Network QDR InfiniBand Gateway switch.

Example

This example shows the VNICs created in a round-robin process for switch 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
.
.
.
```

Executing Commands Across a Cluster 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 [Configuring an Oracle Big Data Appliance Full Rack or Starter Rack](#) (page 7-1).

This chapter contains the following sections:

- [Overview of the dcli Utility](#) (page 14-1)
- [dcli Syntax](#) (page 14-3)
- [dcli Return Values](#) (page 14-5)
- [dcli Examples](#) (page 14-5)

14.1 Overview of the dcli Utility

The `dcli` utility executes commands on multiple Oracle Big Data Appliance servers using the administrative (eth0) interface to make the connections. You can run the utility from any server.

14.1.1 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 14-1](#) (page 14-4) for details about these options.

 Note:

The Mammoth utility requires remote root access. However, Oracle Big Data Appliance does not require it for normal operation.

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
```

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
```



See Also:

Full parameter descriptions of [setup-root-ssh](#) (page 13-32) and [remove-root-ssh](#) (page 13-30).

14.1.2 Basic Use of dcli

This topic identifies some basic options to the dcli command. For a complete list of options, see [Table 14-1](#) (page 14-4).

14.1.2.1 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.

14.1.2.2 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 14-1](#) (page 14-4).

You can manually create 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.

14.1.2.3 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 14-1](#) (page 14-4).

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

14.1.2.4 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 servers.

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

```
# dcli date
bda1node07-adm.example.com: Tue Feb 14 10:22:31 PST 2012
```

The minimal output returns OK for completed execution:

```
# dcli -n date
OK: ['bda1node07.example.com']
```

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

```
dcli -v dateoptions.nodes: Noneoptions.destfile: Noneoptions.file:
Noneoptions.group: dcserveroptions.maxLines: 100000options.listNegatives:
Falseoptions.pushKey: Falseoptions.regex: Noneoptions.sshOptions:
Noneoptions.scpOptions: Noneoptions.dropKey: Falseoptions.serializeOps:
Falseoptions.userID: rootoptions.verbosity 1options.vmstatOps Noneoptions.execfile:
Noneargv: ['/opt/oracle/bda/bin/dcli', '-g', 'dcservers', '-v', 'date']Success
connecting to nodes: ['bda1node07.example.com']...entering thread for
bda1node07.example.com:execute: /usr/bin/ssh -l root bda1node07.example.com '
date' ...exiting thread for bda1node07.example.com status: 0bda1node07.example.com:
Tue Feb 14 10:24:43 PST 2012]
```

14.2 dcli Syntax

```
dcli [option] [command]
```

Parameters

option

An option described in [Table 14-1](#) (page 14-4). You can omit all options to run a command on all servers in the current rack.

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 14-1 dcli Options

Option	Description
-c <i>nodes</i>	Specifies a comma-separated list of Oracle Big Data Appliance servers where the command is executed
-C	Uses the list of servers in /opt/oracle/bda/cluster-rack-infiniband as the target. See " Identifying the Target Servers (page 14-2)".
-d <i>destfile</i>	Specifies a target directory or file name for the -f option
-f <i>file</i>	Specifies files to be copied to the user's home directory on the target servers. The files are not executed. See the -l option.
-g <i>groupfile</i>	Specifies a file containing a list of Oracle Big Data Appliance servers where the command is executed. You can use either server names or IP addresses in the file.
-h, --help	Displays a description of the commands
-k	Pushes the ssh key to each server's /root/.ssh/authorized_keys file. See " Setting Up Passwordless SSH (page 14-1)" for an easier alternative.
-l <i>userid</i>	Identifies the user ID for logging in to another server. The default ID is root.
--maxlines= <i>maxlines</i>	Identifies the maximum lines of output displayed from a command executed on multiple servers. The default is 10,000 lines.
-n	Abbreviates the output for non-error messages. Only the server name is displayed when a server returns normal output (return code 0). You cannot use the -n and -r options together.
-r <i>regexp</i>	Replaces the output with the server name for lines that match the specified regular expression
-s <i>sshoptions</i>	Specifies a string of options that are passed to SSH
--scp= <i>scpoptions</i>	Specifies a string of options that are passed to Secure Copy (SCP), when these options are different from <i>sshoptions</i>
--serial	Serializes execution over the servers. The default is parallel execution.
-t	Lists the target servers
--unkey	Drops the keys from the authorized_key files of the target servers
-v	Displays the verbose version of all messages
--version	Displays the dcli version number

Table 14-1 (Cont.) dcli Options

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

14.3 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.

14.4 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
.
.
.
```

15

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 [Configuring an Oracle Big Data Appliance Full Rack or Starter Rack](#) (page 7-1).

This chapter contains the following sections:

- [Oracle ILOM Overview](#) (page 15-1)
- [Administrative Network Diagram](#) (page 15-2)
- [Oracle ILOM IP Addresses for Oracle Big Data Appliance Components](#) (page 15-3)
- [Connecting to Oracle ILOM Using the Network](#) (page 15-3)
- [Connecting to Oracle ILOM Using a Serial Connection](#) (page 15-5)

See Also:

- For Sun Server X4-2L and Sun Server X3-2L servers, the Oracle Integrated Lights Out Manager 3.1 documentation library at http://docs.oracle.com/cd/E24707_01/index.html
- For Sun Fire X4270 M2 servers, Oracle Integrated Lights Out Manager 3.0 documentation library at <http://docs.oracle.com/cd/E19860-01/>

15.1 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.

15.1.1 Oracle ILOM Interfaces

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

Web Interface

The web interface enables you use a browser to log in to the SP, 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.

15.1.2 Oracle ILOM Users

The Oracle Big Data Appliance servers are configured initially with only the `root` user.


The InfiniBand switches are configured with these additional 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.

As `root`, you can create additional ILOM users if you want.

15.2 Administrative Network Diagram

[Figure 15-1](#) (page 15-3) 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 15-1 Administrative Network in Oracle Big Data Appliance

15.3 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 \(page 3-9\)"](#).

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

15.4 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 (Sun Fire X4270 M2-based racks only).

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 [Configuring an Oracle Big Data Appliance Full Rack or Starter Rack \(page 7-1\)](#)

This section discusses the following topics:

- [Connecting to the Command-Line Interface \(page 15-4\)](#)

- [Connecting to the Web GUI](#) (page 15-4)

15.4.1 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 bda1sw-ib2.
3. Ensure that **SSH** is chosen as the **Connection Type**.
4. Type the user name and the password, when prompted. The default user name is root, and the default password is welcome1.

The CLI command prompt (#) is displayed.

5. Enter help for a list of commands.



See Also:

Oracle ILOM Daily Management -- CLI Procedures Guide at
<http://docs.oracle.com/cd/E19860-01/>

15.4.2 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 your user name and password. See "[Oracle ILOM Users](#) (page 15-2)" for the factory-configured user names. The initial 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 15-2](#) (page 15-5).

Figure 15-2 Oracle ILOM 3.1 Web Interface

Oracle® Integrated Lights Out Manager

System Information	System Monitoring	Power Management	Storage	Configuration	User Management	Remote Control	Maintenance
Overview	Components	Fault Management	Identification Information	Banner Messages	Session Timeout	Versions	

System Overview
View system summary information. You may also change power state and view system status and fault information.

Product Name: SUN FIRE X4270 M2 SERVER
Part/Serial Number: 30057782+1+1 / 1137FMM0C8

Host Power: **On** [Change...](#)
System Status: **Faulted** [View...](#)

BIOS Version: 08120104
SP Hostname: bda1node08-ilom
Uptime: 189 days, 21:57:38
IP Address: 10.128.17.229
ILOM Version: v3.0.16.10.d r74499

See Also:

- *Oracle ILOM 3.1 Quick Start* at http://docs.oracle.com/cd/E24707_01/html/E24528/gentextid-219.html#scrolltoc
- *Oracle ILOM 3.0 Daily Management —Web Procedures Guide* at <http://docs.oracle.com/cd/E19860-01/>

15.5 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

15.5.1 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 as `root` using the default password of `welcome1`.

The Oracle ILOM default command prompt is displayed.

15.5.2 Connecting to the Oracle ILOM of an InfiniBand Switch

To connect to Oracle ILOM of an InfiniBand 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. Log in using the administrative user name and password (`ilom-admin` and `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.

15.6 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:

- For Sun Server X4-2L and Sun Server X3-2L servers, the Oracle Integrated Lights Out Manager 3.1 documentation library at http://docs.oracle.com/cd/E24707_01/index.html
- For Sun Fire X4270 M2 servers, the Oracle Integrated Lights Out Manager 3.0 documentation library at <http://docs.oracle.com/cd/E19860-01/>

15.6.1 Backing Up the Oracle ILOM Configuration Settings

To back up the Oracle ILOM configuration settings on a Sun Server X4-2L or X3-2L server:

1. Open your browser on any system that is on the same network as Oracle Big Data Appliance, and navigate to the Oracle ILOM service processor. This example uses the Oracle ILOM address of a server named node08:

`http://bda1node08-ilom.example.com`

To back up the Oracle ILOM for an InfiniBand switch, enter the switch address, such as bda1sw-ib1.

2. Log in as the root user. The initial password is welcome1.
3. Navigate to the Backup/Restore page:
 - For Oracle ILOM 3.1, select the Maintenance tab, and then select the Backup/Restore subtab.
 - For Oracle ILOM 3.0, expand the ILOM Administration folder in the navigation tree, and select **Configuration Management**.
4. Select the Backup operation and the Browser transfer method, as shown in [Figure 15-3](#) (page 15-8).
5. Enter a pass phrase. The phrase is used to encrypt sensitive information, such as passwords, in the backup file.
6. Click **Run** to start the backup operation. The results are downloaded to your local system in an XML file named config_backup.xml.

7. Save the file in a secure location.
8. Click the **Log Out** button.

Figure 15-3 Oracle ILOM 3.1 Configuration Backup

System Information	System Monitoring	Power Management	Storage	Configuration	User Management	Remote Control	Maintenance
Firmware Upgrade	Backup/Restore	Configuration Management	Reset SP	Snapshot			

Configuration Backup/Restore

Perform system configuration backup or restore from this page. Select Backup or Restore from *Operation* menu. Choose a *Transfer Method* and fill in all required fields. You may choose to supply a *passphrase* to encrypt sensitive data within a backup file or for decrypting such data when restoring a configuration. If a passphrase is not specified, then sensitive data will not be included in the backup file. Click *Run* to start the operation.

Operation:

Transfer Method:

The downloaded file will be saved according to your browser settings.

Passphrase:

Confirm Passphrase:

15.6.2 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-ilom.us.example.com`
2. Log in as the `ilom_admin` user. The default password is `welcome1`.
3. Select the **Maintenance** tab, and then select the **Backup/Restore** subtab.
4. Select the Restore operation and the Browser transfer method.
5. Click **Choose File** and select the `config_backup.xml` file saved previously in a backup operation.
6. Enter the pass phrase that was set during the backup operation.
7. Click **Run** to restore the configuration.

Part V

Appendixes

Part V provides reference information. It contains the following chapters:

- [Error Messages for Oracle Big Data Appliance](#) (page A-1)
- [Replacement Units](#) (page B-1)
- [In-Rack Cabling Tables for Oracle Big Data Appliance X6–2 and X5-2](#) (page C-1)
- [In-Rack Cabling Tables for Oracle Big Data Appliance X4-2 and X3-2](#) (page D-1)
- [In-Rack Cabling Tables for Oracle Big Data Appliance Sun Fire X4270 M2](#) (page E-1)
- [Multirack Cabling Tables](#) (page F-1)

A

Error Messages for Oracle Big Data Appliance

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

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. Follow the instructions in [Replacing a Server Disk](#) (page 12-1) 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. Follow the instructions in [Replacing a Server Disk](#) (page 12-1) 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, follow the instructions in [Replacing a Server Disk](#) (page 12-1) 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. Follow the instructions in [Replacing a Server Disk](#) (page 12-1) 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. Follow the instructions in [Replacing a Server Disk](#) (page 12-1) 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. Follow the instructions in [Replacing a Server Disk](#) (page 12-1) for configuring a disk for either HDFS or Oracle NoSQL Database.

DISKALRT-02007 Solid-state drive disk failure

Cause: A solid-state drive failed.

Action: Replace the drive as soon as possible.

DISKALRT-02008 Solid-state disk predictive failure

Cause: A solid-state drive is in a state of predictive failure.

Action: Replace the drive as soon as possible.

DISKALRT-02009 Solid-state disk poor performance

Cause: A solid-state drive is performing poorly, which indicates that it is failing.

Action: Replace the drive as soon as possible.

B

Replacement Units

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

This appendix contains the following sections:

- [Understanding Repair Categories](#) (page B-1)
- [Disk Controller Batteries](#) (page B-2)
- [LED Status Descriptions](#) (page B-2)
- [Parts for Oracle Big Data Appliance Servers](#) (page B-5)
- [Parts for Sun Datacenter InfiniBand Switch 36](#) (page B-22)
- [Parts for Sun Network QDR InfiniBand Gateway Switch](#) (page B-22)
- [Parts for the Cisco Ethernet Switch](#) (page B-23)
- [Parts for Power Distribution Units](#) (page B-24)
- [Cables and Cable Accessories for an Oracle Big Data Appliance Rack](#) (page B-25)



See Also:

["Spare Parts Kit Components](#) (page 1-4)"

B.1 Understanding Repair Categories

Repair categories describe the potential effects of a repair action. [Table B-1](#) (page 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.	• 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.	• External cables • InfiniBand switch • Ethernet switch

Table B-1 (Cont.) Oracle Big Data Appliance Repair Categories

Category	Description	Scenario Examples
BDA Server Offline (BDA-O)	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. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance (page 11-2)" for additional information. If you monitor the system using Oracle Enterprise Manager Cloud Control, then prevent unnecessary alerts by putting the target in a blackout state. Refer to <i>Oracle Enterprise Manager Administration</i> for additional information.	• System boards • PCIe cards • Memory • Processors

B.2 Disk Controller Batteries

The disk controllers in the 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. The 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.

B.3 LED Status Descriptions

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

- [Sun Server LEDs](#) (page B-3)

- [Sun InfiniBand Switch LEDs](#) (page B-4)

B.3.1 Sun Server LEDs

[Table B-2](#) (page B-3) identifies some LED status indicators on the Sun Server X4-2L and X3-2L that can indicate a fault.



See Also:

"About Server and Component Status Indicators" in the *Sun Server X4-2 Service Manual* at

http://docs.oracle.com/cd/E36975_01/html/E38045/z4000be71405942.html#scrolltoc

Table B-2 Sun Server LED Status Descriptions

Component	LED Status
General Status	The following LEDs are located on the front and back panels of each server: Service Required (amber) • Off indicates normal operation. • Steady on indicates a fault in the server. Power/OK (green) • Steady on indicates that the server is powered on. • A slow blink indicates that the server is starting up. • A steady blink indicates that standby power is on, but chassis power is off. • Off indicates no AC power or an incomplete startup of Oracle ILOM.
Fans	Each fan module has one bicolored LED, which you can see when the top cover fan door is open. Fan Status (amber/green) • Amber indicates that the fan module is faulty. The front and rear panel Service Required LEDs are also lit when the system detects a fan module fault. • Green indicates that the system is powered on, and the fan module is functioning correctly.
Drives	The following LEDs are located on the storage and boot drives: Service Required (amber) • Off indicates normal operation. • Steady on indicates a fault in the drive. OK/Activity (green) • Steady on indicates normal inactivity. • A steady blink indicates normal activity. • Off indicates that the power is off, or the system does not recognize the drive.

Table B-2 (Cont.) Sun Server LED Status Descriptions

Component	LED Status
Power supply	The following LEDs are on each power supply. You can see them from the back of the server. Service Required (amber) • Off indicates normal operation. • Steady on indicates that the system has detected a problem in the power supply regarding the voltage, current, temperature, or fans. AC OK/DC OK (green) • Steady on indicates normal operation for both AC and DC power. • A slow blink indicates that input AC power is normal and output DC power is disabled. • Off indicates no input AC power.

B.3.2 Sun InfiniBand Switch LEDs

Table B-3 (page B-4) identifies some LED status indicators on Sun Datacenter InfiniBand Switch 36 and Sun Network QDR InfiniBand Gateway switches that can indicate a fault.



See Also:

- "Identifying LEDs" in the *Sun Datacenter InfiniBand Switch 36 Service Manual* at http://docs.oracle.com/cd/E26698_01/html/E26434/z40000fe1394411.html
- "Identifying LEDs" in the *Sun Network QDR InfiniBand Gateway Switch Service Manual* at http://docs.oracle.com/cd/E26699_01/html/E26706/z40000fe1394411.html

Table B-3 Sun InfiniBand Switch LED Status Descriptions

Component	LED Status
Chassis	The following LEDs are located on the left side of the rear panel. Attention (amber) • Steady on indicates a fault. • Off indicates normal operation. • Flashing indicates no function. OK (green) • Steady on indicates normal operation. • Off indicates that the switch is off or initializing. • Flashing indicates no function.

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

Component	LED Status
Link status	The link status LEDs are located by the data cable connectors on the rear panel. Link (green) - Steady on indicates that a link is established. - Off indicates that the link is down. - Flashing indicates that there are symbol errors.
Network management ports	The following LEDs are located on the network management connector on the rear panel. Link Speed (amber/green): - Green steady on indicates a 1000BASE-T link. - Amber steady on indicates a 100BASE-T link. - Off indicates no link. - Flashing indicates no function. Activity (green): - Flashing indicates packet activity. - On indicates no function. - Off indicates no activity.
Power supply	The following LEDs are located on the power supply at the front of the chassis. Attention (amber) - Steady on indicates a fault, and 12 VDC is shut down. - Off indicates normal operation. - Flashing indicates no function. OK (green) - Steady on indicates 12 VDC. - Off indicates no DC voltage. - Flashing indicates no function. AC LED is green - Steady on indicates AC power is supplied. - Off indicates that AC power is not present. - Flashing indicates no function.
Fan	The following LED is located on the air inlet on the front of the chassis. FT - On indicates a problem with the fan. - Off indicates normal operation

B.4 Parts for Oracle Big Data Appliance Servers

The following topics identify the replacement parts for the Oracle Big Data Appliance storage servers:

- [Oracle Server X6-2L and X5-2L Parts](#) (page B-6)
- [Sun Server X4-2L Parts](#) (page B-9)
- [Sun Server X3-2L Parts](#) (page B-13)
- [Sun Fire X4270 M2 Parts](#) (page B-17)



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.

B.4.1 Oracle Server X6-2L and X5-2L Parts

Figure B-1 (page B-6) shows the front view of the Oracle Server X6-2L and X5-2L. The hard drives are numbered from left to right, starting in the lower left. The drives in the bottom row are numbers 0, 1, 2, and 3. The drives in the middle row are numbers 4, 5, 6, and 7. The drives in the top row are numbers 8, 9, 10, and 11.

Figure B-1 Front View of Oracle Servers X6-2L and X5-2L



Figure B-2 (page B-6) shows the rear view of a Oracle Server X6-2L or X5-2L storage server.

Figure B-2 Rear View of Oracle Servers X6-2L and X5-2L



Table B-4 (page B-7) lists the replaceable parts for the Oracle Server X6-2L and X5-2L, which are covered under the Oracle Premier Support for Systems warranty. The parts list for X6-2L and X5-2L are nearly the same except for CPUs and DIMMs. Unless otherwise noted, the repair procedures are in the *Oracle Server X5-2L Service Manual*.



Note:

All replacement parts listed in the following table are FRUs when covered under Premier Support of Systems warranty.

Table B-4 Replacement Parts for Oracle Server X6-2L and X5-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
350-1719	Assembly, 1U-2U Snap-in, slim rail rack	EBO	1. If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance (page 11-2)". 2. See "Reinstall the Server in the Rack" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40004191003899.html
370-5621	3V lithium battery	EBO	See "Servicing the Battery" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40000f91037456.html
7014226	Cable, interlock, 125 mm	EBO	See "Servicing SAS Cables" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.baecfjbc.html
7017231	Indicator assembly, right	EBO	See "Servicing the Front LED/USB Indicator Modules" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40001d31014236.html
7023515	Disk backplane assembly, 2 rear	EBO	See "Servicing the Front and Rear Storage Drive Backplanes" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40007ac1419009.html
7044130	Power Supply, A258, 1000W, AC Input	HS	1. Remove rear left side of the CMA, and then swing it back and out to the right. The telescopic cross member supporting the cables directly behind the server can be retracted to access both power supplies. 2. See "Servicing Power Supplies" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40000091014153.html 3. Replace the CMA.
7045042	Fan assembly, 2U, 80 x 80, Ferrite, center rotate	HS	See "Servicing Fan Modules" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40000091014194.html

Table B-4 (Cont.) Replacement Parts for Oracle Server X6-2L and X5-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
7046442	Assembly, dual port, QDR CX3, host channel adapter	EBO	1. If the server is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)". 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40000f91037409.html 3. Verify that the new hardware meets the hardware and firmware specifications: # bdacheckhw
7064884	Heat sink, CPU (Oracle Server X5-2L)	EBO	See "Servicing Storage Drives (CRU)" in <i>Sun Server X5-2 Service Manual</i> at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40001d31037480.html
7067728	Pre-Greased 1U CPU Heatsink (Oracle Server X6-2L)		See "Servicing Storage Drives (CRU)" in <i>Sun Server X5-2 Service Manual</i> at
7066824	Assembly, drive, 8 TB, 3.5 inch, SAS/7200rpm	HS	See "Servicing Storage Drives and Rear Drives" in <i>Sun Server X5-2 Service Manual</i> at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40000091011460.html
7069337	Assembly, backplane, 12X3.5, W/SAS expansion, 2U	EBO	See "Servicing the Front and Rear Storage Drive Backplanes" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40007ac1419009.html
7078071	DIMM 16 GB (Oracle Server X5-2L)	EBO	See "Servicing the DIMMs" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.ceicjagi.html
7310210	32GB DDR4-2400 DIMM (Oracle Server X6-2L)		See "Servicing the DIMMs" at http://docs.oracle.com/cd/E62172_01/html/E62184/ceicjagi.html#scrolltoc
7085209	HBA, 12 Gb SAS PCIe 8-port RAID internal with flash module	EBO	See "Servicing the Oracle Storage 12 Gb/s SAS PCIe RAID HBA Card" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.gocfy.html
7086345	Assembly, super cap, 13.5 volt, 6.4F, SAS3 HBA	EBO	See "Servicing the Oracle Storage 12 Gb/s SAS PCIe RAID HBA Card" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.gocfy.html

Table B-4 (Cont.) Replacement Parts for Oracle Server X6-2L and X5-2L Storage Servers

Part Number	Description	Repair Category	Repair Procedure
7086346	Cable, super cap, 620M, 6-pin, straight, SAS3 HBA	EBO	See "Servicing SAS Cables" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.baecfjbc.html
7090170	Flash drive, MLC, OSA ID, USB, 8 GB	EBO	See "Servicing the Internal USB Flash Drives" at http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.ceiicfbd.html
7095106	CPU, 2.3GHz 18-Core E5-2699 (Oracle Server X5-2L)	EBO	 Caution: Processors should be removed and replaced only by authorized Oracle Service personnel.
7313670	CPU, 22-Core 2.2GHz Intel E5-2699 v4, 145W (Oracle Server X6-2L)		
			http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40001d31037480.html
7096185	Assembly, indicator, left	EBO	 Caution: Processors should be removed and replaced only by authorized Oracle Service personnel.
7098504	ASM, MOBO tray	EBO	
			http://docs.oracle.com/cd/E62172_01/html/E62184/z40001d31037480.html#scrolltoc http://docs.oracle.com/cd/E41033_01/html/E48325/cnpsm.z40001d31014236.html

B.4.2 Sun Server X4-2L Parts

Table B-6 (page B-13) lists the replaceable parts for the Sun Server X4-2L, which are covered under the Oracle Premier Support for Systems warranty. Unless otherwise noted, the repair procedures are in the *Sun Server X4-2 Service Manual*.

Table B-5 Replacement Parts for a Sun Server X4-2L Server

Part Number	Description	Repair Category	Repair Procedure
150-3993	Battery, 3V, 390 MAH, LITH, COIN	BDA-O	See "Servicing the Battery" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000f91037456.html#scrolltoc
7090170	8 GB USB flash drive	BDA-O	See "Servicing the Internal USB Flash Drives" at http://docs.oracle.com/cd/E36974_01/html/E38145/ceiicfbd.html#scrolltoc
371-4985	6 Gigabit SAS RAID PCI battery module, RoHS:Y (LION), BBU-08	BDA-O	1. Power down the server gracefully. See "Powering Off the Server for Orderly Shutdown" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40002401396307.html#scrolltoc 2. 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
7047503	Niwot card	BDA-O	1. Power down the server gracefully. See "Powering Off the Server for Orderly Shutdown" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40002401396307.html#scrolltoc 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000f91037409.html#scrolltoc
7057184	Assembly Battery Carrier Remote	HS	Oracle internal service documentation

Table B-5 (Cont.) Replacement Parts for a Sun Server X4-2L Server

Part Number	Description	Repair Category	Repair Procedure
7082397	InfiniBand dual port, QDR CX3, HCA	BDA-O	1. Power down the server gracefully. See "Powering Off the Server for Orderly Shutdown" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40002401396307.html#scrolltoc 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000f91037409.html#scrolltoc
375-3701	6 Gbps SAS RAID disk PCIe HBA, RoHS:Y	BDA-O	1. Power down the server gracefully. See "Powering Off the Server for Orderly Shutdown" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40002401396307.html#scrolltoc 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000f91037409.html#scrolltoc Note: Do not run any controller commands when replacing the disk.
7066826	Drive, 8 TB, 3.5", SAS/7200 RPM, CORAL	HS	See "Servicing Storage Drives" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000091011460.html#scrolltoc Do not run any controller commands when replacing the disk.
7014391	SAS expander assembly	BDA-O	See "Servicing the SAS Expander Module" at http://docs.oracle.com/cd/E36974_01/html/E38145/ceieidja.html#scrolltoc
7020486	DIMM, 8 GB, DDR3, 1600 MHz, 1.3V	BDA-O	See "Servicing the DIMMs" at http://docs.oracle.com/cd/E36974_01/html/E38145/ceicjagi.html#scrolltoc
7018704	DIMM, 16 GB, DDR3, 1600 MHz, 1.35V	BDA-O	See "Servicing the DIMMs" at http://docs.oracle.com/cd/E36974_01/html/E38145/ceicjagi.html#scrolltoc

Table B-5 (Cont.) Replacement Parts for a Sun Server X4-2L Server

Part Number	Description	Repair Category	Repair Procedure
7071891	LRDIMM, 32GB,DDR3,1600,4Rx4,1.35V	BDA-O	See "Servicing the DIMMs" at http://docs.oracle.com/cd/E36974_01/html/E38145/ceicjagi.html#scrolltoc
7020774	Pregreased CPU heatsink	BDA-O	See "Servicing Processors" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40001d31037480.html#scrolltoc
7075289	CPU, Intel E5-2650 v2 8-core 2.6 GHz	BDA-O	See "Servicing Processors" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40001d31037480.html#scrolltoc
7044130	Power supply, AC, A258, F, 12V, 1000W	HS	See "Servicing Power Supplies" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000091014153.html#scrolltoc
7045042	Dual counter rotating fan module, 80x80 mm	HS	See "Servicing Fan Modules" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40000091014194.html#scrolltoc
7047257	12-slot backplane cable kit with the following: • Disk drive flex cable • Dual backplane power cable • Single backplane power cable • Disk data cable, iPass + HD to iPass+ HD, 600 mm	BDA-O	1. Shut down the rack. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)". 2. See "Servicing the Front and Rear Storage Drive Backplanes" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40007ac1419009.html#scrolltoc
7047768	12-slot 3.5" disk cage, including fan door, and left and right ear assembly	BDA-O	1. Shut down the rack. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)". 2. See "Servicing the Front and Rear Storage Drive Backplanes" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40007ac1419009.html#scrolltoc

Table B-5 (Cont.) Replacement Parts for a Sun Server X4-2L Server

Part Number	Description	Repair Category	Repair Procedure
7049382	System board assembly, RoHS:Y	BDA-O	1. Shut down the rack. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)". 2. "Servicing the Motherboard Assembly" at http://docs.oracle.com/cd/E36974_01/html/E38145/z40001d31037512.html#scrolltoc

B.4.3 Sun Server X3-2L Parts

[Table B-6](#) (page B-13) lists the replaceable parts for the Sun Server X3-2L, which are covered under the Oracle Premier Support for Systems warranty. Unless otherwise noted, the repair procedures are in the *Sun Server X3-2L Service Manual*.

Table B-6 Replacement Parts for a Sun Server X3-2L Server

Part Number	Description	Repair Category	Repair Procedure
150-3993	Battery, 3V, 390 MAH, LITH, COIN	BDA-O	See "Servicing the Battery" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000f91037456.html#scrolltoc Note: Rear storage drive 1 might be populated with a remote battery module for the host bus adapter (HBA) card. See "Remove a Rear Storage Drive" at http://docs.oracle.com/cd/E23393_01/html/E27229/bbggjgea.html#scrolltoc
7090170	8 GB USB flash drive	BDA-O	See "Servicing the Internal USB Flash Drives" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceicfbd.html#scrolltoc

Table B-6 (Cont.) Replacement Parts for a Sun Server X3-2L Server

Part Number	Description	Repair Category	Repair Procedure
371-4985	6 Gigabit SAS RAID PCI battery module, RoHS:Y (LION), BBU-08	BDA-O	1. Power down the server gracefully. See "Powering Down the Server" at http://docs.oracle.com/cd/E23393_01/html/E27229/z400061b1011735.html 2. 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
7047503	Niwot card	BDA-O	1. Power down the server gracefully. See "Powering Down the Server" at http://docs.oracle.com/cd/E23393_01/html/E27229/z400061b1011735.html 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000f91037409.html
7057184	Assembly Battery Carrier Remote	HS	See "Remove a Rear Storage Drive" at http://docs.oracle.com/cd/E23393_01/html/E27229/bbggjgea.html#scrolltoc
375-3696	InfiniBand (CX2) dual port CX2 4x QDR PCI-e	BDA-O	1. Power down the server gracefully. See "Powering Down the Server" at http://docs.oracle.com/cd/E23393_01/html/E27229/z400061b1011735.html 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000f91037409.html

Table B-6 (Cont.) Replacement Parts for a Sun Server X3-2L Server

Part Number	Description	Repair Category	Repair Procedure
375-3701	6 Gbps SAS RAID disk PCIe HBA, RoHS:Y	BDA-O	1. Power down the server gracefully. See "Powering Down the Server" at http://docs.oracle.com/cd/E23393_01/html/E27229/z400061b1011735.html 2. See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000f91037409.html Note: Do not run any controller commands when replacing the disk.
7010036	DR 3 TB, 3.5-inch, SAS/ 7200, CORAL	HS	See "Servicing Storage Drives and Rear Drives" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000091011460.html Do not run any controller commands when replacing the disk. Note: Rear storage drive 1 might be populated with a remote battery module for the host bus adapter (HBA) card. See "Remove a Rear Storage Drive" at http://docs.oracle.com/cd/E23393_01/html/E27229/bbggjgea.html#scrolltoc
7014391	SAS expander assembly	BDA-O	See "Servicing the SAS Expander Module" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceieidja.html
7020486	DIMM, 8 GB, DDR3, 1600 MHz, 1.3V	BDA-O	See "Servicing the DIMMs" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceicjagi.html#scrolltoc
7018704	DIMM, 16 GB, DDR3, 1600 MHz, 1.35V	BDA-O	See "Servicing the DIMMs" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceicjagi.html#scrolltoc
7071891	LRDIMM, 32GB,DDR3,1600,4Rx4,1. 35V	BDA-O	See "Servicing the DIMMs" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceicjagi.html#scrolltoc

Table B-6 (Cont.) Replacement Parts for a Sun Server X3-2L Server

Part Number	Description	Repair Category	Repair Procedure
7020774	Pregreased CPU heatsink	BDA-O	See "Servicing Processors" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40001d31037480.html#scrolltoc
7026881	CPU, Intel, E5-2660, 2.2 GHz, 8-core	BDA-O	See "Servicing Processors" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40001d31037480.html#scrolltoc
7044130	Power supply, AC, A258, F, 12V, 1000W	HS	See "Servicing Power Supplies" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000091014153.html
7045042	Dual counter rotating fan module, 80x80 mm	HS	See "Servicing Fan Modules" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000091014194.html
7047257	12-slot backplane cable kit with the following: - Disk drive flex cable - Dual backplane power cable - Single backplane power cable - Disk data cable, iPass + HD to iPass+ HD, 600 mm	BDA-O	1. Shut down the rack. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)". 2. See "Servicing the Front and Rear Storage Drive Backplanes" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40007ac1419009.html
7047768	12-slot 3.5" disk cage, including fan door, and left and right ear assembly	BDA-O	1. Shut down the rack. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)". 2. See "Servicing the Front and Rear Storage Drive Backplanes" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40007ac1419009.html
7049382	System board assembly, RoHS:Y	BDA-O	1. Shut down the rack. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)". 2. "Servicing the Motherboard Assembly" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40001d31037512.html

B.4.4 Sun Fire X4270 M2 Parts

[Table B-7](#) (page B-17) lists the replaceable parts for the Sun Fire X4270 M2, which are covered under the Oracle Premier Support for Systems warranty. Unless otherwise noted, the repair procedures are in the *Sun Fire X4270 M2 Server Service Manual*.

Table B-7 Replacement Parts for Sun Fire X4270 M2 Servers

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	1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)." 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
7090170	8 GB USB flash drive	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-7 (Cont.) Replacement Parts for Sun Fire X4270 M2 Servers

Part Number	Description	Repair Category	Repair Procedure
371-4982	6 Gigabit SAS RAID PCI battery module, RoHS:Y (LION), BBU-08	BDA-O	1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)." 2. 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	1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)." 2. See Section 4.5, "Servicing PCIe Cards" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_86021
375-3701	6 Gbps SAS RAID disk PCIe HBA, RoHS:Y	BDA-O	1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)." 2. See Section 4.5, "Servicing PCIe Cards" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_86021 Note: Do not run any controller commands in the service manual when replacing the disk.
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

Table B-7 (Cont.) Replacement Parts for Sun Fire X4270 M2 Servers

Part Number	Description	Repair Category	Repair Procedure
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
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: - x8/x8 switch PCIe riser - bracket	BDA-O	1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)." 2. See Section 4.4, "Servicing PCIe Risers" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_53252

Table B-7 (Cont.) Replacement Parts for Sun Fire X4270 M2 Servers

Part Number	Description	Repair Category	Repair Procedure
541-4081	System board assembly, RoHS:Y	BDA-O	1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" (page 11-2)." 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-7 (Cont.) Replacement Parts for Sun Fire X4270 M2 Servers

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: • 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

B.5 Parts for Sun Datacenter InfiniBand Switch 36

Table B-8 (page B-22) 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 (page B-25)" for information about InfiniBand cables.

Table B-8 Replacement Parts for the Sun Datacenter InfiniBand Switch 36

Part Number	Description	Repair Category	Repair Procedure
7081595	A247A, 12V, 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	1. See "Replacing a Failed InfiniBand Switch (page 11-7)". 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

B.6 Parts for Sun Network QDR InfiniBand Gateway Switch

Table B-9 (page B-23) 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 (page B-25)" for information about InfiniBand cables.

Table B-9 Replacement Parts for the Sun Network QDR InfiniBand Gateway Switch

Part Number	Description	Repair Category	Repair Procedure
7081595	A247A, 12V, 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
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	1. See "Replacing a Failed InfiniBand Switch" (page 11-7)". 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

B.7 Parts for the Cisco Ethernet Switch

Table B-10 (page B-23) 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-10 Replacement Parts for Cisco Catalyst Ethernet

Part Number	Description	Repair Category	Repair
7055121	Cooling fan tray	HS	See Chapter 4 of the <i>Cisco Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> .

Table B-10 (Cont.) Replacement Parts for Cisco Catalyst Ethernet

Part Number	Description	Repair Category	Repair
7055122	Power supply, redundant	HS	See Chapter 4 of the <i>Cisco Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> .
7055123	SNET, WS-C4948E-F	IR	See Chapter 3 of the <i>Cisco Catalyst 4948E and Catalyst 4948E-F Switch Installation Guide</i> .

B.8 Parts for Power Distribution Units

Table B-11 (page B-24) lists the replacement parts for the power distribution units (PDUs), which are covered under the Oracle Premier Support for Systems warranty. The repair procedures references the *Sun Rack II Power Distribution Units User's Guide* at

<http://docs.oracle.com/cd/E19844-01/html/E23956/index.html>

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.

Table B-11 Replacement Parts for Power Distribution Units

Part Number	Description	Repair Category	Repair Procedures
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 "Removing and Replacing PDUs" at http://docs.oracle.com/cd/E19844-01/html/E23956/z4000546165586.html#scrolltoc
371-4976	Three-phase PDU: 15 kVA, 5-pin, IEC309 32A, 5-pin 230/400V, three-phase IP44, international, RoHS:Y	IR	See "Removing and Replacing PDUs" at http://docs.oracle.com/cd/E19844-01/html/E23956/z4000546165586.html#scrolltoc

Table B-11 (Cont.) Replacement Parts for Power Distribution Units

Part Number	Description	Repair Category	Repair Procedures
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 "Removing and Replacing PDUs" at http://docs.oracle.com/cd/E19844-01/html/E23956/z4000546165586.html#scrolltoc
371-5040	Single-phase PDU: 22 kVA, with three blue 32A/240V splash-proof 2-pole/3-wire IEC 60309 plugs, international, RoHS:Y	IR	See "Removing and Replacing PDUs" at http://docs.oracle.com/cd/E19844-01/html/E23956/z4000546165586.html#scrolltoc

B.9 Cables and Cable Accessories for an Oracle Big Data Appliance Rack

Table B-12 (page B-26) 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/gentextid-26783.html#scrolltoc
- *Sun Datacenter InfiniBand Switch 36 Service Manual* at http://docs.oracle.com/cd/E26698_01/html/E26434/gentextid-21904.html#scrolltoc
- 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-12 Replacement Cables for an Oracle Big Data Appliance Rack

Part Number	Description	Repair Category	Repair
350-1287	Slim rail, CMA, 1U-2U for Sun servers	BDA-O	1. Power down the server gracefully. See "Powering Down the Server" at http://docs.oracle.com/cd/E23393_01/html/E27229/z400061b1011735.html 2. See "Install Cable Management Arm" in the <i>Sun Server X3-2L Installation Guide</i> at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000401014951.html
350-1519	Serial cable kit, RoHS:Y with the following: • USB to DB 9-pin M serial cable • DB 9-pin F to DB 9-pin F null modem cable	HS	See "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
530-4567	4x QSFP copper QDR InfiniBand cable, 2 m, RoHS:Y	HS	1. See "Maintaining the InfiniBand Network (page 11-7)". 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 Note: Redundancy is lost during cable replacement.
530-4445	4x QSFP copper QDR InfiniBand cable, 3 m, RoHS:Y	HS	1. See "Maintaining the InfiniBand Network (page 11-7)". 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
530-4446	4x QSFP copper QDR InfiniBand cable, 5 m, RoHS:Y	HS	1. See "Maintaining the InfiniBand Network (page 11-7)". 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-12 (Cont.) Replacement Cables for an Oracle Big Data Appliance Rack

Part Number	Description	Repair Category	Repair
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> .

C

In-Rack Cabling Tables for Oracle Big Data Appliance X6–2 and X5-2

The tables in this appendix show the in-rack cable connections for Oracle Big Data Appliance X6-2 and X5-2 racks. This appendix contains the following sections:

- [Cable Color Coding](#) (page C-1)
- [Oracle Big Data Appliance Rack Layout](#) (page C-1)
- [Administrative Gigabit Ethernet Cable Connections](#) (page C-4)
- [Oracle Integrated Lights Out Manager Cable Connections](#) (page C-5)
- [Single-Phase Power Distribution Unit Cable Connections](#) (page C-6)
- [Three-Phase Power Distribution Unit Cable Connections](#) (page C-7)
- [In-Rack InfiniBand Switch-to-Switch Cable Connections](#) (page C-8)
- [In-Rack InfiniBand Switch-to-Server Cable Connections](#) (page C-9)

C.1 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

C.2 Oracle Big Data Appliance Rack Layout

The following figures show the front and rear layouts for three configurations of Oracle Big Data Appliance X6-2 and X5-2: full rack, starter rack, and starter rack with High Capacity nodes plus InfiniBand Infrastructure.

[Figure C-1](#) (page C-2) shows the layout of a full rack for Oracle Big Data Appliance X6-2 and X5-2.

Figure C-1 Oracle Big Data Appliance X6-2 and X5-2 Full Rack Layout

Rear View					Front View				
U#	PDU A	Big Data Appliance	PDU B	U#	U#	PDU B	Big Data Appliance	PDU A	U#
42	INP2	1U Solid Filler	INP0	42	42	INP0	1U Solid Filler	INP2	42
41		1U Solid Filler		41	41		1U Solid Filler		41
40		BDA Server 18		40	40		BDA Server 18		40
39		BDA Server 17		39	39		BDA Server 17		39
38		BDA Server 16		38	38		BDA Server 16		38
37		BDA Server 15		37	37		BDA Server 15		37
36		BDA Server 14		36	36		BDA Server 14		36
35		BDA Server 13		35	35		BDA Server 13		35
34		BDA Server 12		34	34		BDA Server 12		34
33		BDA Server 11		33	33		BDA Server 11		33
32	INP1	BDA Server 10	INP1	32	32	INP1	BDA Server 10	INP1	32
31		Sun Network QDR Infiniband GW Switch		31	31		1U Vented Filler		31
30		Cisco Ethernet "Admin" Switch		30	30		1U Vented Filler		30
29		Sun Network QDR Infiniband GW Switch		29	29		1U Vented Filler		29
28		BDA Server 9		28	28		BDA Server 9		28
27		BDA Server 8		27	27		BDA Server 8		27
26		BDA Server 7		26	26		BDA Server 7		26
25		BDA Server 6		25	25		BDA Server 6		25
24		BDA Server 5		24	24		BDA Server 5		24
23		BDA Server 4		23	23		BDA Server 4		23
22	INP0	BDA Server 3	INP2	22	22	INP2	BDA Server 3	INP0	22
21		BDA Server 2		21	21		BDA Server 2		21
20		BDA Server 1		20	20		BDA Server 1		20
19		Sun Datacenter QDR Infiniband Switch		19	19		1U Vented Filler		19
18		Power Distribution Unit-A		18	18		Power Distribution Unit-A		18
17		Power Distribution Unit-B		17	17		Power Distribution Unit-B		17
16				16	16				16
15				15	15				15
14				14	14				14
13				13	13				13
12				12	12				12
11				11	11				11
10				10	10				10
9				9	9				9
8				8	8				8
7				7	7				7
6				6	6				6
5				5	5				5
4				4	4				4
3				3	3				3
2				2	2				2
1				1	1				1

Figure C-2 (page C-3) shows the layout of a starter rack for Oracle Big Data Appliance X6-2 or X5-2.

Figure C-2 Oracle Big Data Appliance X6-2 and X5-2 Starter Rack Layout

Rear View					Front View				
U#	PDU A	Big Data Appliance	PDU B	U#	U#	PDU B	Big Data Appliance	PDU A	U#
42	INP2	1U Solid Filler	INP0	42	42	INP0	1U Solid Filler	INP2	42
41		1U Solid Filler		41	41		1U Solid Filler		41
40		1U Solid Filler		40	40		1U Solid Filler		40
39		1U Solid Filler		39	39		1U Solid Filler		39
38		1U Solid Filler		38	38		1U Solid Filler		38
37		1U Solid Filler		37	37		1U Solid Filler		37
36		1U Solid Filler		36	36		1U Solid Filler		36
35		1U Solid Filler		35	35		1U Solid Filler		35
34		1U Solid Filler		34	34		1U Solid Filler		34
33		1U Solid Filler		33	33		1U Solid Filler		33
32		1U Solid Filler		32	32		1U Solid Filler		32
31		1U Solid Filler		31	31		1U Solid Filler		31
30	INP1	1U Solid Filler	INP1	30	30	INP1	1U Solid Filler	INP1	30
29		1U Solid Filler		29	29		1U Solid Filler		29
28		1U Solid Filler		28	28		1U Solid Filler		28
27		1U Solid Filler		27	27		1U Solid Filler		27
26		1U Solid Filler		26	26		1U Solid Filler		26
25		1U Solid Filler		25	25		1U Solid Filler		25
24		1U Solid Filler		24	24		1U Solid Filler		24
23		1U Solid Filler		23	23		1U Solid Filler		23
22		Sun Network QDR Infiniband GW Switch		22	22		1U Vented Filler		22
21		Cisco Ethernet "Admin" Switch		21	21		1U Vented Filler		21
20		Sun Network QDR Infiniband GW Switch		20	20		1U Vented Filler		20
19	INP0	1U Solid Filler	INP2	19	19	INP2	1U Solid Filler	INP0	19
18		1U Solid Filler		18	18		1U Solid Filler		18
17		1U Solid Filler		17	17		1U Solid Filler		17
16		1U Solid Filler		16	16		1U Solid Filler		16
15		1U Solid Filler		15	15		1U Solid Filler		15
14		1U Solid Filler		14	14		1U Solid Filler		14
13		BDA Server 6		13	13		BDA Server 6		13
12		BDA Server 5		12	12		BDA Server 5		12
11		BDA Server 4		11	11		BDA Server 4		11
10		BDA Server 3		10	10		BDA Server 3		10
9		BDA Server 2		9	9		BDA Server 2		9
8		BDA Server 1		8	8		BDA Server 1		8
7		Sun Datacenter QDR Infiniband Switch		7	7		1U Vented Filler		7
6		Power Distribution Unit-A		6	6		Power Distribution Unit-A		6
5		Power Distribution Unit-B		5	5		Power Distribution Unit-B		5
4				4	4				4
3				3	3				3
2				2	2				2
1				1	1				1

Figure C-3 (page C-4) shows the layout of an Oracle Big Data Appliance X6-2 or X5-2 starter rack with High Capacity Nodes plus InfiniBand Infrastructure.

Figure C-3 Oracle Big Data Appliance X6-2 and X5-2 Starter Rack With Big Data Appliance High Capacity Nodes Plus InfiniBand Infrastructure

Rear View					Front View				
U#	PDU A	Big Data Appliance	PDU B	U#	U#	PDU B	Big Data Appliance	PDU A	U#
42	INP2	1U Solid Filler	INP0	42	42	INP0	1U Solid Filler	INP2	42
41		1U Solid Filler		41	41		1U Solid Filler		41
40		1U Solid Filler		40	40		1U Solid Filler		40
39		1U Solid Filler		39	39		1U Solid Filler		39
38		1U Solid Filler		38	38		1U Solid Filler		38
37		1U Solid Filler		37	37		1U Solid Filler		37
36		1U Solid Filler		36	36		1U Solid Filler		36
35		1U Solid Filler		35	35		1U Solid Filler		35
34		1U Solid Filler		34	34		1U Solid Filler		34
33		1U Solid Filler		33	33		1U Solid Filler		33
32		1U Solid Filler		32	32		1U Solid Filler		32
31		1U Solid Filler		31	31		1U Solid Filler		31
30	INP1	1U Solid Filler	INP1	30	30	INP1	1U Solid Filler	INP1	30
29		1U Solid Filler		29	29		1U Solid Filler		29
28		BDA Server 12		28	28		BDA Server 12		28
27		BDA Server 11		27	27		BDA Server 11		27
26		BDA Server 10		26	26		BDA Server 10		26
25		BDA Server 10		25	25		BDA Server 10		25
24		BDA Server 10		24	24		BDA Server 10		24
23		Sun Network QDR Infiniband GW Switch		23	23		1U Vented Filler		23
22		Cisco Ethernet "Admin" Switch		22	22		1U Vented Filler		22
21		Sun Network QDR Infiniband GW Switch		21	21		1U Vented Filler		21
20		BDA Server 9		20	20		BDA Server 9		20
19	INP0	BDA Server 8	INP2	19	19	INP2	BDA Server 8	INP0	19
18		BDA Server 7		18	18		BDA Server 7		18
17		BDA Server 6		17	17		BDA Server 6		17
16		BDA Server 5		16	16		BDA Server 5		16
15		BDA Server 4		15	15		BDA Server 4		15
14		BDA Server 3		14	14		BDA Server 3		14
13		BDA Server 2		13	13		BDA Server 2		13
12		BDA Server 1		12	12		BDA Server 1		12
11		Sun Datacenter QDR Infiniband Switch		11	11		1U Vented Filler		11
10		Power Distribution Unit-A		10	10		Power Distribution Unit-A		10
9		Power Distribution Unit-B		9	9		Power Distribution Unit-B		9
8				8	8				8
7				7	7				7
6				6	6				6
5				5	5				5
4				4	4				4
3				3	3				3
2				2	2				2
1				1	1				1

C.3 Administrative Gigabit Ethernet Cable Connections

Table C-1 (page C-5) 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, with the exception of the PDU cables, which are 1 meter long and white.

Table C-1 Administrative Gigabit Ethernet Cabling

From Rack Unit ¹	Component	To Ethernet Port
U39	Server 18	1
U37	Server 17	3
U35	Server 16	5
U33	Server 15	7
U31	Server 14	9
U29	Server 13	11
U27	Server 12	13
U25	Server 11	17
U23	Server 10	21
U22	NM2-IB Switch	45
U20	NM2-IB Switch	46
U18	Server 9	25
U16	Server 8	29
U14	Server 7	31
U12	Server 6	33
U10	Server 5	35
U8	Server 4	37
U6	Server 3	39
U4	Server 2	41
U2	Server 1	43
U1	NM2-IB Switch	47
PDU-A	PDU-A	15
PDU-B	PDU-B	19

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

C.4 Oracle Integrated Lights Out Manager Cable Connections

[Table C-2](#) (page C-5) 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-2 Oracle ILOM Cabling

From Rack Unit ¹	Component	Ethernet Port	Cable Length	Cable Color
U39	Server 18	2	10 feet	Red
U37	Server 17	4	10 feet	Red
U35	Server 16	6	10 feet	Red

Table C-2 (Cont.) Oracle ILOM Cabling

From Rack Unit ¹	Component	Ethernet Port	Cable Length	Cable Color
U33	Server 15	8	10 feet	Red
U31	Server 14	10	10 feet	Red
U29	Server 13	12	10 feet	Red
U27	Server 12	14	10 feet	Red
U25	Server 11	18	10 feet	Red
U23	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
Not applicable	Service	48	7 feet	Blue

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

C.5 Single-Phase Power Distribution Unit Cable Connections

[Table C-3](#) (page C-6) 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-3 Single-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U39	G5-6	G0-0	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-0	G0-6	2 meters
U33	G4-6	G1-0	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G3-6	G2-0	2 meters
U25	G3-5	G2-1	2 meters
U23	G3-3	G2-3	2 meters

Table C-3 (Cont.) Single-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U22	G3-1	G2-5	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.

C.6 Three-Phase Power Distribution Unit Cable Connections

[Table C-4](#) (page C-7) 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-4 Three-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U39	G5-6	G2-0	2 meters
U37	G5-3	G2-3	2 meters
U35	G5-0	G2-6	2 meters
U33	G4-6	G1-0	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G3-6	G0-0	2 meters
U25	G3-5	G0-1	2 meters
U23	G3-3	G0-3	2 meters
U22	G5-1	G2-5	2 meters
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

Table C-4 (Cont.) Three-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
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.

C.7 In-Rack InfiniBand Switch-to-Switch Cable Connections

Table C-5 (page C-8) lists the location, ports, and cables for the InfiniBand connections from switch to switch.


Table C-5 InfiniBand Switch-to-Switch Cabling

From InfiniBand Switch Rack Unit ¹	Port	To InfiniBand Switch Rack Unit ¹	Port	Cable Length	Description
U20	9B	U22	9A	2 meters	Leaf to leaf
U20	10B	U22	10A	2 meters	Leaf to leaf
U20	11B	U22	11A	2 meters	Leaf to leaf
U20	8A	U22	8A	2 meters	Leaf to leaf
U20	9A	U22	9B	2 meters	Leaf to leaf
U20	10A	U22	10B	2 meters	Leaf to leaf
U20	11A	U22	11B	2 meters	Leaf to leaf
U1	1B	U20	8B	3 meters	Spine to leaf
U1	0B	U22	8B	3 meters	Spine to leaf

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

Figure C-4 (page C-9) identifies the locations of the ports on a Sun Network QDR InfiniBand Gateway switch. Ports 5A, 5B, 6A, 6B, and 7A are available to connect to another device, such as Oracle Exalytics In-Memory Machine. Oracle recommends that you start with port 5A. Port 7B is reserved for future use by Oracle Big Data Appliance.

Figure C-4 Sun Network QDR InfiniBand Gateway Ports



C.8 In-Rack InfiniBand Switch-to-Server Cable Connections

Table C-6 (page C-9) lists the location, ports, and cables for the InfiniBand connections from the leaf switches to the servers.

Table C-6 InfiniBand Switch-to-Server Cabling

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit ¹	Port ²	Cable Length
U22	0A	U39	PCIe 3 P1	3 meters
U22	0B	U37	PCIe 3 P1	3 meters
U22	1A	U35	PCIe 3 P1	3 meters
U22	1B	U33	PCIe 3 P1	3 meters
U22	2A	U31	PCIe 3 P1	3 meters
U22	2B	U29	PCIe 3 P1	3 meters
U22	3A	U27	PCIe 3 P1	3 meters
U22	3B	U25	PCIe 3 P1	3 meters
U22	4A	U23	PCIe 3 P1	3 meters
U20	0A	U39	PCIe 3 P2	3 meters
U20	0B	U37	PCIe 3 P2	3 meters
U20	1A	U35	PCIe 3 P2	3 meters
U20	1B	U33	PCIe 3 P2	3 meters
U20	2A	U31	PCIe 3 P2	3 meters
U20	2B	U29	PCIe 3 P2	3 meters
U20	3A	U27	PCIe 3 P2	3 meters
U20	3B	U25	PCIe 3 P2	3 meters
U20	4A	U23	PCIe 3 P2	3 meters
U20	4B	U18	PCIe 3 P1	3 meters

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

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit ¹	Port ²	Cable Length
U20	12A	U16	PCle 3 P1	3 meters
U20	12B	U14	PCle 3 P1	3 meters
U20	13A	U12	PCle 3 P1	3 meters
U20	13B	U10	PCle 3 P1	3 meters
U20	14A	U8	PCle 3 P1	3 meters
U20	14B	U6	PCle 3 P1	3 meters
U20	15A	U4	PCle 3 P1	3 meters
U20	15B	U2	PCle 3 P1	3 meters
U22	4B	U18	PCle 3 P2	3 meters
U22	12A	U16	PCle 3 P2	3 meters
U22	12B	U14	PCle 3 P2	3 meters
U22	13A	U12	PCle 3 P2	3 meters
U22	13B	U10	PCle 3 P2	3 meters
U22	14A	U8	PCle 3 P2	3 meters
U22	14B	U6	PCle 3 P2	3 meters
U22	15A	U4	PCle 3 P2	3 meters
U22	15B	U2	PCle 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

In-Rack Cabling Tables for Oracle Big Data Appliance X4-2 and X3-2

The tables in this appendix show the in-rack cable connections for Oracle Big Data Appliance X4-2 and Oracle Big Data Appliance X3-2 racks. This appendix contains the following sections:

- [Cable Color Coding](#) (page D-1)
- [Oracle Big Data Appliance Rack Layout](#) (page D-1)
- [Administrative Gigabit Ethernet Cable Connections](#) (page D-4)
- [Oracle Integrated Lights Out Manager Cable Connections](#) (page D-5)
- [Single-Phase Power Distribution Unit Cable Connections](#) (page D-6)
- [Three-Phase Power Distribution Unit Cable Connections](#) (page D-7)
- [In-Rack InfiniBand Switch-to-Switch Cable Connections](#) (page D-8)
- [In-Rack InfiniBand Switch-to-Server Cable Connections](#) (page D-9)

D.1 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

D.2 Oracle Big Data Appliance Rack Layout

The following figures show the front and rear layouts for three configurations of Oracle Big Data Appliance X4-2 and X3-2: full rack, starter rack, and starter rack with High Capacity nodes plus InfiniBand Infrastructure.

[Figure D-1](#) (page D-2) shows the layout of a full rack for Oracle Big Data Appliance X4-2 or Oracle Big Data Appliance X3-2.

Figure D-1 Oracle Big Data Appliance X4-2 Full Rack Layout

Rear View					Front View				
U#	PDU A	Big Data Appliance	PDU B	U#	U#	PDU B	Big Data Appliance	PDU A	U#
42	INP2	1U Solid Filler	INP0	42	42	INP0	1U Solid Filler	INP2	42
41		1U Solid Filler		41	41		1U Solid Filler		41
40		BDA Server 18		40	40		BDA Server 18		40
39		BDA Server 17		39	39		BDA Server 17		39
38		BDA Server 16		38	38		BDA Server 16		38
37		BDA Server 15		37	37		BDA Server 15		37
36		BDA Server 14		36	36		BDA Server 14		36
35		BDA Server 13		35	35		BDA Server 13		35
34		BDA Server 12		34	34		BDA Server 12		34
33		BDA Server 11		33	33		BDA Server 11		33
28	INP1	BDA Server 10	INP1	28	28	INP1	BDA Server 10	INP1	28
27		Sun Network QDR Infiniband GW Switch		27	27		1U Vented Filler		27
26		Cisco Ethernet "Admin" Switch		26	26		1U Vented Filler		26
25		Sun Network QDR Infiniband GW Switch		25	25		1U Vented Filler		25
24		BDA Server 9		24	24		BDA Server 9		24
23		BDA Server 8		23	23		BDA Server 8		23
22		BDA Server 7		22	22		BDA Server 7		22
21		BDA Server 6		21	21		BDA Server 6		21
20		BDA Server 5		20	20		BDA Server 5		20
19		BDA Server 4		19	19		BDA Server 4		19
18	INP0	BDA Server 3	INP2	18	18	INP2	BDA Server 3	INP0	18
17		BDA Server 2		17	17		BDA Server 2		17
16		BDA Server 1		16	16		BDA Server 1		16
15		Sun Datacenter QDR Infiniband Switch		15	15		1U Vented Filler		15
14		Power Distribution Unit-A		14	14		Power Distribution Unit-A		14
13		Power Distribution Unit-B		13	13		Power Distribution Unit-B		13
12				12	12				12
11				11	11				11
10				10	10				10
9				9	9				9
8				8	8				8
7				7	7				7
6				6	6				6
5				5	5				5
4				4	4				4
3				3	3				3
2				2	2				2
1				1	1				1

Figure D-2 (page D-3) shows the layout of a starter rack for Oracle Big Data Appliance X4-2 or Oracle Big Data Appliance X3-2.

Figure D-2 Oracle Big Data Appliance X4-2 Starter Rack Layout

Rear View					Front View				
U#	PDU A	Big Data Appliance	PDU B	U#	U#	PDU B	Big Data Appliance	PDU A	U#
42	INP2	1U Solid Filler	INP0	42	42	INP0	1U Solid Filler	INP2	42
41		1U Solid Filler		41	41		1U Solid Filler		41
40		1U Solid Filler		40	40		1U Solid Filler		40
39		1U Solid Filler		39	39		1U Solid Filler		39
38		1U Solid Filler		38	38		1U Solid Filler		38
37		1U Solid Filler		37	37		1U Solid Filler		37
36		1U Solid Filler		36	36		1U Solid Filler		36
35		1U Solid Filler		35	35		1U Solid Filler		35
34		1U Solid Filler		34	34		1U Solid Filler		34
33		1U Solid Filler		33	33		1U Solid Filler		33
32		1U Solid Filler		32	32		1U Solid Filler		32
31		1U Solid Filler		31	31		1U Solid Filler		31
30	INP1	1U Solid Filler	INP1	30	30	INP1	1U Solid Filler	INP1	30
29		1U Solid Filler		29	29		1U Solid Filler		29
28		1U Solid Filler		28	28		1U Solid Filler		28
27		1U Solid Filler		27	27		1U Solid Filler		27
26		1U Solid Filler		26	26		1U Solid Filler		26
25		1U Solid Filler		25	25		1U Solid Filler		25
24		1U Solid Filler		24	24		1U Solid Filler		24
23		1U Solid Filler		23	23		1U Solid Filler		23
22		Sun Network QDR Infiniband GW Switch		22	22		1U Vented Filler		22
21		Cisco Ethernet "Admin" Switch		21	21		1U Vented Filler		21
20		Sun Network QDR Infiniband GW Switch		20	20		1U Vented Filler		20
19	INP0	1U Solid Filler	INP2	19	19	INP2	1U Solid Filler	INP0	19
18		1U Solid Filler		18	18		1U Solid Filler		18
17		1U Solid Filler		17	17		1U Solid Filler		17
16		1U Solid Filler		16	16		1U Solid Filler		16
15		1U Solid Filler		15	15		1U Solid Filler		15
14		1U Solid Filler		14	14		1U Solid Filler		14
13		BDA Server 6		13	13		BDA Server 6		13
12		BDA Server 5		12	12		BDA Server 5		12
11		BDA Server 4		11	11		BDA Server 4		11
10		BDA Server 3		10	10		BDA Server 3		10
9		BDA Server 2		9	9		BDA Server 2		9
8		BDA Server 1		8	8		BDA Server 1		8
7		Sun Datacenter QDR Infiniband Switch		7	7		1U Vented Filler		7
6		Power Distribution Unit-A		6	6		Power Distribution Unit-A		6
5		Power Distribution Unit-B		5	5		Power Distribution Unit-B		5
4				4	4				4
3				3	3				3
2				2	2				2
1				1	1				1

Figure D-3 (page D-4) shows the layout of an Oracle Big Data Appliance X4-2 or Oracle Big Data Appliance X3-2 starter rack with High Capacity Nodes plus InfiniBand Infrastructure.

Figure D-3 Oracle Big Data Appliance X4-2 Starter Rack With Big Data Appliance X5-2 High Capacity Nodes Plus InfiniBand Infrastructure

Rear View					Front View				
U#	PDU A	Big Data Appliance	PDU B	U#	U#	PDU B	Big Data Appliance	PDU A	U#
42	INP2	1U Solid Filler	INP0	42	42	INP0	1U Solid Filler	INP2	42
41		1U Solid Filler		41	41		1U Solid Filler		41
40		1U Solid Filler		40	40		1U Solid Filler		40
39		1U Solid Filler		39	39		1U Solid Filler		39
38		1U Solid Filler		38	38		1U Solid Filler		38
37		1U Solid Filler		37	37		1U Solid Filler		37
36		1U Solid Filler		36	36		1U Solid Filler		36
35		1U Solid Filler		35	35		1U Solid Filler		35
34		1U Solid Filler		34	34		1U Solid Filler		34
33		1U Solid Filler		33	33		1U Solid Filler		33
32		1U Solid Filler		32	32		1U Solid Filler		32
31		1U Solid Filler		31	31		1U Solid Filler		31
30	INP1	1U Solid Filler	INP1	30	30	INP1	1U Solid Filler	INP1	30
29		1U Solid Filler		29	29		1U Solid Filler		29
28		BDA Server 12		28	28		BDA Server 12		28
27		BDA Server 11		27	27		BDA Server 11		27
26		BDA Server 10		26	26		BDA Server 10		26
25		BDA Server 10		25	25		BDA Server 10		25
24		BDA Server 10		24	24		BDA Server 10		24
23		Sun Network QDR Infiniband GW Switch		23	23		1U Vented Filler		23
22		Cisco Ethernet "Admin" Switch		22	22		1U Vented Filler		22
21		Sun Network QDR Infiniband GW Switch		21	21		1U Vented Filler		21
20		BDA Server 9		20	20		BDA Server 9		20
19	INP0	BDA Server 8	INP2	19	19	INP2	BDA Server 8	INP0	19
18		BDA Server 7		18	18		BDA Server 7		18
17		BDA Server 6		17	17		BDA Server 6		17
16		BDA Server 5		16	16		BDA Server 5		16
15		BDA Server 4		15	15		BDA Server 4		15
14		BDA Server 3		14	14		BDA Server 3		14
13		BDA Server 2		13	13		BDA Server 2		13
12		BDA Server 1		12	12		BDA Server 1		12
11		Sun Datacenter QDR Infiniband Switch		11	11		1U Vented Filler		11
10		Power Distribution Unit-A		10	10		Power Distribution Unit-A		10
9		Power Distribution Unit-B		9	9		Power Distribution Unit-B		9
8				8	8				8
7				7	7				7
6				6	6				6
5				5	5				5
4				4	4				4
3				3	3				3
2				2	2				2
1				1	1				1

D.3 Administrative Gigabit Ethernet Cable Connections

Table D-1 (page D-5) 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 D-1 Administrative Gigabit Ethernet Cabling

From Rack Unit ¹	Component	To Ethernet Port
U39	Server 18	1
U37	Server 17	3
U35	Server 16	5
U33	Server 15	7
U31	Server 14	9
U29	Server 13	11
U27	Server 12	13
U25	Server 11	17
U23	Server 10	21
U22	NM2-IB Switch	45
U20	NM2-IB Switch	46
U18	Server 9	25
U16	Server 8	29
U14	Server 7	31
U12	Server 6	33
U10	Server 5	35
U8	Server 4	37
U6	Server 3	39
U4	Server 2	41
U2	Server 1	43
U1	NM2-IB Switch	47

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

D.4 Oracle Integrated Lights Out Manager Cable Connections

[Table D-2](#) (page D-5) 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 D-2 Oracle ILOM Cabling

From Rack Unit ¹	Component	Ethernet Port	Cable Length	Cable Color
U39	Server 18	2	10 feet	Red
U37	Server 17	4	10 feet	Red
U35	Server 16	6	10 feet	Red
U33	Server 15	8	10 feet	Red
U31	Server 14	10	10 feet	Red

Table D-2 (Cont.) Oracle ILOM Cabling

From Rack Unit ¹	Component	Ethernet Port	Cable Length	Cable Color
U29	Server 13	12	10 feet	Red
U27	Server 12	14	10 feet	Red
U25	Server 11	18	10 feet	Red
U23	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	19	1 meter	White
Not applicable	Service	48	10 feet	Blue

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

D.5 Single-Phase Power Distribution Unit Cable Connections

[Table D-3](#) (page D-6) 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 D-3 Single-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U39	G5-6	G0-0	2 meters
U37	G5-3	G0-3	2 meters
U35	G5-0	G0-6	2 meters
U33	G4-6	G1-0	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G3-6	G2-0	2 meters
U25	G3-5	G2-1	2 meters
U23	G3-3	G2-3	2 meters

Table D-3 (Cont.) Single-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
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.

D.6 Three-Phase Power Distribution Unit Cable Connections

[Table D-4](#) (page D-7) 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 D-4 Three-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U39	G5-6	G2-0	2 meters
U37	G5-3	G2-3	2 meters
U35	G5-0	G2-6	2 meters
U33	G4-6	G1-0	2 meters
U31	G4-4	G1-2	2 meters
U29	G4-2	G1-4	2 meters
U27	G3-6	G0-0	2 meters
U25	G3-5	G0-1	2 meters
U23	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

Table D-4 (Cont.) Three-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
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.

D.7 In-Rack InfiniBand Switch-to-Switch Cable Connections

Table D-5 (page D-8) lists the location, ports, and cables for the InfiniBand connections from switch to switch.


Table D-5 InfiniBand Switch-to-Switch Cabling

From InfiniBand Switch Rack Unit ¹	Port	To InfiniBand Switch Rack Unit ¹	Port	Cable Length	Description
U20	9B	U22	9A	2 meters	Leaf to leaf
U20	10B	U22	10A	2 meters	Leaf to leaf
U20	11B	U22	11A	2 meters	Leaf to leaf
U20	8A	U22	8A	2 meters	Leaf to leaf
U20	9A	U22	9B	2 meters	Leaf to leaf
U20	10A	U22	10B	2 meters	Leaf to leaf
U20	11A	U22	11B	2 meters	Leaf to leaf
U1	0B	U20	8B	3 meters	Spine to leaf
U1	1B	U22	8B	3 meters	Spine to leaf

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

Figure D-4 (page D-9) identifies the locations of the ports on a Sun Network QDR InfiniBand Gateway switch. Ports 5A, 5B, 6A, 6B, and 7A are available to connect to another device, such as Oracle Exalytics In-Memory Machine. Oracle recommends that you start with port 5A. Port 7B is reserved for future use by Oracle Big Data Appliance.

Figure D-4 Sun Network QDR InfiniBand Gateway Ports



D.8 In-Rack InfiniBand Switch-to-Server Cable Connections

Table D-6 (page D-9) lists the location, ports, and cables for the InfiniBand connections from the leaf switches to the servers.

Table D-6 InfiniBand Switch-to-Server Cabling

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit ¹	Port ²	Cable Length
U22	0A	U39	PCle 3 P1	3 meters
U22	0B	U37	PCle 3 P1	3 meters
U22	1A	U35	PCle 3 P1	3 meters
U22	1B	U33	PCle 3 P1	3 meters
U22	2A	U31	PCle 3 P1	3 meters
U22	2B	U29	PCle 3 P1	3 meters
U22	3A	U27	PCle 3 P1	3 meters
U22	3B	U25	PCle 3 P1	2 meters
U22	4A	U23	PCle 3 P1	2 meters
U20	0A	U39	PCle 3 P2	3 meters
U20	0B	U37	PCle 3 P2	3 meters
U20	1A	U35	PCle 3 P2	3 meters
U20	1B	U33	PCle 3 P2	3 meters
U20	2A	U31	PCle 3 P2	3 meters
U20	2B	U29	PCle 3 P2	3 meters
U20	3A	U27	PCle 3 P2	3 meters
U20	3B	U25	PCle 3 P2	2 meters
U20	4A	U23	PCle 3 P2	2 meters
U20	4B	U18	PCle 3 P1	2 meters

Table D-6 (Cont.) InfiniBand Switch-to-Server Cabling

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit ¹	Port ²	Cable Length
U20	12A	U16	PCle 3 P1	2 meters
U20	12B	U14	PCle 3 P1	3 meters
U20	13A	U12	PCle 3 P1	3 meters
U20	13B	U10	PCle 3 P1	3 meters
U20	14A	U8	PCle 3 P1	3 meters
U20	14B	U6	PCle 3 P1	3 meters
U20	15A	U4	PCle 3 P1	3 meters
U20	15B	U2	PCle 3 P1	3 meters
U22	4B	U18	PCle 3 P2	2 meters
U22	12A	U16	PCle 3 P2	2 meters
U22	12B	U14	PCle 3 P2	3 meters
U22	13A	U12	PCle 3 P2	3 meters
U22	13B	U10	PCle 3 P2	3 meters
U22	14A	U8	PCle 3 P2	3 meters
U22	14B	U6	PCle 3 P2	3 meters
U22	15A	U4	PCle 3 P2	3 meters
U22	15B	U2	PCle 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.

E

In-Rack Cabling Tables for Oracle Big Data Appliance Sun Fire X4270 M2

The tables in this appendix show the in-rack cable connections for Sun Fire X4270 M2-based Oracle Big Data Appliance racks. This appendix contains the following sections:

- [Cable Color Coding](#) (page E-1)
- [Oracle Big Data Appliance Rack Layout](#) (page E-1)
- [KVM Network Cable Connections](#) (page E-2)
- [Administrative Gigabit Ethernet Cable Connections](#) (page E-3)
- [Oracle Integrated Lights Out Manager Cable Connections](#) (page E-4)
- [Single-Phase Power Distribution Unit Cable Connections](#) (page E-5)
- [Three-Phase Power Distribution Unit Cable Connections](#) (page E-6)
- [In-Rack InfiniBand Switch-to-Switch Cable Connections](#) (page E-7)
- [In-Rack InfiniBand Switch-to-Server Cable Connections](#) (page E-8)

E.1 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

E.2 Oracle Big Data Appliance Rack Layout

[Figure E-1](#) (page E-2) shows the front and rear layout of an Oracle Big Data Appliance rack.

Figure E-1 Rack Layout of Oracle Big Data Appliance

Rear View					Front View				
U#	PDU A	Big Data Appliance	PDU B	U#	U#	PDU B	Big Data Appliance	PDU A	U#
42	INP2	BDA Server 18	INP0	42	42	INP0	BDA Server 18	INP2	42
41				41	41				41
40		BDA Server 17		40	40		BDA Server 17		40
39				39	39				39
38		BDA Server 16		38	38		BDA Server 16		38
37				37	37				37
36		BDA Server 15		36	36		BDA Server 15		36
35				35	35				35
34		BDA Server 14		34	34		BDA Server 14		34
33				33	33				33
32	INP1	BDA Server 13	INP1	32	32	INP1	BDA Server 13	INP1	32
31				31	31				31
30		BDA Server 12		30	30		BDA Server 12		30
29				29	29				29
28		BDA Server 11		28	28		BDA Server 11		28
27				27	27				27
26		BDA Server 10		26	26		BDA Server 10		26
25				25	25				25
24		Sun Network QDR Infiniband GW Switch		24	24		1U Vented Filler		24
23		KVM LCD		23	23		KVM LCD		23
22		KVM Switch		22	22		1U Vented Filler		22
21		Cisco Ethernet "Admin" Switch		21	21		1U Vented Filler		21
20	INP0	Sun Network QDR Infiniband GW Switch	INP2	20	20	INP2	1U Vented Filler		20
19		BDA Server 9		19	19		BDA Server 9		19
18				18	18				18
17		BDA Server 8		17	17		BDA Server 8		17
16				16	16				16
15		BDA Server 7		15	15		BDA Server 7		15
14				14	14				14
13		BDA Server 6		13	13		BDA Server 6		13
12				12	12				12
11		BDA Server 5		11	11		BDA Server 5		11
10	INP0			10	10				10
9		BDA Server 4		9	9		BDA Server 4		9
8				8	8				8
7		BDA Server 3		7	7		BDA Server 3		7
6				6	6				6
5		BDA Server 2		5	5		BDA Server 2		5
4				4	4				4
3		BDA Server 1		3	3		BDA Server 1		3
2				2	2				2
1		Sun Datacenter QDR Infiniband Switch		1	1		1U Vented Filler		1
		Power Distribution Unit-A					Power Distribution Unit-A		
		Power Distribution Unit-B					Power Distribution Unit-B		

E.3 KVM Network Cable Connections

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

[Table E-1](#) (page E-3) shows the keyboard, video, and mouse (KVM) network cabling to the servers. The cables are orange.

Table E-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
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

E.4 Administrative Gigabit Ethernet Cable Connections

[Table E-2](#) (page E-3) 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 E-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

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

From Rack Unit ¹	Component	To Ethernet Port
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
U14	Server8	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.

E.5 Oracle Integrated Lights Out Manager Cable Connections

[Table E-3](#) (page E-4) 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 E-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

Table E-3 (Cont.) Oracle ILOM Cabling

From Rack Unit ¹	Component	Ethernet Port	Cable Length	Cable Color
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.

E.6 Single-Phase Power Distribution Unit Cable Connections

[Table E-4](#) (page E-5) 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 E-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

Table E-4 (Cont.) Single-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U14	G2-0	G3-6	2 meters
U12	G1-4	G4-2	2 meters
U10	G1-2	G4-4	2 meters
U8	G1-0	G4-6	2 meters
U6	G0-6	G5-0	2 meters
U4	G0-4	G5-2	2 meters
U2	G0-2	G5-4	2 meters
U1	G0-0	G5-6	2 meters

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

E.7 Three-Phase Power Distribution Unit Cable Connections

Table E-5 (page E-6) 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 E-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

Table E-5 (Cont.) Three-Phase PDU Cabling

Rack Unit ¹	PDU-A/PS-00	PDU-B/PS-01	Cable Length
U8	G1-0	G4-6	2 meters
U6	G0-6	G3-0	2 meters
U4	G0-4	G3-2	2 meters
U2	G0-2	G3-4	2 meters
U1	G0-0	G3-6	2 meters

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

E.8 In-Rack InfiniBand Switch-to-Switch Cable Connections

Table E-6 (page E-7) lists the location, ports, and cables for the InfiniBand connections from switch to switch. Figure E-2 (page E-8) identifies the locations of the ports on a Sun Network QDR InfiniBand Gateway switch.


Table E-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
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 E-2 (page E-8) identifies the locations of the ports on a Sun Network QDR InfiniBand Gateway switch. Ports 5A, 5B, 6A, 6B, and 7A are available to connect to another device, such as Oracle Exalytics In-Memory Machine. Oracle recommends that you start with port 5A. Port 7B is reserved for future use by Oracle Big Data Appliance.

Figure E-2 Sun Network QDR InfiniBand Gateway Switch Ports



E.9 In-Rack InfiniBand Switch-to-Server Cable Connections

Table E-7 (page E-8) lists the location, ports, and cables for the InfiniBand connections from the leaf switches to the servers.

Table E-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
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

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

From InfiniBand Switch Rack Unit ¹	Port	To Rack Unit ¹	Port ²	Cable Length
U20	12A	U16	PCle 3 P1	2 meters
U20	12B	U14	PCle 3 P1	3 meters
U20	13A	U12	PCle 3 P1	3 meters
U20	13B	U10	PCle 3 P1	3 meters
U20	14A	U8	PCle 3 P1	3 meters
U20	14B	U6	PCle 3 P1	3 meters
U20	15A	U4	PCle 3 P1	3 meters
U20	15B	U2	PCle 3 P1	3 meters
U24	4B	U18	PCle 3 P2	2 meters
U24	12A	U16	PCle 3 P2	2 meters
U24	12B	U14	PCle 3 P2	3 meters
U24	13A	U12	PCle 3 P2	3 meters
U24	13B	U10	PCle 3 P2	3 meters
U24	14A	U8	PCle 3 P2	3 meters
U24	14B	U6	PCle 3 P2	3 meters
U24	15A	U4	PCle 3 P2	3 meters
U24	15B	U2	PCle 3 P2	3 meters

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

² P_n is the InfiniBand port, where n is the port number.

F

Multirack Cabling Tables

This appendix describes the hardware connections and contains the tables for multirack cabling. It contains the following sections:

- [Understanding Multirack Cabling](#) (page F-1)
- [Key to Cabling Table Abbreviations](#) (page F-4)
- [Two-Rack Cabling](#) (page F-5)
- [Three-Rack Cabling](#) (page F-6)
- [Four-Rack Cabling](#) (page F-8)
- [Five-Rack Cabling](#) (page F-10)
- [Six-Rack Cabling](#) (page F-13)
- [Seven-Rack Cabling](#) (page F-17)
- [Eight-Rack Cabling](#) (page F-21)



See Also:

[Connecting Multiple Oracle Big Data Appliance Racks](#) (page 9-1)

F.1 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 racks.

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 E-1](#) (page E-2).

These switches attach to standard quad small form-factor pluggable (QSFP) connectors at the end of the InfiniBand cables. This document describes how to cable together up to 8 racks without external InfiniBand switches. For systems larger than 8 racks, contact Oracle Support for assistance with formulating a plan.

The cable lengths in this appendix are for adjacent racks with the cables running underneath the racks. In this scenario, the cables are generally 3 feet long to connect within the rack, 5 feet long to connect up to two racks away, and 10 feet long to connect three or more racks away. If the racks are apart or the cables run above the racks, then longer cables may be required. The cable lengths listed here are provided only as a guideline. You should always verify the actual cable lengths required.

F.1.1 Cabling Together Oracle Exadata Database Machine and Oracle Big Data Appliance

Oracle Exadata Database Machine full racks and half racks must have a spine switch. If the rack does not have a spine switch, then be sure to order an Exadata Expansion Switch Kit for the installation.

An Oracle Exadata Database Machine eighth or quarter rack does not require a spine switch to connect to a Oracle Big Data Appliance, whether it is a starter rack, a starter rack with High Capacity nodes plus InfiniBand Infrastructure, or a full rack.




See Also:

Oracle Exadata Database Machine Extending and Multi-Rack Cabling Guide

F.1.2 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 F-1](#) (page F-2).

Figure F-1 Connections Between Spine Switch and Leaf Switches



The servers connect to the leaf switches as shown in [Figure F-2](#) (page F-3).

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



F.1.3 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 F-3](#) (page F-4).


[Figure F-3](#) (page F-4) 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 F-3 Connecting Switches Across Two Racks




F.1.4 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 F-4](#) (page F-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 F-4 Connecting Switches Across Multiple Racks



F.2 Key to Cabling Table Abbreviations

The following abbreviations are used in the tables:

Abbreviation	Description
R <i>n</i>	Rack <i>n</i> , where <i>n</i> is the number of the rack, such as R1.
IB <i>n</i>	Unit location in rack, where <i>n</i> is the number, such as IB3.
P <i>n</i>	InfiniBand port <i>n</i> , where <i>n</i> is port number, such as P8A.

F.3 Two-Rack Cabling

Table F-1 (page F-5) 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 F-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	3 meters
	R1-IB3-P8B to R1-U1-P4A	
	R1-IB3-P9A to R1-U1-P5A	
	R1-IB3-P9B to R1-U1-P6A	
R1 IB3 to Rack 2	R1-IB3-P10A to R2-U1-P7A	5 meters
	R1-IB3-P10B to R2-U1-P8A	
	R1-IB3-P11A to R2-U1-P9A	
	R1-IB3-P11B to R2-U1-P10A	
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
	R1-IB2-P8B to R1-U1-P4B	
	R1-IB2-P9A to R1-U1-P5B	
	R1-IB2-P9B to R1-U1-P6B	
R1 IB2 to Rack 2	R1-IB2-P10A to R2-U1-P7B	5 meters
	R1-IB2-P10B to R2-U1-P8B	
	R1-IB2-P11A to R2-U1-P9B	
	R1-IB2-P11B to R2-U1-P10B	

Table F-2 (page F-6) shows the cable connections for the second spine switch (R2-U1) when two racks are cabled together.

Table F-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	3 meters
	R2-IB3-P8B to R2-U1-P4A	
	R2-IB3-P9A to R2-U1-P5A	
	R2-IB3-P9B to R2-U1-P6A	
R2 IB3 to Rack 1	R2-IB3-P10A to R1-U1-P7A	5 meters
	R2-IB3-P10B to R1-U1-P8A	
	R2-IB3-P11A to R1-U1-P9A	
	R2-IB3-P11B to R1-U1-P10A	
R2 IB2 within Rack 2	R2-IB2-P8A to R2-U1-P3B	3 meters
	R2-IB2-P8B to R2-U1-P4B	
	R2-IB2-P9A to R2-U1-P5B	
	R2-IB2-P9B to R2-U1-P6B	
R2 IB2 to Rack 1	R2-IB2-P10A to R1-U1-P7B	5 meters
	R2-IB2-P10B to R1-U1-P8B	
	R2-IB2-P11A to R1-U1-P9B	
	R2-IB2-P11B to R1-U1-P10B	

F.4 Three-Rack Cabling

Table F-3 (page F-6) 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 F-3 Leaf Switch Connections for the First Rack in a Three-Rack System

Leaf Switch	Connection	Cable Length
R1 IB3 within Rack 1	R1-IB3-P8A to R1-U1-P3A	3 meters
	R1-IB3-P8B to R1-U1-P4A	
	R1-IB3-P9A to R1-U1-P5A	
R1 IB3 to Rack 2	R1-IB3-P9B to R2-U1-P6A	5 meters
	R1-IB3-P10A to R2-U1-P7A	
	R1-IB3-P10B to R2-U1-P8A	
R1 U24 to Rack 3	R1-IB3-P11A to R3-U1-P9A	5 meters
	R1-IB3-P11B to R3-U1-P10A	

Table F-3 (Cont.) Leaf Switch Connections for the First Rack in a Three-Rack System

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

[Table F-4](#) (page F-7) shows the cable connections for the second spine switch (R2-U1) when three racks are cabled together.

Table F-4 Leaf Switch Connections for the Second Rack in a Three-Rack System

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

[Table F-5](#) (page F-7) shows the cable connections for the third spine switch (R3-U1) when three racks are cabled together.

Table F-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	3 meters
	R3-IB3-P8B to R3-U1-P4A	
	R3-IB3-P9A to R3-U1-P5A	

Table F-5 (Cont.) Leaf Switch Connections for the Third Rack in a Three-Rack System

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

F.5 Four-Rack Cabling

[Table F-6](#) (page F-8) 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 F-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	3 meters
	R1-IB3-P8B to R1-U1-P4A	
R1 IB3 to Rack 2	R1-IB3-P9A to R2-U1-P5A	5 meters
	R1-IB3-P9B to R2-U1-P6A	
R1 IB3 to Rack 3	R1-IB3-P10A to R3-U1-P7A	5 meters
	R1-IB3-P10B to R3-U1-P8A	
R1 IB3 to Rack 4	R1-IB3-P11A to R4-U1-P9A	10 meters
	R1-IB3-P11B to R4-U1-P10A	
R1 IB2 within Rack 1	R1-IB2-P8A to R1-U1-P3B	3 meters
	R1-IB2-P8B to R1-U1-P4B	

Table F-6 (Cont.) Leaf Switch Connections for the First Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
R1 IB2 to Rack 2	R1-IB2-P9A to R2-U1-P5B R1-IB2-P9B to R2-U1-P6B	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 F-7](#) (page F-9) shows the cable connections for the second spine switch (R2-U1) when four racks are cabled together.

Table F-7 Leaf Switch Connections for the Second Rack in a Four-Rack System

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

[Table F-8](#) (page F-9) shows the cable connections for the third spine switch (R3-U1) when four racks are cabled together.

Table F-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

Table F-8 (Cont.) Leaf Switch Connections for the Third Rack in a Four-Rack System

Leaf Switch	Connection	Cable Length
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 F-9](#) (page F-10) shows the cable connections for the fourth spine switch (R4-U1) when four racks are cabled together

Table F-9 Leaf Switch Connections for the Fourth Rack in a Four-Rack System

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

F.6 Five-Rack Cabling

[Table F-10](#) (page F-11) 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 F-10 Leaf Switch Connections for the First Rack in a Five-Rack System

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

[Table F-11](#) (page F-11) shows the cable connections for the second spine switch (R2-U1) when five racks are cabled together.

Table F-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

Table F-11 (Cont.) Leaf Switch Connections for the Second Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
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 F-12](#) (page F-12) shows the cable connections for the third spine switch (R3-U1) when five racks are cabled together.

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

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

[Table F-13](#) (page F-12) shows the cable connections for the fourth spine switch (R4-U1) when five racks are cabled together.

Table F-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

Table F-13 (Cont.) Leaf Switch Connections for the Fourth Rack in a Five-Rack System

Leaf Switch	Connection	Cable Length
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 F-14](#) (page F-13) shows the cable connections for the fifth spine switch (R5-U1) when five racks are cabled together.

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

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

F.7 Six-Rack Cabling

[Table F-15](#) (page F-14) 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 F-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 F-16](#) (page F-14) shows the cable connections for the second spine switch (R2-U1) when six racks are cabled together.

Table F-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

Table F-16 (Cont.) Leaf Switch Connections for the Second Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
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 F-17](#) (page F-15) shows the cable connections for the third spine switch (R3-U1) when six racks are cabled together.

Table F-17 Leaf Switch Connections for the Third Rack in a Six-Rack System

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

[Table F-18](#) (page F-15) shows the cable connections for the fourth spine switch (R4-U1) when six racks are cabled together.

Table F-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

Table F-18 (Cont.) Leaf Switch Connections for the Fourth Rack in a Six-Rack System

Leaf Switch	Connection	Cable Length
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 F-19](#) (page F-16) shows the cable connections for the fifth spine switch (R5-U1) when six racks are cabled together.

Table F-19 Leaf Switch Connections for the Fifth Rack in a Six-Rack System

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

[Table F-20](#) (page F-17) shows the cable connections for the sixth spine switch (R6-U1) when six racks are cabled together.

Table F-20 Leaf Switch Connections for the Sixth Rack in a Six-Rack System

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

F.8 Seven-Rack Cabling

[Table F-21](#) (page F-17) 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 F-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

Table F-21 (Cont.) Leaf Switch Connections for the First Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
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	3 meters
	R1-IB2-P8B to R1-U1-P4B	
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 F-22](#) (page F-18) shows the cable connections for the second spine switch (R2-U1) when cabling seven racks.

Table F-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	3 meters
	R2-IB3-P8B to R2-U1-P4A	
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	3 meters
	R2-IB2-P8B to R2-U1-P4B	
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 F-23](#) (page F-19) shows the cable connections for the third spine switch (R3-U1) when seven racks are cabled together.

Table F-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 F-24](#) (page F-19) shows the cable connections for the fourth spine switch (R4-U1) when seven racks are cabled together.

Table F-24 Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 within Rack 4	R4-IB3-P8A to R4-U1-P3A R4-IB3-P8B to R4-U1-P4A	3 meters
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P9A to R5-U1-P5A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9B to R6-U1-P6A	5 meters
R4 IB3 to Rack 7	R4-IB3-P10A to R7-U1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B R4-IB2-P8B to R4-U1-P4B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P9A to R5-U1-P5B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9B to R6-U1-P6B	5 meters

Table F-24 (Cont.) Leaf Switch Connections for the Fourth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R4 IB2 to Rack 7	R4-IB2-P10A to R7-U1-P7B	10 meters

Table F-25 (page F-20) shows the cable connections for the fifth spine switch (R5-U1) when seven racks are cabled together.

Table F-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 F-26 (page F-20) shows the cable connections for the sixth spine switch (R6-U1) when seven racks are cabled together.

Table F-26 Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R6 IB3 within Rack 6	R6-IB3-P8A to R6-U1-P3A R6-IB3-P8B to R6-U1-P4A	3 meters
R6 IB3 to Rack 1	R6-IB3-P9B to R1-U1-P6A	10 meters
R6 IB3 to Rack 2	R6-IB3-P10A to R2-U1-P7A	10 meters
R6 IB3 to Rack 3	R6-IB3-P10B to R3-U1-P8A	10 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-U1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-U1-P10A	5 meters
R6 IB3 to Rack 7	R6-IB3-P9A to R7-U1-P5A	5 meters

Table F-26 (Cont.) Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R6 IB2 within Rack 6	R6-IB2-P8A to R6-U1-P3B R6-IB2-P8B to R6-U1-P4B	3 meters
R6 IB2 to Rack 1	R6-IB2-P9B to R1-U1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB2-P10A to R2-U1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-U1-P8B	10 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters
R6 IB2 to Rack 7	R6-IB2-P9A to R7-U1-P5B	5 meters

[Table F-27](#) (page F-21) shows the cable connections for the seventh spine switch (R7-U1) when seven racks are cabled together.

Table F-27 Leaf Switch Connections for the Seventh Rack in a Seven-Rack System

Leaf Switch	Connection	Cable Length
R7 IB3 within Rack 7	R7-IB3-P8A to R7-U1-P3A R7-IB3-P8B to R7-U1-P4A	3 meters
R7 IB3 to Rack 1	R7-IB3-P9A to R1-U1-P5A	10 meters
R7 IB3 to Rack 2	R7-IB3-P9B to R2-U1-P6A	10 meters
R7 IB3 to Rack 3	R7-IB3-P10A to R3-U1-P7A	10 meters
R7 IB3 to Rack 4	R7-IB3-P10B to R4-U1-P8A	10 meters
R7 IB3 to Rack 5	R7-IB3-P11A to R5-U1-P9A	5 meters
R7 IB3 to Rack 6	R7-IB3-P11B to R6-U1-P10A	5 meters
R7 IB2 within Rack 7	R7-IB2-P8A to R7-U1-P3B R7-IB2-P8B to R7-U1-P4B	3 meters
R7 IB2 to Rack 1	R7-IB2-P9A to R1-U1-P5B	10 meters
R7 IB2 to Rack 2	R7-IB2-P9B to R2-U1-P6B	10 meters
R7 IB2 to Rack 3	R7-IB2-P10A to R3-U1-P7B	10 meters
R7 IB2 to Rack 4	R7-IB2-P10B to R4-U1-P8B	10 meters
R7 IB2 to Rack 5	R7-IB2-P11A to R5-U1-P9B	5 meters
R7 IB2 to Rack 6	R7-IB2-P11B to R6-U1-P10B	5 meters

F.9 Eight-Rack Cabling

[Table F-28](#) (page F-22) 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 F-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-P9B	10 meters
R1 IB2 to Rack 8	R1-IB2-P11B to R8-U1-P10B	10 meters

[Table F-29](#) (page F-22) shows the cable connections for the second spine switch (R2-U1) when eight racks are cabled together.

Table F-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

Table F-29 (Cont.) Leaf Switch Connections for the Second Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
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 F-30](#) (page F-23) shows the cable connections for the third spine switch (R3-U1) when eight racks are cabled together.

Table F-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	10 meters
R3 IB3 to Rack 7	R3-IB3-P10A to R7-U1-P7A	10 meters
R3 IB3 to Rack 8	R3-IB3-P10B to R8-U1-P8A	10 meters
R3 IB2 within Rack 3	R3-IB2-P8A to R3-U1-P3B	3 meters
R3 IB2 to Rack 1	R3-IB2-P11A to R1-U1-P9B	5 meters
R3 IB2 to Rack 2	R3-IB2-P11B to R2-U1-P10B	5 meters
R3 IB2 to Rack 4	R3-IB2-P8B to R4-U1-P4B	5 meters
R3 IB2 to Rack 5	R3-IB2-P9A to R5-U1-P5B	5 meters
R3 IB2 to Rack 6	R3-IB2-P9B to R6-U1-P6B	10 meters
R3 IB2 to Rack 7	R3-IB2-P10A to R7-U1-P7B	10 meters
R3 IB2 to Rack 8	R3-IB2-P10B to R8-U1-P8B	10 meters

[Table F-31](#) (page F-23) shows the cable connections for the fourth spine switch (R4-U1) when eight racks are cabled together.

Table F-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

Table F-31 (Cont.) Leaf Switch Connections for the Fourth Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
R4 IB3 to Rack 1	R4-IB3-P10B to R1-U1-P8A	10 meters
R4 IB3 to Rack 2	R4-IB3-P11A to R2-U1-P9A	5 meters
R4 IB3 to Rack 3	R4-IB3-P11B to R3-U1-P10A	5 meters
R4 IB3 to Rack 5	R4-IB3-P8B to R5-U1-P4A	5 meters
R4 IB3 to Rack 6	R4-IB3-P9A to R6-U1-P5A	5 meters
R4 IB3 to Rack 7	R4-IB3-P9B to R7-U1-P6A	10 meters
R4 IB3 to Rack 8	R4-IB3-P10A to R8-U1-P7A	10 meters
R4 IB2 within Rack 4	R4-IB2-P8A to R4-U1-P3B	3 meters
R4 IB2 to Rack 1	R4-IB2-P10B to R1-U1-P8B	10 meters
R4 IB2 to Rack 2	R4-IB2-P11A to R2-U1-P9B	5 meters
R4 IB2 to Rack 3	R4-IB2-P11B to R3-U1-P10B	5 meters
R4 IB2 to Rack 5	R4-IB2-P8B to R5-U1-P4B	5 meters
R4 IB2 to Rack 6	R4-IB2-P9A to R6-U1-P5B	5 meters
R4 IB2 to Rack 7	R4-IB2-P9B to R7-U1-P6B	10 meters
R4 IB2 to Rack 8	R4-IB2-P10A to R8-U1-P7B	10 meters

[Table F-32](#) (page F-24) shows the cable connections for the fifth spine switch (R5-U1) when eight racks are cabled together.

Table F-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 F-33 (page F-25) shows the cable connections for the sixth spine switch (R6-U1) when eight racks are cabled together.

Table F-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	10 meters
R6 IB3 to Rack 4	R6-IB3-P11A to R4-U1-P9A	5 meters
R6 IB3 to Rack 5	R6-IB3-P11B to R5-U1-P10A	5 meters
R6 IB3 to Rack 7	R6-IB3-P8B to R7-U1-P4A	5 meters
R6 IB3 to Rack 8	R6-IB3-P9A to R8-U1-P5A	5 meters
R6 IB2 within Rack 6	R6-IB2-P8A to R6-U1-P3B	3 meters
R6 IB2 to Rack 1	R6-IB2-P9B to R1-U1-P6B	10 meters
R6 IB2 to Rack 2	R6-IB2-P10A to R2-U1-P7B	10 meters
R6 IB2 to Rack 3	R6-IB2-P10B to R3-U1-P8B	10 meters
R6 IB2 to Rack 4	R6-IB2-P11A to R4-U1-P9B	5 meters
R6 IB2 to Rack 5	R6-IB2-P11B to R5-U1-P10B	5 meters
R6 IB2 to Rack 7	R6-IB2-P8B to R7-U1-P4B	5 meters
R6 IB2 to Rack 8	R6-IB2-P9A to R8-U1-P5B	5 meters

Table F-34 (page F-25) shows the cable connections for the seventh spine switch (R7-U1) when eight racks are cabled together.

Table F-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

Table F-34 (Cont.) Leaf Switch Connections for the Seventh Rack in an Eight-Rack System

Leaf Switch	Connection	Cable Length
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-IB2-P8B to R8-U1-P4B	5 meters

[Table F-35](#) (page F-26) shows the cable connections for the eighth spine switch (R8-U1) when eight racks are cabled together.

Table F-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
R8 IB3 to Rack 4	R8-IB3-P10A to R4-U1-P7A	10 meters
R8 IB3 to Rack 5	R8-IB3-P10B to R5-U1-P8A	10 meters
R8 IB3 to Rack 6	R8-IB3-P11A to R6-U1-P9A	5 meters
R8 IB3 to Rack 7	R8-IB3-P11B to R7-U1-P10A	5 meters
R8 IB2 within Rack 8	R8-IB2-P8A to R8-U1-P3B	3 meters
R8 IB2 to Rack 1	R8-IB2-P8B to R1-U1-P4B	10 meters
R8 IB2 to Rack 2	R8-IB2-P9A to R2-U1-P5B	10 meters
R8 IB2 to Rack 3	R8-IB2-P9B to R3-U1-P6B	10 meters
R8 IB2 to Rack 4	R8-IB2-P10A to R4-U1-P7B	10 meters
R8 IB2 to Rack 5	R8-IB2-P10B to R5-U1-P8B	10 meters
R8 IB2 to Rack 6	R8-IB2-P11A to R6-U1-P9B	5 meters
R8 IB2 to Rack 7	R8-IB2-P1B to R7-U1-P10B	5 meters

Index

Symbols

, Client Network Config Summary page, [4-25](#)

A

AC breaker panels, [2-10](#)
AC power, location, [2-10](#)
access route requirements, [2-5](#)
accessing servers from KVM, [7-3](#)
accessory kits, [1-4](#)
acoustic levels, [2-1](#), [2-3](#)
adding racks, [10-8](#)
adding services, [13-7](#)
adding to an existing cluster configuration
 options, [4-25](#)
administration software overview, [1-6](#)
administrative network, [3-2](#)
 cabling requirements, [2-14](#)
 connections, [2-14](#)
 ports, [3-1](#)
administrative network diagram, [15-2](#)
Administrative Network page (configuration), [4-8](#)
administrative utilities, [1-6](#)
African PDU specifications, [2-9](#)
air conditioning, [2-12](#), [2-16](#), [6-2](#)
airflow, [2-2–2-4](#), [2-11](#)
airflow requirements, [2-13](#)
alert notifications, [4-24](#)
alerts, [15-2](#)
altitude requirements, [2-11](#)
ambient relative humidity range, [2-12](#)
ambient temperature range, [2-12](#), [11-1](#)
American PDU specifications, [2-7](#)
Asia Pacific PDU specifications, [2-9](#)
ASR
 about, [5-1](#)
 activating assets, [5-4](#)
 adding support, [10-16](#)
 configuration, [5-2](#)
 installation prerequisites, [10-3](#)
 troubleshooting, [5-6](#)
 validation, [5-5](#), [10-7](#)
ASR Manager, [4-23](#)
 installation, [5-3](#)

ASR Manager (*continued*)
 port number, [4-23](#)
ASR network connections, [5-2](#)
ASR trap destinations, [5-4](#)
asrexacheck script, [10-7](#)
asset verification (ASR), [5-4](#)
Audit Vault, [10-3](#)
 installation prerequisites, [10-3](#)
 plug-in configuration, [4-22](#)
Auto Service Request, [5-1](#)
 configuration, [4-23](#)
 installing, [10-29](#)
 port assignments, [3-10](#)

B

base image reinstallation, [10-19](#)
base image utilities, [10-32](#)
batteries
 life expectancy, [B-2](#)
 replacement policy, [B-2](#)
BDA server offline repair category, [B-2](#)
BDA_REBOOT files, [6-14](#), [13-22](#)
bda-checkip.out file, [3-12](#)
bda-install-review.html file, [4-3](#)
BDA-O repair category, [B-2](#)
bda-preinstall-checkip.sh file, [4-3](#)
bda.zip configuration archive, [4-3](#)
bdacheckcluster utility, [13-2](#)
bdacheckhw utility, [7-21](#), [11-9](#), [13-3](#)
bdacheckib utility, [13-3](#)
bdachecknet utility, [8-7](#), [13-5](#)
bdachecksw utility, [7-22](#), [11-10](#), [13-5](#)
bdacli utility, [13-7](#)
bdaclustersynctime utility, [13-6](#)
BdaDeploy.json file, [4-3](#)
bdadiag utility, [13-17](#)
bdagetclientconfig utility, [13-20](#)
bdaibstatus utility, [13-21](#)
bdaid utility, [13-21](#)
bdaimagevalidate utility, [13-22](#)
BDAMammoth directory, [10-27](#)
bdaredoclientnet utility, [13-23](#)
bdaserials utility, [13-25](#)
bdaswapoff utility, [13-26](#)

bdaswapon utility, [13-26](#)
 bdaupdatefw utility, [13-27](#)
 Big Data Appliance
 See Oracle Big Data Appliance
 Big Data Connectors
 See Oracle Big Data Connectors
 Big Data SQL
 starting, stopping, restarting, [13-10](#)
 black cables, [1-5](#), [C-1](#), [D-1](#), [E-1](#)
 blue cables, [1-5](#), [C-1](#), [D-1](#), [E-1](#)
 bondeth0 interface
 network connections, [3-1](#)
 bondib0 interface
 description, [3-3](#)
 maintenance, [11-7](#)
 bondib0 network
 verifying operations, [11-9](#)
 boot order, resetting, [13-31](#)
 breaker panels, [2-11](#)

C

cable tie,
 removing, [6-10](#)
 cables, [1-4](#)
 color coding, [C-1](#), [D-1](#), [E-1](#)
 spare parts, [1-4](#)
 spares, [1-4](#)
 splitter, [3-7](#)
 cabling
 client network, [3-5](#)
 eight racks, [F-21](#)
 existing network requirements, [2-14](#)
 five racks, [F-10](#)
 four racks, [F-8](#)
 multirack, [9-2](#), [F-1](#)
 seven racks, [F-17](#)
 six racks, [F-13](#)
 spare parts, [B-25](#)
 three racks, [F-6](#)
 two racks, [F-5](#)
 cabling tables
 with KVM (Sun Fire servers), [E-1](#)
 without KVM (X4-2L and X3-2L servers), [D-1](#)
 without KVM (X6-2 and X5-2 servers), [C-1](#)
 cage nuts, [6-3](#)
 casters, [6-6](#)
 cautions, [11-4](#)
 CDH
 health checks, [13-2](#)
 CDH cluster configuration, [4-16](#)
 CDH cluster logs, [13-17](#)
 CDH installation, [10-29](#)
 CDH ports, [3-10](#)
 CDH services, starting, [10-29](#)
 ceiling height requirements, [2-4](#)
 checking firmware, [7-22](#)
 checking network configuration, [13-5](#)
 checking the server software, [13-5](#)
 circuit breakers, [2-11](#), [6-12](#)
 Cisco switch, [1-4](#), [7-3](#)
 replacement parts, [B-23](#)
 See also Ethernet switch
 CleanupInstall step, [10-29](#)
 CleanupSSHroot step, [10-29](#)
 CLI utility, [13-7](#)
 CLI, Oracle ILOM, [15-2](#)
 client access network
 cabling requirements, [2-14](#)
 description, [3-2](#)
 requirements, [2-14](#)
 client configuration distribution, [13-20](#)
 client network, [3-6](#)
 bandwidth, [3-6](#)
 cabling, [3-5](#)
 configuration, [4-7](#)
 Client network
 configuration, [4-13](#), [4-17](#)
 LACP/LAG, [4-13](#), [4-17](#)
 client network connections, [3-9](#)
 clock configuration
 Ethernet switch, [7-9](#)
 InfiniBand switches, [7-14](#)
 Cloud Control, [1-6](#)
 See Also Oracle Enterprise Manager Cloud
 Control, [4-24](#)
 Cloudera's Distribution including Apache Hadoop
 See CDH, [1-1](#)
 cluster expansion, [10-8](#)
 cluster name configuration, [4-11](#)
 cluster naming conventions, [4-6](#)
 Cluster page (configuration), [4-12](#)
 cluster reimaging, [10-33](#)
 cluster_name-config.json, [4-3](#)
 cmf.keytab file, [10-4](#)
 collecting diagnostics, [13-17](#)
 color coding, cables, [C-1](#), [D-1](#), [E-1](#)
 command files, executing across servers, [14-5](#)
 command-line interface, [13-7](#)
 commands, remote execution across servers,
 [1-6](#), [14-1](#)
 Complete page, [4-25](#)
 component configuration, full or starter rack, [7-1](#)
 component failure, environmental causes, [2-11](#)
 component names, [4-6](#)
 components
 location in the rack, [C-1](#), [D-1](#), [E-1](#)
 server, [1-2](#)
 spare parts kit, [1-4](#)
 configuration

configuration (*continued*)

- disk controller, [11-6](#)
- editing in the Configuration Generation utility, [4-3](#)
- full or starter rack components, [7-1](#)
- physical drives, [11-7](#)
- querying, [13-7](#)
- virtual drive, [11-6](#)

Configuration, [4-25](#)

configuration files, [4-3](#)

- copying to new servers, [8-5](#)
- generating, [4-3](#), [4-25](#)
- ntp.conf file, [11-13](#)

Configuration Generation Utility, [4-1](#)

- input file, [4-3](#)
- installing, [4-3](#)
- overview, [4-1](#)

configuration worksheets

- See Configuration Generation Utility

configure terminal command, [7-4](#)

configuring a new disk, [12-7](#)

configuring ASR, [5-2](#)

configuring Ethernet switch, [7-3](#), [7-5](#)

configuring KVM, [7-1](#), [7-3](#)

configuring servers, [7-20](#)

connecting KVM to the network, [7-2](#)

connectors

- See Oracle Big Data Connectors

cooling requirements, [2-1–2-4](#)

copper cables, [1-5](#)

copying files, [14-2](#)

CopyLicenseFiles step, [10-29](#)

corrosion problems, [2-12](#)

CreateUsers step, [10-29](#)

Customer Details Page (configuration), [4-5](#)

D

data center switch configuration, [2-15](#)

dba group, [4-13](#), [4-17](#), [10-30](#)

dcli utility

- description, [1-6](#)
- examples, [14-5](#)
- option descriptions, [14-3](#)
- return values, [14-5](#)
- syntax, [14-3](#)
- target servers, [14-2](#)
- using, [14-1](#)

Define Clusters page (configuration), [4-11](#)

diagnostics collection, [13-17](#)

dimensions

- rack, [2-4](#)
- shipping, [2-4](#)

directories, copying files across servers, [14-4](#)

disablesm command, [7-15](#), [9-2](#), [9-4](#), [9-5](#), [11-8](#)

disk controllers, [11-6](#)

disk drives

- physical slots, [12-2](#)
- server status indicators, [B-3](#)
- spares, [1-4](#), [B-6](#)
- symbolic links, [12-2](#)

disks

- configuring, [12-7](#)
- repairing, [11-6](#)
- replacing, [12-1](#)

DNS network registration, [2-15](#)

DNS server requirement, [3-2](#), [4-9](#)

documentation, third-party hardware, [1-4](#)

domain name system configuration, [4-9](#)

door dimensions, [2-5](#)

downgrading the factory image, [10-8](#)

drive configuration, physical, [11-7](#)

E

earth ground cable, [6-8](#)

electric receptacles, [2-6](#)

electrical access, [2-5](#)

electrical requirements, [2-6](#), [6-2](#), [6-13](#)

electrical storms, [2-6](#)

electronic noise, [6-2](#)

electrostatic discharge, [2-12](#)

elevator requirements, [2-5](#)

email accounts

- alerts (Cloudera Manager), [4-24](#)

email accounts, My Oracle Support (ASR), [5-1](#)

email alerts configuration, [4-24](#)

emergency shutdown, [11-3](#)

enablesm command, [7-15](#), [9-4](#), [9-5](#)

engineered systems, connecting to Oracle Big Data Appliance, [C-8](#), [D-8](#), [E-7](#)

Enterprise Management Cloud Control

- configuration, [4-24](#)

Enterprise Manager

- See Also Oracle Enterprise Manager Cloud Control, [4-24](#)

env_test command, [7-15](#)

environmental requirements

- full rack, [2-3](#)

environmental specifications, [2-1](#), [2-3](#)

eth0 interface, [3-2](#)

Ethernet bridge ports, [13-30](#)

Ethernet cabling, [2-14](#), [3-5](#)

Ethernet interface, [3-2](#)

Ethernet over InfiniBand protocol., [3-6](#)

Ethernet ports, [2-2–2-4](#), [3-1](#)

Ethernet switch, [7-3](#)

- configuration, [7-3](#), [7-4](#)
- configuring host name, [7-5](#)
- IP routing, [7-5](#)

Ethernet switch (*continued*)
 port speed, [7-5](#)
 replacement parts, [B-23](#)
 setting remote password, [7-8](#)
 spare parts, [1-4](#)
 time settings, [7-9](#)
 validating and saving configuration, [7-11](#)
 European PDU specifications, [2-9](#)
 ext4 file system, [12-16](#)
 external filer configuration, [10-29](#)

F

Fabric Management, [7-15](#)
 factory image reinstallation, [10-19](#)
 factory patches, [10-29](#)
 failover connections, [3-1](#), [3-2](#), [3-5](#)
 fan modules
 server status indicators, [B-3](#)
 switch status indicators, [B-5](#)
 files, copying to multiple servers, [14-4](#)
 filler panels, [2-13](#)
 firewall configuration, [2-15](#)
 Firmware State status, [12-8](#)
 firmware updates, [13-27](#)
 first server location, [7-23](#)
 floor load, [2-5](#), [2-6](#)
 floor tiles, [2-13](#)
 flooring
 raised, [2-6](#)
 raised pattern, [2-5](#)
 Foreign State status, [12-8](#)
 forklift, [2-5](#)
 FRUs, [B-1](#)
 full rack component configuration, [7-1](#)

G

gateway configuration
 Ethernet switch, [7-5](#)
 gateway switch factory IP addresses, [3-9](#)
 gateway switch ports, [E-7](#)
 gateway switches
 assigning to VNICs, [3-6](#)
 cabling, [3-5](#)
 configuring, [7-12](#)
 connecting to Oracle ILOM, [15-6](#)
 guidelines for connecting, [3-7](#)
 replacing, [11-7](#)
 See also InfiniBand switches, [3-6](#)
 using splitter cables, [3-7](#)
 General Network Properties page (configuration),
 [4-9](#)
 getmaster command, [9-2](#), [9-4](#), [9-5](#)
 getting help, [13-21](#)

ground cable, [6-8](#)
 grounding resistance, [6-1](#)
 grounding to earth, [2-11](#)
 groups, hadoop, [10-30](#)
 GRUB master boot records, [12-17](#)

H

Hadoop cluster configuration, [4-16](#)
 hadoop group, [10-30](#)
 Hadoop installation, [10-29](#)
 Hadoop ports, [3-10](#)
 Hadoop services, starting, [10-29](#)
 hardware configuration, [4-7](#)
 querying, [13-7](#)
 hardware diagnostics, [15-1](#)
 hardware failures, automatic reporting, [5-1](#)
 hardware profiles, acquiring from servers, [7-21](#)
 Hardware Selection page (configuration), [4-7](#)
 hardware status indicators, [B-2](#)
 hardware validation, [13-2](#), [13-3](#), [13-22](#)
 HDFS diagnostics, [13-17](#)
 HDFS disk formatting, [12-19](#)
 HDFS partitions, [12-3](#)
 HDFS Transparent Encryption
 configuration, [4-22](#)
 hdfs user, [10-30](#)
 health checks on InfiniBand switches, [7-15](#)
 High Capacity node plus InfiniBand Infrastructure
 cables, [8-2](#)
 components, [8-2](#)
 configuring servers, [8-4](#)
 installing servers, [8-3](#)
 network configuration, [8-6](#)
 PDU firmware requirement, [8-2](#)
 high-voltage PDUs, [2-9](#), [11-14](#), [11-15](#)
 host factory IP addresses, [3-9](#)
 hot swap repair category, [B-1](#)
 HS repair category, [B-1](#)
 humidity range, [2-11](#), [2-12](#)

I

ibdiagnet utility, [11-9](#)
 iblinkinfo utility, [13-28](#)
 ibqueryerrors utility, [11-9](#)
 ibstatus utility, [7-20](#)
 ILOM
 See Oracle ILOM
 ilom-admin default password, [7-13](#)
 image validation, [13-22](#)
 image version
 downgrading, [10-8](#)
 upgrading, [10-8](#)
 imagehistory utility, [13-29](#)

imageinfo utility, [10-9](#), [13-29](#)
 incline, maximum, [2-5](#)
 InfiniBand configuration, [7-12](#)
 InfiniBand Gateway switch ports, [E-7](#)
 InfiniBand network
 cabling checks, [13-3](#)
 factory IP addresses, [3-9](#)
 network checks, [13-21](#)
 InfiniBand network interface card, [3-1](#)
 InfiniBand Network page (configuration), [4-14](#),
 [4-18](#)
 InfiniBand network verification, [11-9](#)
 InfiniBand ports, connecting to other devices,
 [C-8](#), [D-8](#), [E-7](#)
 InfiniBand private network, [3-3](#)
 InfiniBand Subnet Manager, [11-11](#)
 InfiniBand switches
 checking status, [7-20](#)
 configuring, [7-12](#)
 health checks, [7-15](#)
 priority settings, [7-16](#)
 replacing, [11-7](#)
 setting time zone, [7-14](#)
 spare parts, [B-22](#)
 infrastructure repair category, [B-1](#)
 installation of new servers over a starter rack, [8-3](#)
 installation prerequisites, [10-3](#)
 InstallBDASoftware step, [10-29](#)
 InstallHadoop step, [10-29](#)
 installing base image, [10-19](#)
 installing patches, [10-29](#)
 installing the software, [10-6](#)
 International Electrotechnical Commission (IEC),
 [2-11](#)
 IP address requirements, [2-2–2-4](#)
 IP addresses
 administrative network configuration, [4-8](#)
 editing individual, [4-10](#)
 factory settings, [3-9](#)
 InfiniBand network configuration, [4-14](#), [4-18](#)
 IP routing, [7-5](#)
 ipmitool command, [7-21](#)
 IR repair category, [B-1](#)

J

Japanese PDU specifications, [2-7](#)
 jumper brackets, [6-3](#)

K

Kerberos authentication
 adding support, [10-18](#)
 configuration, [4-20](#), [4-21](#)
 NoSQL, [4-15](#)

Kerberos authentication (*continued*)
 installation prerequisites (remote KDC only),
 [10-3](#)
 Kerberos configuration, [10-29](#)
 kernel device names (disk drives), [12-2](#)
 key distribution center, [4-20](#), [10-3](#)
 KVM switch
 accessing servers, [7-3](#)
 cabling, [E-2](#)
 configuration procedures, [7-1](#)
 connecting to the network, [7-2](#)
 MAC address, [7-3](#)
 naming, [7-3](#)
 starting, [7-2](#)

L

leaf switches
 configuring, [7-12](#)
 See also InfiniBand switches, [7-12](#)
 LED color codes, [B-3](#), [B-4](#)
 LED status indicators, [B-2](#)
 leveling bolts, [6-5](#)
 leveling feet, [6-6](#), [6-8](#)
 license location, [10-29](#)
 lift gate, [2-5](#)
 lightning arresters, [2-6](#)
 lights-out management, [15-1](#)
 Link Aggregation Group, [4-13](#), [4-17](#)
 listlinkup command, [13-30](#)
 loading dock, [2-5](#)
 locating first server, [7-23](#)
 locks, panel, [6-3](#)
 low-voltage PDUs, [2-7](#), [11-14](#)

M

MAC address of KVM switch, [7-3](#)
 MAC addresses, [13-25](#)
 maintenance access requirements, [2-5](#)
 makebdaimage utility, [10-32](#)
 Mammoth download, [10-4](#)
 Mammoth options, [10-28](#)
 Mammoth Utility
 description, [10-6](#)
 execution steps, [10-29](#)
 location, [10-27](#)
 reasons for failure, [10-9](#)
 running, [10-6](#)
 SetupHadoopSecurity step
 enabling Kerberos, [10-29](#)
 enabling Network Encryption, [10-29](#)
 enabling Sentry, [10-29](#)
 enabling Transparent Encryption, [10-29](#)
 setting mount points, [10-29](#)

Mammoth Utility (*continued*)
 synchronizing clocks, [7-15](#)
 syntax, [10-27](#)
 mammoth-rack_name.params file, [4-3](#)
 management network
 cabling requirements, [2-14](#)
 connecting the KVM, [7-2](#)
 connections, [2-14](#)
 factory IP addresses, [3-9](#)
 ports, [3-1](#)
 management utilities, [1-6](#), [15-2](#)
 mapped user, [10-30](#)
 master.xml configuration file, [4-3](#)
 MegaCli64 command, [7-22](#), [11-6](#)
 memory, adding to servers, [11-4](#)
 Middle Eastern PDU specifications, [2-9](#)
 minimum ceiling height, [2-4](#)
 monitoring PDUs, [11-13](#)
 mount points, [12-3](#)
 mount_hadoop_dirs utility, [13-30](#)
 mounting a device, [7-24](#)
 mounting tool, [6-3](#)
 multirack cabling
 procedures, [9-1](#)
 three to eight, [9-3](#)
 two racks, [9-2](#)
 multirack clusters, [3-9](#)
 multirack software configuration, [10-8](#)
 My Oracle Support email account, [5-1](#)
 MySQL Database installation, [10-29](#)

N

naming conventions, [4-6](#)
 NET0, [3-2](#)
 network
 connecting the KVM, [7-2](#)
 factory IP addresses, [3-9](#)
 network configuration, [4-8](#)
 network configuration file, [4-3](#)
 network connections, [2-14](#), [7-26](#)
 ASR, [5-2](#)
 reconfiguring 10GbE, [13-23](#)
 network diagrams, [3-3](#), [15-2](#)
 network errors, checking for, [3-11](#)
 network interfaces, InfiniBand card, [3-1](#)
 network operations, verifying InfiniBand, [11-9](#)
 network registration, [2-15](#)
 network requirements, [2-2](#)–[2-4](#)
 network settings
 preinstall validation, [4-5](#)
 network suffix configuration, [4-5](#)
 network time protocol, [4-9](#), [11-12](#)
 network topology, [11-11](#)
 network validation, [13-2](#), [13-5](#)

network.json configuration file, [4-3](#)
 networkexpansion.json configuration file, [8-6](#)
 Networking page (configuration), [4-8](#)
 networks
 administrative, [3-2](#)
 client access, [2-14](#), [3-2](#)
 descriptions, [3-2](#)
 InfiniBand, [3-3](#)
 networksetup-one script, [8-6](#)
 networksetup-two script, [8-7](#)
 Node Migration, [11-15](#)
 noise, electronic, [6-2](#)
 North American PDU specifications, [2-7](#)
 NTP server requirement, [3-2](#), [4-10](#)
 NTP servers, [11-12](#)
 configuring, [7-10](#)
 setting on InfiniBand switches, [7-15](#)
 ntp.conf file, [11-13](#)
 nuts, [6-3](#)

O

ODI

 See Oracle Data Integrator
 oinstall group, [4-13](#), [4-17](#), [10-30](#)
 operating cautions and warnings, [11-4](#)
 operating system disks, [12-11](#)
 operating system history, [13-29](#)
 operating system image, [13-29](#)
 Oracle Audit Vault and Database Firewall, [10-3](#)
 plug-in configuration, [4-22](#)
 See also Audit Vault
 Oracle Big Data Appliance,
 general description, [1-1](#)
 rack components, [1-1](#)
 server components, [1-2](#)
 Oracle Big Data Connectors,
 configuration, [4-20](#)
 installation, [10-9](#), [10-29](#)
 installing server components, [4-20](#)
 software installation, [10-9](#)
 Oracle Data Integrator, [4-20](#)
 Oracle Enterprise Manager
 troubleshooting, [10-4](#)
 Oracle Enterprise Manager Cloud Control, [1-6](#),
 [4-24](#)
 adding support on Oracle Big Data
 Appliance, [10-16](#)
 rediscovery, [10-9](#)
 Oracle Enterprise Manager plug-in installation
 prerequisites, [10-3](#)
 Oracle Exadata Database Machine, [4-13](#), [4-17](#)
 Oracle ILOM, [15-1](#)
 backing up configuration, [15-7](#)
 connecting to CLI, [15-4](#)

Oracle ILOM (*continued*)
 connecting to web interface, [15-4](#)
 description, [15-1](#)
 Ethernet connections, [15-3](#)
 factory IP addresses, [3-9](#)
 opening serial connections, [15-5](#)
 restoring configuration, [15-8](#)
 user interfaces, [15-2](#)
 Oracle ILOM port, [3-1](#)
 Oracle Inventory Group, [4-13](#), [4-17](#)
 Oracle Linux, [1-1](#)
 Oracle NoSQL Database
 configuration, [4-12](#), [4-16](#)
 disk formatting, [12-19](#)
 installation, [10-29](#)
 Oracle Server X6-2L
 replacement parts, [B-6](#)
 Oracle Support, [13-21](#)
 oracle user name, [4-13](#), [4-17](#)
 orange cables, [1-5](#), [C-1](#), [D-1](#), [E-1](#)
 OS watcher logs, [13-17](#)

P

pallet jack, [2-5](#)
 panel locks, [6-3](#)
 parallel execution of commands, [14-1](#)
 partitions
 creating on a new disk, [12-11](#), [12-13](#)
 dismounting, [12-4](#)
 parts kit, [1-4](#)
 passwordless SSH, [13-30](#), [13-32](#), [14-1](#)
 passwords (default)
 Ethernet switch, [7-5](#)
 ilom-admin, [7-13](#)
 PDU metering unit, [7-19](#)
 servers, [7-20](#)
 patch installation, [10-29](#)
 PatchFactoryImage step, [10-29](#)
 path to data center, [2-5](#)
 PDAs, location in the rack, [C-1](#), [D-1](#), [E-1](#)
 PDU configuration, [7-18](#)
 PDU metering unit password, [7-19](#)
 PDU requirements, [2-6](#)
 PDU specifications
 for America, Japan, Taiwan, [2-7](#)
 for Americas, Japan, Taiwan, [2-7](#)
 for Europe, Middle East, Africa, Asia Pacific, [2-9](#)
 PDU thresholds, [11-13](#)
 PDUs
 circuit breakers, [6-12](#)
 network connections, [7-18](#)
 powering on, [6-12](#)
 replacement parts, [B-24](#)

PDUs (*continued*)
 single-phase, [6-2](#), [6-13](#)
 perforated tiles, [2-13](#)
 photocopiers, [6-2](#)
 physical drive configuration, [11-7](#)
 port assignments, [3-10](#)
 port speed of Ethernet switch, [7-5](#)
 ports
 ASR Manager, [4-23](#)
 connecting to 40 GbE, [3-7](#)
 ports, open gateway, [C-8](#), [D-8](#), [E-7](#)
 power connectors, [2-6](#)
 power cord length requirements, [2-10](#)
 power cords, [2-11](#)
 connecting to rack, [6-10](#)
 removing cable ties, [6-10](#)
 routing, [6-10](#)
 power LED locations, [6-13](#)
 power off switch, [11-3](#)
 power outlet grounding, [2-11](#)
 power receptacle location, [2-10](#)
 power requirements, [2-1–2-4](#), [2-6](#)
 power supplies
 server status indicators, [B-4](#)
 switch status indicators, [B-5](#)
 uninterruptable, [2-11](#)
 powering off, [11-2](#)
 emergency, [11-3](#)
 powering on, [6-12](#), [11-2](#)
 preinstall-checkip script
 running, [3-12](#)
 preinstall-checkip.sh file, [4-3](#)
 preinstallation tasks, [2-14](#)
 prerequisites to software installation, [10-3](#)
 private network
 verifying operations, [11-9](#)
 problem reporting, [5-1](#)
 Puppet utility, [10-29](#)

Q

QSFP cables, [1-5](#)
 QSFP ports, splitting, [3-7](#)
 querying configuration, [13-7](#)

R

rack base name configuration, [4-5](#)
 rack components, [1-1](#)
 Rack Details page (configuration), [4-7](#)
 rack dimensions, [2-4](#)
 rack layout illustration, [C-1](#), [D-1](#), [E-1](#)
 rack layouts
 with KVM (Sun Fire servers), [E-1](#)
 without KVM, [C-1](#), [D-1](#)

rack reimaging, [10-21](#)
 rack weight, [2-1–2-3](#)
 racks
 adding another, [9-1](#)
 connecting power cords, [6-10](#)
 inspecting, [6-9](#)
 installing Oracle Big Data Appliance
 software, [10-6](#)
 moving instructions, [6-6](#)
 reimaging, [10-34](#)
 reinstalling base image, [10-19](#)
 RAID arrays, repairing logical, [12-14](#)
 RAID configuration, [7-22](#)
 RAID logical drives, [12-2](#)
 red cables, [1-5](#), [C-1](#), [D-1](#), [E-1](#)
 reimagecluster utility, [10-33](#)
 reimagerack utility, [10-34](#)
 reimaging a rack, [10-21](#)
 reimaging the servers, [10-19](#)
 reinstalling the software, [10-7](#)
 relative humidity, [2-12](#), [2-13](#)
 remove-root-ssh script, [13-30](#), [14-2](#)
 removing services, [13-7](#)
 repair categories, [B-1](#)
 replacement parts, [B-1](#)
 replacing spine switch, [11-7](#)
 reset-boot-order utility, [13-31](#)
 resolv.conf configuration file, [8-7](#)
 restart validation, [6-14](#)
 restarting Big Data SQL, [13-10](#)
 Review and Edit Details page (configuration),
 [4-10](#)
 Rolling upgrades, [10-2](#)
 root user, passwordless SSH, [14-1](#)
 route requirements, [2-5](#)
 running-config output, [7-11](#)

S

safety guidelines, [6-1](#)
 scp file copy utility, [14-2](#)
 screwdrivers, [6-4](#)
 screws, [6-3](#)
 search domain configuration, [4-9](#)
 securing the rack, [6-8](#)
 sensor monitoring, [15-2](#)
 serial numbers, [13-25](#)
 server boot order, resetting, [13-31](#)
 server components, [1-2](#)
 server configuration, [7-20](#)
 High Capacity node plus InfiniBand
 Infrastructure, [8-4](#)
 server diagnostics, [13-17](#)
 server disks
 configuring, [12-7](#)
 server disks (*continued*)
 replacing, [12-1](#)
 server factory IP addresses, [3-9](#)
 server LEDs, [B-2](#)
 server locations in the rack, [C-1](#), [D-1](#), [E-1](#)
 server software validation, [13-5](#)
 servers
 adding memory, [11-4](#)
 adding to a cluster, [10-8](#)
 allocating into clusters, [4-11](#)
 connecting to Oracle ILOM, [15-5](#)
 default passwords, [7-20](#)
 LED color codes, [B-3](#), [B-4](#)
 location in the rack, [C-1](#), [D-1](#), [E-1](#)
 monitoring, [15-1](#)
 spare parts, [B-6](#), [B-9](#), [B-13](#)
 upgrading base image, [10-19](#)
 services, starting, [10-29](#)
 setsmpriority command, [7-15](#)
 setup-root-ssh script, [11-12](#), [13-32](#), [14-2](#)
 SetupASR step, [10-29](#)
 SetupInstall step, [10-29](#)
 SetupKerberos step, [10-29](#)
 SetupMySQL step, [10-29](#)
 SetupPuppet step, [10-29](#)
 shipping brackets, [6-5](#)
 shipping dimensions, [2-4](#), [6-3](#)
 shipping kit, [6-3](#)
 shipping weight, [2-4](#)
 show command, [7-11](#)
 showunhealthy command, [7-15](#)
 showusb utility, [7-24](#), [13-33](#)
 showvlan utility, [13-34](#)
 showvnics utility, [13-34](#)
 shutdown, [11-2](#)
 shutdown command, [11-2](#)
 single-phase PDUs, [2-7](#), [6-2](#), [6-13](#), [11-14](#)
 site readiness tasks, [2-14](#)
 slot numbers (disk drives), [12-3](#)
 slots, disk drives, [12-2](#)
 SMTP alerts, [5-1](#)
 SNMP alerts, [5-1](#)
 software configuration
 querying, [13-7](#)
 software image
 downgrading, [10-8](#)
 upgrading, [10-8](#)
 software installation
 multiple racks, [10-8](#)
 upgrading, [10-10](#)
 using Mammoth, [10-6](#)
 software port assignments, [3-10](#)
 software profiles, acquiring from servers, [7-22](#)
 software source code, third party, [10-30](#)
 software validation, [13-2](#), [13-22](#)

sound levels, [2-1](#), [2-3](#)
 South American PDU specifications, [2-7](#)
 spares kit, [1-4](#), [6-5](#)
 spine switches
 cabling, [F-3](#)
 configuring, [7-12](#)
 replacing, [11-7](#)
 splitter cables, [3-7](#)
 spring nuts, [6-3](#)
 SSH
 removing passwordless, [13-30](#)
 setting up passwordless, [13-32](#)
 SSH keys, checking rack distribution, [7-20](#)
 stabilizing feet, [6-6](#)
 starter rack
 component configuration, [7-1](#)
 expanding, [8-1](#)
 physical layout, [C-1](#), [D-1](#)
 StartHadoopServices step, [10-29](#)
 starting KVM, [7-2](#)
 starting or stopping Big Data SQL, [13-10](#)
 startup order, resetting, [13-31](#)
 static discharges, [2-12](#)
 static electricity, [6-2](#)
 status indicators, [B-2](#)
 storage servers
 replacement parts, [B-6](#)
 Subnet Manager, [9-2](#), [9-4](#), [11-11](#)
 Subnet Manager Master, [9-5](#)
 Subnet Managers, [11-11](#)
 subnet mask factory settings, [3-9](#)
 Sun Network 10 GbE Switch 72p, [3-7](#)
 Sun Network QDR InfiniBand Gateway Switch,
 [11-12](#), [13-23](#), [13-30](#)
 Sun Server components, [1-2](#)
 support accounts, [5-2](#)
 support services, [13-21](#)
 swap partition, restoring, [12-16](#)
 swapping, [12-4](#), [13-26](#)
 switch LEDs, InfiniBand switch LEDs, [B-2](#)
 switch ports, [E-7](#)
 switches
 location in the rack, [C-1](#), [D-1](#), [E-1](#)
 monitoring, [15-1](#)
 naming conventions, [4-6](#)
 replacement parts, [B-23](#)
 replacing, [11-7](#)
 setting host names, [7-14](#)
 spare parts, [B-22](#)
 symbolic link names (disk drives), [12-2](#)
 synchronizing servers, [13-6](#)

T

Taiwanese PDU specifications, [2-7](#)

tee utility, [8-7](#)
 telnet access, configuring Ethernet switch, [7-8](#)
 temperature
 acceptable range, [2-12](#)
 operating range, [2-11](#)
 test events (ASR), [5-5](#)
 testrule setting (ASR), [5-5](#)
 third-party license location, [10-29](#)
 third-party software source code, [10-30](#)
 three-phase PDUs, [2-7](#), [2-9](#), [11-14](#), [11-15](#)
 thresholds for PDUs, [11-13](#)
 time synchronization, [13-6](#)
 time zone
 configuration setting, [4-5](#)
 configuring Ethernet switch, [7-9](#)
 configuring InfiniBand switches, [7-14](#)
 setting summer time, [7-9](#)
 top clearance (rack), [2-4](#)
 trap destinations (ASR), [5-4](#)
 turning off, [11-2](#)
 turning on, [11-2](#)

U

unpacking instructions, [6-4](#)
 upgrade history, [13-29](#)
 upgrading base image, [10-19](#)
 upgrading the factory image, [10-8](#)
 upgrading the software, [10-10](#)
 UPS, [2-11](#)
 USB flash drive, [7-23](#), [13-33](#)
 users, hdfs and mapred, [10-30](#)

V

validating a cluster, [13-2](#)
 validating hardware and software, [13-22](#)
 validating network configuration, [13-5](#)
 validating the server software, [13-5](#)
 validation checks after reboot, [6-14](#)
 ventilation requirements, [2-13](#)
 verifying
 database server configuration, [11-6](#)
 firmware, [B-8](#)
 hardware, [B-8](#)
 version command, [7-15](#)
 virtual drive configuration, [11-6](#)
 VLAN, [4-13](#), [4-17](#)
 VLANs
 configuration, [13-34](#)
 configuration example, [7-5](#)
 default configuration, [3-1](#)
 VNICS
 assigning to gateway switches, [3-6](#)
 configuration, [13-34](#)

VNICs (*continued*)
 default configuration, [3-1](#)
 re-creating, [13-23](#)

W

warnings, [11-4](#)
watcher logs (OS), [13-17](#)
weight

weight (*continued*)
 rack, [2-1–2-3](#)
 shipping, [2-4](#)
wrenches, [6-3](#)
wrist strap, [6-4](#)

X

X6-2 and X5-2 rack cabling tables, [C-1](#)