

Oracle® Big Data Appliance

Owner's Guide

Release 2 (2.1)

E40655-03

July 2013

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

Oracle Big Data Appliance Owner's Guide, Release 2 (2.1)

E40655-03

Copyright © 2012, 2013, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

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

Contents

| | |
|--|-------|
| Preface | xv |
| Audience | xv |
| Documentation Accessibility | xv |
| Related Documentation | xv |
| Conventions | xvi |
| Backus-Naur Form Syntax | xvi |
| | |
| Changes in This Release for Oracle Big Data Appliance Owner's Guide | xvii |
| Changes in Oracle Big Data Appliance Release 2 (2.1) | xvii |
| Changes in Oracle Big Data Appliance Release 2 (2.0.1) | xviii |
| Changes in Oracle Big Data Appliance Release 2 (2.0) | xviii |
| | |
| Part I Preinstallation | |
| | |
| 1 Introduction to Oracle Big Data Appliance | |
| About Oracle Big Data Appliance | 1-1 |
| Oracle Big Data Appliance Rack Components | 1-1 |
| Server Components | 1-2 |
| Spare Parts Kit Components | 1-3 |
| Oracle Big Data Appliance Management Software | 1-3 |
| Oracle Big Data Appliance Restrictions on Use | 1-4 |
| | |
| 2 Site Requirements for Oracle Big Data Appliance | |
| General Environmental Requirements | 2-1 |
| Space Requirements | 2-3 |
| Receiving and Unpacking Requirements | 2-3 |
| Maintenance Access Requirements | 2-4 |
| Flooring Requirements | 2-4 |
| Electrical Power Requirements | 2-4 |
| PDUs for Oracle Big Data Appliance | 2-5 |
| Facility Power Requirements | 2-8 |
| Circuit Breaker Requirements | 2-9 |
| Grounding Guidelines | 2-9 |
| Temperature and Humidity Requirements | 2-9 |

| | |
|---|------|
| Ventilation and Cooling Requirements..... | 2-10 |
| Network Connection Requirements..... | 2-12 |
| Ensuring That the Site Is Ready..... | 2-12 |

3 Understanding the Network Requirements

| | |
|---|-----|
| Overview of Network Requirements..... | 3-1 |
| Default Network Connections..... | 3-1 |
| Minimum Requirements for the Networks..... | 3-2 |
| Network Diagram..... | 3-3 |
| Cabling the Client Network..... | 3-4 |
| How the Servers Connect to the Gateway Switches..... | 3-4 |
| Using Splitter Cables for Connecting to the 40 Gbps Physical Ports..... | 3-5 |
| What About Data Centers Without a 10 GbE Infrastructure?..... | 3-5 |
| Configuring Multiple Racks into a Single Hadoop Cluster..... | 3-6 |
| Factory Network Settings..... | 3-6 |
| Port Assignments for Oracle Big Data Appliance Software..... | 3-7 |
| Preparing the Network for Oracle Big Data Appliance..... | 3-9 |

4 Using the Oracle Big Data Appliance Configuration Utility

| | |
|--|-----|
| Overview of the Oracle Big Data Appliance Configuration Utility..... | 4-1 |
| Using the Oracle Big Data Appliance Configuration Utility..... | 4-1 |
| Using OpenOffice Spreadsheet..... | 4-2 |
| Generating the Configuration Files..... | 4-2 |
| About the Configuration Files..... | 4-3 |
| Welcome Page..... | 4-5 |
| Network Configuration..... | 4-5 |
| Software Configuration..... | 4-8 |

5 Setting Up Auto Service Request

| | |
|---|-----|
| Understanding Auto Service Request..... | 5-1 |
| Getting Ready to Install ASR..... | 5-2 |
| Installing ASR Manager..... | 5-3 |
| Verifying the ASR Manager Installation..... | 5-3 |
| Configuring the Trap Destinations on Oracle Big Data Appliance..... | 5-4 |
| Verifying ASR Assets..... | 5-4 |
| Validating Auto Service Request on Oracle Big Data Appliance..... | 5-5 |
| Troubleshooting ASR..... | 5-6 |

Part II Hardware Installation and Configuration

6 Installing Oracle Big Data Appliance at the Site

| | |
|---|-----|
| Reviewing Safety Guidelines..... | 6-1 |
| Unpacking Oracle Big Data Appliance..... | 6-2 |
| Contents of the Shipping Kit..... | 6-3 |
| Removing Oracle Big Data Appliance from the Shipping Crate..... | 6-4 |
| Placing Oracle Big Data Appliance in Its Allocated Space..... | 6-5 |

| | |
|--|-------------|
| Moving Oracle Big Data Appliance..... | 6-5 |
| Stabilizing Oracle Big Data Appliance..... | 6-6 |
| Attaching a Ground Cable (Optional)..... | 6-9 |
| Powering On the System for the First Time..... | 6-10 |
| Inspecting the Oracle Big Data Appliance Rack After It Is in Place | 6-10 |
| Connecting Power Cords | 6-10 |
| Powering On Oracle Big Data Appliance | 6-12 |

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

| | |
|---|-------------|
| Configuring the KVM Switch..... | 7-1 |
| Starting the KVM Switch..... | 7-2 |
| Connecting the KVM Switch to the Management Network..... | 7-2 |
| Checking the KVM Firmware Version..... | 7-3 |
| Configuring the KVM Switch to Access the Servers..... | 7-3 |
| Accessing a Server by Using the KVM Switch..... | 7-3 |
| Configuring the Cisco Ethernet Switch | 7-4 |
| Scope of the Configuration | 7-4 |
| Prerequisites for Configuring the Ethernet Switch | 7-4 |
| Configuring the Ethernet Switch on the Customer Network..... | 7-5 |
| Setting Up Telnet Access on the Ethernet Switch..... | 7-6 |
| Configuring SSH on the Ethernet Switch | 7-7 |
| Setting the Clock and Time Zone on the Ethernet Switch | 7-8 |
| Configuring the NTP Servers for the Ethernet Switch | 7-9 |
| Disabling the Spanning Tree | 7-9 |
| Verifying the Ethernet Configuration | 7-10 |
| Configuring the InfiniBand Leaf and Spine Switches | 7-11 |
| Configuring an InfiniBand Switch..... | 7-11 |
| Setting the Time Zone and Clock on an InfiniBand Switch..... | 7-13 |
| Checking the Health of an InfiniBand Switch..... | 7-13 |
| Checking the Sun Network QDR InfiniBand Gateway Switches | 7-16 |
| Configuring the Power Distribution Units | 7-17 |
| Connecting the PDUs to the Network..... | 7-17 |
| Updating the PDU Firmware Version | 7-17 |
| Configuring the Oracle Big Data Appliance Servers | 7-18 |
| Configuring the Network | 7-20 |
| Copying the Configuration Files to Oracle Big Data Appliance..... | 7-21 |
| Starting the Network Configuration | 7-22 |
| Connecting to the Administrative Network | 7-24 |
| Connecting to the Client Network..... | 7-24 |
| Completing the Network Configuration | 7-25 |
| Checking the Virtual Interface Cards (VNICs)..... | 7-26 |
| Verifying Access to the Network Services..... | 7-27 |
| Configuring Multiple Oracle Big Data Appliance Racks | 7-28 |
| Configuring and Cabling Multiple Racks..... | 7-28 |
| Validating a Multirack InfiniBand Network..... | 7-30 |

8 Connecting Multiple Oracle Big Data Appliance Racks

| | |
|---|-----|
| Extending a Rack by Adding Another Rack | 8-1 |
| Cabling Two Racks Together..... | 8-2 |
| Cabling Several Racks Together | 8-3 |

Part III Software Installation and Configuration

9 Installing the Oracle Big Data Appliance Software

| | |
|---|------|
| About the Mammoth Utility | 9-1 |
| Installing the Software on a New Rack | 9-2 |
| Installing the Software on a Single or Primary Rack | 9-2 |
| Adding Servers to a Cluster | 9-3 |
| What If an Error Occurs During the Installation?..... | 9-4 |
| Upgrading the Software on Oracle Big Data Appliance | 9-5 |
| About the Changes from Version 1.1 to Version 2.1 | 9-5 |
| Upgrading from Software Version 2.0 or 1.1 | 9-5 |
| Changing the Configuration of Optional Software | 9-7 |
| Reinstalling the Base Image | 9-9 |
| Reimaging a Single Oracle Big Data Appliance Server | 9-9 |
| Reimaging an Oracle Big Data Appliance Rack | 9-10 |
| Reimaging an Oracle Big Data Appliance Cluster | 9-12 |
| Mammoth Utility Syntax | 9-13 |
| Mammoth Utility Options | 9-14 |
| Mammoth Installation Steps..... | 9-14 |
| Mammoth Reconfiguration Utility Syntax..... | 9-18 |
| Oracle Big Data Appliance Base Imaging Utilities | 9-19 |
| makebdamage..... | 9-19 |
| reimagecluster | 9-19 |
| reimagerack..... | 9-20 |

Part IV Maintenance

10 Maintaining Oracle Big Data Appliance

| | |
|---|------|
| Monitoring the Ambient Temperature of Servers | 10-1 |
| Powering On and Off Oracle Big Data Appliance | 10-2 |
| Nonemergency Power Procedures | 10-2 |
| Emergency Power-Off Considerations | 10-3 |
| Cautions and Warnings..... | 10-3 |
| Adding Memory to the Servers | 10-4 |
| Adding Memory to Sun Server X3-2L Servers..... | 10-4 |
| Adding Memory to Sun Fire X4270 M2 Servers | 10-5 |
| Maintaining the Physical Disks of Servers..... | 10-5 |
| Verifying the Server Configuration | 10-5 |
| Replacing a Server Disk..... | 10-6 |
| Overview of the Disk Replacement Process..... | 10-7 |
| About Disk Drive Identifiers | 10-7 |

| | |
|---|-------|
| Prerequisites for Replacing a Working Disk | 10-9 |
| What If a Server Fails to Restart? | 10-10 |
| Replacing a Disk Drive | 10-11 |
| Identifying the Function of a Disk Drive | 10-13 |
| Configuring an Operating System Disk..... | 10-14 |
| Configuring an Oracle NoSQL Database Disk | 10-20 |
| Configuring an HDFS Disk..... | 10-23 |
| Verifying the Disk Configuration | 10-24 |
| Maintaining the InfiniBand Network | 10-25 |
| Backing Up and Restoring Oracle ILOM Settings..... | 10-25 |
| Replacing a Failed InfiniBand Switch | 10-26 |
| Verifying InfiniBand Network Operation | 10-27 |
| Understanding the Network Subnet Manager Master | 10-29 |
| Changing the Number of Connections to a Gateway Switch | 10-30 |
| Changing the NTP Servers | 10-31 |
| Monitoring the PDU Current | 10-32 |

11 Monitoring the Health of Oracle Big Data Appliance

12 Using the dcli Utility

| | |
|---|------|
| Overview of the dcli Utility | 12-1 |
| Setting Up Passwordless SSH | 12-1 |
| Basic Use of dcli..... | 12-2 |
| dcli Syntax | 12-3 |
| dcli Return Values | 12-5 |
| dcli Examples | 12-5 |

13 Using Oracle Integrated Lights Out Manager

| | |
|--|------|
| Oracle ILOM Overview | 13-1 |
| Oracle ILOM Interfaces | 13-2 |
| Oracle ILOM Users | 13-2 |
| Administrative Network Diagram | 13-2 |
| Oracle ILOM IP Addresses for Oracle Big Data Appliance Components | 13-3 |
| Connecting to Oracle ILOM Using the Network | 13-3 |
| Connecting to the Command-Line Interface..... | 13-3 |
| Connecting to the Web GUI..... | 13-3 |
| Connecting to Oracle ILOM Using a Serial Connection | 13-4 |
| Connecting to the Oracle ILOM of a Server | 13-4 |
| Connecting to the Oracle ILOM of an InfiniBand Switch | 13-5 |

Part V Appendixes

A Error Messages for Oracle Big Data Appliance

| | |
|---|-----|
| Auto Service Request Alerts: DISKALRT-02001 to DISKALRT-02006 | A-1 |
|---|-----|

B Replacement Units

| | |
|--|------|
| Understanding Repair Categories..... | B-1 |
| Disk Controller Batteries | B-2 |
| LED Status Descriptions | B-2 |
| Sun Server LEDs..... | B-3 |
| Sun InfiniBand Switch LEDs | B-4 |
| Parts for Oracle Big Data Appliance Servers | B-5 |
| Parts for Sun Datacenter InfiniBand Switch 36 | B-12 |
| Parts for Sun Network QDR InfiniBand Gateway Switches | B-13 |
| Parts for the Cisco Ethernet Switch | B-14 |
| Parts for Power Distribution Units | B-14 |
| Cables and Cable Accessories for an Oracle Big Data Appliance Rack | B-15 |

C In-Rack Cabling Tables for Oracle Big Data Appliance X3-2

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

D In-Rack Cabling Tables for Oracle Big Data Appliance

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

E Multirack Cabling Tables

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

Eight-Rack Cabling E-19

Index

List of Examples

| | | |
|-----|---|------|
| 3-1 | Output from preinstall-checkip.sh | 3-10 |
| 7-1 | Sample Output from networksetup-one | 7-23 |
| 7-2 | Sample Output from networksetup-two | 7-25 |
| 7-3 | Multirack Configuration File sample-multi-rack.json | 7-31 |
| 9-1 | Mammoth Utility Syntax Examples | 9-13 |

List of Figures

| | | |
|------|---|-------|
| 2-1 | Low-Voltage, Single-Phase Power Connector | 2-6 |
| 2-2 | Low-Voltage Three-Phase Power Connector | 2-7 |
| 2-3 | High-Voltage Single-Phase Power Connector | 2-7 |
| 2-4 | High-Voltage Three-Phase Power Connector | 2-8 |
| 2-5 | Typical Data Center Configuration for Perforated Floor Tiles | 2-12 |
| 3-1 | Network Diagram for Oracle Big Data Appliance X3-2 | 3-3 |
| 3-2 | Network Diagram for Oracle Big Data Appliance | 3-4 |
| 4-1 | Installation Template | 4-5 |
| 5-1 | Auto Service Request Network Connections | 5-3 |
| 6-1 | Oracle Big Data Appliance in the Shipping Crate | 6-3 |
| 6-2 | Carefully Push Oracle Big Data Appliance from Behind | 6-5 |
| 6-3 | Stabilizing the Oracle Big Data Appliance Rack Using Leveling Feet | 6-6 |
| 6-4 | Location of Mounting Brackets on the Rack | 6-7 |
| 6-5 | Bottom View of the Oracle Big Data Appliance Rack | 6-7 |
| 6-6 | Base Position of the Rack to Measure (in mm) | 6-8 |
| 6-7 | Securing the Oracle Big Data Appliance Rack Using Mounting Brackets | 6-9 |
| 6-8 | Earth Ground Attachment Bolt Location | 6-10 |
| 6-9 | Location of Power Cord Cable Ties | 6-11 |
| 6-10 | Power Cord Routing Out the Bottom | 6-12 |
| 6-11 | Power Cord Routing Out the Top | 6-12 |
| 6-12 | PDU Switch Locations | 6-13 |
| 10-1 | Oracle ILOM Configuration Backup | 10-26 |
| 13-1 | Administrative Network in Oracle Big Data Appliance | 13-2 |
| 13-2 | Oracle ILOM Web Interface on an Oracle Big Data Appliance X3-2 Server | 13-4 |
| C-1 | Oracle Big Data Appliance X3-2 Full Rack Layout | C-2 |
| C-2 | Oracle Big Data Appliance X3-2 Starter Rack Layout | C-3 |
| C-3 | Oracle Big Data Appliance X3-2 Starter Rack With In-Rack Expansion | C-4 |
| C-4 | Sun Network QDR InfiniBand Gateway Ports | C-8 |
| D-1 | Rack Layout of Oracle Big Data Appliance | D-2 |
| D-2 | Sun Network QDR InfiniBand Gateway Ports | D-7 |
| E-1 | Connections Between Spine Switch and Leaf Switches | E-2 |
| E-2 | Connections Between the Servers and the Leaf Switches | E-2 |
| E-3 | Connecting Switches Across Two Racks | E-3 |
| E-4 | Connecting Switches Across Multiple Racks | E-3 |

List of Tables

| | | |
|------|--|-------|
| 1-1 | Oracle Big Data Appliance Rack Components | 1-2 |
| 1-2 | Sun Server X3-2L Server Components..... | 1-2 |
| 1-3 | Sun Fire X4270 M2 Server Components | 1-2 |
| 1-4 | Bundled Components and Accessories | 1-3 |
| 1-5 | Bundled Cables | 1-3 |
| 1-6 | In-Rack Spare Cables | 1-3 |
| 2-1 | Environmental Requirements for Oracle Big Data Appliance Full Rack | 2-1 |
| 2-2 | Environmental Requirements for Oracle Big Data Appliance Starter Rack..... | 2-2 |
| 2-3 | Access Route Requirements for Oracle Big Data Appliance | 2-4 |
| 2-4 | Maintenance Access Requirements for Oracle Big Data Appliance..... | 2-4 |
| 2-5 | Floor Load Requirements for Oracle Big Data Appliance | 2-4 |
| 2-6 | Low-Voltage 15 kVA Single-Phase PDUs | 2-5 |
| 2-7 | Low-Voltage 15 kVA Three-Phase PDUs | 2-6 |
| 2-8 | High-Voltage 15 kVA Single-Phase PDUs | 2-7 |
| 2-9 | High-Voltage 15 kVA Three-Phase PDUs | 2-8 |
| 2-10 | Temperature, Humidity, and Altitude Requirements..... | 2-10 |
| 3-1 | Default IP Addresses for Oracle Big Data Appliance..... | 3-6 |
| 3-2 | Oracle Big Data Appliance Port Numbers | 3-7 |
| 3-3 | CDH Port Numbers | 3-7 |
| 4-1 | General Properties for the Installation..... | 4-5 |
| 4-2 | General Properties for the Rack..... | 4-5 |
| 4-3 | Network Properties for the Rack | 4-6 |
| 4-4 | Network Properties for Servers | 4-7 |
| 4-5 | Network Properties for Switches..... | 4-7 |
| 4-6 | Network Properties for Power Distribution Units..... | 4-7 |
| 4-7 | Advanced Network Configuration Properties | 4-7 |
| 4-8 | Installed Components | 4-8 |
| 4-9 | Oracle Auto Service Request..... | 4-8 |
| 4-10 | Users, Groups, and Passwords | 4-8 |
| 4-11 | Cloudera Manager Email Alerts | 4-9 |
| 4-12 | Oracle Enterprise Manager Cloud Control | 4-9 |
| 10-1 | Disk Drive Identifiers | 10-8 |
| 10-2 | Mount Points | 10-8 |
| 10-3 | Threshold Values for Single-Phase, Low-Voltage PDU | 10-32 |
| 10-4 | Threshold Values for Three-Phase, Low-Voltage PDU..... | 10-32 |
| 10-5 | Threshold Values for Single-Phase, High-Voltage PDU | 10-32 |
| 10-6 | Threshold Values for Three-Phase, High-Voltage PDU..... | 10-33 |
| 12-1 | dcli Options..... | 12-4 |
| B-1 | Oracle Big Data Appliance Repair Categories..... | B-1 |
| B-2 | LED Status Descriptions | B-3 |
| B-3 | Sun Datacenter InfiniBand Switch 36 LED Status Descriptions | B-4 |
| B-4 | Replacement Parts for a Sun Server X3-2L Server | B-6 |
| B-5 | Replacement Parts for Sun Fire X4270 M2 Servers | B-9 |
| B-6 | Replacement Parts for the Sun Datacenter InfiniBand Switch 36 | B-13 |
| B-7 | Replacement Parts for the Sun Network QDR InfiniBand Gateway Switch..... | B-13 |
| B-8 | Replacement Parts for Cisco Catalyst Ethernet Switch | B-14 |
| B-9 | Replacement Parts for Power Distribution Units..... | B-15 |
| B-10 | Replacement Cables for an Oracle Big Data Appliance Rack..... | B-16 |
| C-1 | Administrative Gigabit Ethernet Cabling | C-4 |
| 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-8 |
| D-1 | KVM Cabling | D-2 |
| D-2 | Administrative Gigabit Ethernet Cabling | D-3 |
| D-3 | Oracle ILOM Cabling | D-4 |
| D-4 | Single-Phase PDU Cabling | D-5 |
| D-5 | Three-Phase PDU Cabling | D-6 |
| D-6 | InfiniBand Switch-to-Switch Cabling | D-6 |
| D-7 | InfiniBand Switch-to-Server Cabling | D-7 |
| E-1 | Leaf Switch Connections for the First Rack in a Two-Rack System | E-4 |
| E-2 | Leaf Switch Connections for the Second Rack in a Two-Rack System | E-4 |
| E-3 | Leaf Switch Connections for the First Rack in a Three-Rack System | E-5 |
| E-4 | Leaf Switch Connections for the Second Rack in a Three-Rack System | E-6 |
| E-5 | Leaf Switch Connections for the Third Rack in a Three-Rack System | E-6 |
| E-6 | Leaf Switch Connections for the First Rack in a Four-Rack System..... | E-7 |
| E-7 | Leaf Switch Connections for the Second Rack in a Four-Rack System | E-7 |
| E-8 | Leaf Switch Connections for the Third Rack in a Four-Rack System..... | E-8 |
| E-9 | Leaf Switch Connections for the Fourth Rack in a Four-Rack System..... | E-8 |
| E-10 | Leaf Switch Connections for the First Rack in a Five-Rack System..... | E-9 |
| E-11 | Leaf Switch Connections for the Second Rack in a Five-Rack System..... | E-10 |
| E-12 | Leaf Switch Connections for the Third Rack in a Five-Rack System..... | E-10 |
| E-13 | Leaf Switch Connections for the Fourth Rack in a Five-Rack System..... | E-11 |
| E-14 | Leaf Switch Connections for the Fifth Rack in a Five-Rack System | E-11 |
| E-15 | Leaf Switch Connections for the First Rack in a Six-Rack System..... | E-12 |
| E-16 | Leaf Switch Connections for the Second Rack in a Six-Rack System | E-13 |
| E-17 | Leaf Switch Connections for the Third Rack in a Six-Rack System..... | E-13 |
| E-18 | Leaf Switch Connections for the Fourth Rack in a Six-Rack System..... | E-14 |
| E-19 | Leaf Switch Connections for the Fifth Rack in a Six-Rack System | E-14 |
| E-20 | Leaf Switch Connections for the Sixth Rack in a Six-Rack System..... | E-15 |
| E-21 | Leaf Switch Connections for the First Rack in a Seven-Rack System..... | E-15 |
| E-22 | Leaf Switch Connections for the Second Rack in a Seven-Rack System..... | E-16 |
| E-23 | Leaf Switch Connections for the Third Rack in a Seven-Rack System..... | E-17 |
| E-24 | Leaf Switch Connections for the Fourth Rack in a Seven-Rack System..... | E-17 |
| E-25 | Leaf Switch Connections for the Fifth Rack in a Seven-Rack System | E-18 |
| E-26 | Leaf Switch Connections for the Sixth Rack in a Seven-Rack System..... | E-18 |
| E-27 | Leaf Switch Connections for the Seventh Rack in a Seven-Rack System | E-19 |
| E-28 | Leaf Switch Connections for the First Rack in an Eight-Rack System | E-19 |
| E-29 | Leaf Switch Connections for the Second Rack in an Eight-Rack System..... | E-20 |
| E-30 | Leaf Switch Connections for the Third Rack in an Eight-Rack System | E-21 |
| E-31 | Leaf Switch Connections for the Fourth Rack in an Eight-Rack System | E-21 |
| E-32 | Leaf Switch Connections for the Fifth Rack in an Eight-Rack System | E-22 |
| E-33 | Leaf Switch Connections for the Sixth Rack in an Eight-Rack System | E-22 |
| E-34 | Leaf Switch Connections for the Seventh Rack in an Eight-Rack System | E-23 |
| E-35 | Leaf Switch Connections for the Eighth Rack in an Eight-Rack System..... | E-23 |

Preface

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

This preface contains the following topics:

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

Audience

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

Documentation Accessibility

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

Access to Oracle Support

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

Related Documentation

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

http://docs.oracle.com/cd/E40622_01/index.htm

- Sun Server X3-2L library:
http://docs.oracle.com/cd/E23393_01/index.html

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

Conventions

The following text conventions are used in this document:

| Convention | Meaning |
|-----------------|--|
| boldface | Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary. |
| <i>italic</i> | Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values. |
| 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 This Release for Oracle Big Data Appliance Owner's Guide

This preface contains:

- [Changes in Oracle Big Data Appliance Release 2 \(2.1\)](#)
- [Changes in Oracle Big Data Appliance Release 2 \(2.0.1\)](#)
- [Changes in Oracle Big Data Appliance Release 2 \(2.0\)](#)

Changes in Oracle Big Data Appliance Release 2 (2.1)

The following are changes in *Oracle Big Data Appliance Owner's Guide* for Oracle Big Data Appliance release 2 (2.1).

New Features

- Starter Rack and In-Rack Upgrade Kit

Oracle Big Data Appliance is now available as a starter rack with six servers instead of the full rack of 18 servers. An upgrade kit expands the starter rack to 12 servers, and a second upgrade kit expands it to 18 servers.

See "[Oracle Big Data Appliance Rack Components](#)" on page 1-1 and "[Oracle Big Data Appliance Rack Layout](#)" on page C-1.

- Multiple Clusters in a Rack

You can now configure a full Oracle Big Data Appliance rack with multiple Hadoop clusters. The smallest cluster size is six servers. Larger clusters must have multiples of three servers.

This support affects the *Oracle Big Data Appliance Configuration Worksheets*, the network configuration, and the Mammoth software installation.

See [Chapter 7](#).

- Cloudera Software Upgrade

Mammoth 2.1.1 installs Cloudera's Distribution including Apache Hadoop (CDH) 4.2.0 and Cloudera Manager 4.5.1.

You can upgrade from Oracle Big Data Appliance 2.0 or 1.1 software.

See "[Upgrading the Software on Oracle Big Data Appliance](#)" on page 9-5.

- Oracle NoSQL Database Enterprise Edition

If you have a license, you can have Mammoth install Oracle NoSQL Database Enterprise Edition instead of the Community Edition.

See ["Software Configuration"](#) on page 4-8.

- The `dcli` utility uses the Ethernet IP addresses that are stored in the `BdaDeploy.json` and `BdaShip.json` files. The `rack-hosts-infiniband` file is obsolete. See [Chapter 12](#).

Changes in Oracle Big Data Appliance Release 2 (2.0.1)

The following are changes in *Oracle Big Data Appliance Owner's Guide* for Oracle Big Data Appliance release 2 (2.0.1).

New Features

- Mammoth Upgrade Support
The Mammoth Utility supports upgrades from software version 1.1 to 2.0.1. See ["Upgrading the Software on Oracle Big Data Appliance"](#) on page 9-5.
- Oracle Enterprise Manager Support
Agents run on Oracle Big Data Appliance to support the Enterprise Manager system monitoring plug-in. See ["Software Configuration"](#) on page 4-8.

Changes in Oracle Big Data Appliance Release 2 (2.0)

The following are changes in *Oracle Big Data Appliance Owner's Guide* for Oracle Big Data Appliance release 2 (2.0).

New Features

- Oracle Big Data Appliance X3-2
Oracle Big Data Appliance racks are now shipped with a different hardware configuration named Oracle Big Data Appliance X3-2. The Sun Server X3-2L-based racks supersede the Sun Fire X4270 M2-based racks. All versions of the Oracle Big Data Appliance software run on both hardware configurations. See ["Oracle Big Data Appliance Rack Components"](#) on page 1-1.
- Cloudera's Distribution including Apache Hadoop Version 4 (CDH4)
CDH4 replaces CDH3. CDH4 provides automatic failover of the NameNode and uses different services than CDH3. See ["Port Assignments for Oracle Big Data Appliance Software"](#) on page 3-7 for the ports used by CDH4. See *Oracle Big Data Appliance Software User's Guide* for more information about using CDH4.

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:

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

Introduction to Oracle Big Data Appliance

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

This chapter contains these sections:

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

About Oracle Big Data Appliance

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

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

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

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

Oracle Big Data Appliance Rack Components

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

[Table 1-1](#) lists the components in an Oracle Big Data Appliance rack. An in-rack expansion kit contains six Sun Server X3-2L servers.

Table 1–1 Oracle Big Data Appliance Rack Components

| Full Rack | Starter Rack | Description |
|-----------|--------------|---|
| 18 | 6 | 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 with Sun Fire X4270 M2) |
| 2 | 2 | Power distribution units |
| 0 or 1 | 0 | KVM (only in Sun Fire X4270 M2-based racks) |

Server Components

Table 1–2 lists the components of a Sun Server X3-2L.

Table 1–2 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 |
| 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 1–3 lists the components of a Sun Fire X4270 M2.

Table 1–3 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 |

Table 1–3 (Cont.) Sun Fire X4270 M2 Server Components

| Quantity | Description |
|----------|--|
| 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 |

Spare Parts Kit Components

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

Table 1–4 Bundled Components and Accessories

| Part Number | Quantity | Description |
|--------------|----------|---|
| 390-0476-03 | 1 or 2 | 3 TB 3.5-inch 7200 RPM drives (1 for a starter rack or 2 for a full rack) |
| 590-896-501B | 1 | LCD Accessory Kit (Sun Fire X4270 M2-based racks only) |
| | 1 | Sun Server X3-2L Documentation Kit |
| | 1 | Sun Rack II Accessory Kit |
| 53-2332-xx | 1 | Cisco Switch Documents and Accessory Kit |

Table 1–5 Bundled Cables

| Part Number | Length | Quantity | Description |
|-------------|----------|----------|---|
| 530-4446-01 | 5 meters | 10 | QSFP passive copper cable for multiracking |
| 530-4445-01 | 3 meters | 6 | QSFP passive copper cable for multiracking |
| 350-1519-01 | -- | 3 | ASSY, NM2 serial cable sets for the InfiniBand switches |

[Table 1–6](#) lists the spare parts that are secured to the center section of an Oracle Big Data Appliance rack.

Table 1–6 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 |

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 [Chapter 11](#) except where directed otherwise.

- Oracle Integrated Lights Out Manager (Oracle ILOM) consists of preinstalled, dedicated hardware and software that you can use to manage and monitor the servers and switches in an Oracle Big Data Appliance rack. See [Chapter 13](#).
- The `setup-root-ssh` utility sets up passwordless SSH for the `root` user among all the servers in an Oracle Big Data Appliance rack. See "[Setting Up Passwordless SSH](#)" on page 12-1.
- The `dcli` utility executes commands across a group of servers on Oracle Big Data Appliance and returns the output. See [Chapter 12](#).
- The `bdacheckcluster` utility checks the health of the cluster.
- The `bdacheckib` utility checks the private InfiniBand network.
- The `bdachecknet` utility checks the network configuration.
- The `iblinkinfo` utility lists the connections in the InfiniBand network.
- The `listlinkup` utility identifies the active Ethernet ports.
- The `showvlan` utility lists the virtual local area networks (VLANs) configured on a Sun Network QDR InfiniBand Gateway switch.
- The `shownics` utility lists the virtual network interface cards (VNICs) created on a Sun Network QDR InfiniBand Gateway switch.

All end-user software is installed on site using the Mammoth Utility. See [Chapter 9](#).

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

Oracle Big Data Appliance Restrictions on Use

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

- Oracle Big Data Appliance hardware cannot be modified or customized, with two exceptions:
 1. Replace the administrative 48-port Cisco 4948 Ethernet switch. Customers can do the following:
 - Replace the Gigabit Ethernet switch with an equivalent 1U 48-port Gigabit Ethernet switch that conforms to their internal data center network standards. This replacement must be performed by customers at their own expense and labor, after delivery of Oracle Big Data Appliance. If customers choose to make this change, then Oracle cannot make or assist with this change given the numerous possible scenarios involved, and it is not included as part of the standard installation. Customers must supply the replacement hardware, and make or arrange for this change through other means.
 - Remove the Cat 5 cables connected to the Ethernet switch and connect them to the customer network through an external switch or patch panel. Customers must perform these changes at their expense and labor. In this case, the Cisco 4948 Ethernet switch in the rack can be turned off and disconnected from the data center network.

2. In a Sun Server X3-2L-based rack, install a top-of-rack (TOR) switch that supports both 10 GbE (or 40 GbE) and 1 GbE, such as the Sun Network 10 GbE Switch 72p, in rack slots 41 and 42. These slots are empty in the factory configuration. Customers must perform this installation at their expense and labor. It is not included as part of the standard installation.

This modification is not permitted in Sun Fire X4270 M2-based racks, which have no empty slots.

- Customers cannot update the firmware directly on Oracle Big Data Appliance servers. Oracle Big Data Appliance patches update the firmware of the server ILOM/BIOS, InfiniBand host channel adapters (HCA), disk controller host bus adapters (HBA), and hard drives (HDD).

The only exception is if the customer replaces a faulty server component (such as a disk drive) with a unit installed with firmware at a lower version than the faulty component. The upgrade must be performed using the [bdupdatefw](#) utility.

- Customers can update the firmware of the other components of Oracle Big Data Appliance:
 - Customers can update the IOS and firmware versions on the Cisco 4948 Ethernet switch to meet their data center requirements.
 - Customers can update the firmware of the InfiniBand switches if the updates comply with the validated versions documented in My Oracle Support ID 1528190.1.
 - Customers can update the firmware of the power distribution units (PDUs) if the updates comply with the validated versions documented in My Oracle Support ID 1528190.1.
 - Customers can update the firmware of the KVM switch and the KMM (keyboard, monitor, mouse) as needed (Sun Fire X4270 M2-based rack only).
- Customers can install additional software on Oracle Big Data Appliance servers. Oracle does not support questions or issues with the nonstandard modules. If a server fails, and Oracle suspects that the failure may have been caused by a nonstandard module, then Oracle Support may refer the customer to the vendor of the nonstandard module or ask that the issue be reproduced without the nonstandard module.

Modifying the server operating system other than by applying official patches and upgrades is not supported.

InfiniBand-related packages must always be maintained at the officially supported release.

You can install third-party software that complies with the following restrictions:

- The software must support the version of 64-bit Oracle Linux running on Oracle Big Data Appliance.
- The software must be related to the use of Cloudera's Distribution including Apache Hadoop (CDH) or Oracle NoSQL Database installed on the cluster, or to Information Technologies (IT) or data center requirements, such as monitoring software.
- The software must not overwrite, update, replace, or invalidate the software components already installed on Oracle Big Data Appliance. Run `bdachcksww` on each server where new software is installed, to ensure that the original software components are still valid. See [bdachcksww](#) on page 11-7.

- The software must not consume CPU, disk, or memory resources to the extent that it adversely affects the operating system, the CDH cluster, or Oracle NoSQL Database.

Customers are responsible for supporting all software on Oracle Big Data Appliance that they install in addition to the software installed by the Mammoth utility.

- Customers can use Cloudera Manager to stop and start services, but cannot use it to move services from one server to another. The Hadoop services such as NameNode and JobTracker must remain on the servers where they were installed by the Mammoth Utility.

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

- Customers cannot connect USB devices to Oracle Big Data Appliance servers except as documented in this guide. In those documented situations, the USB device should not draw more than 100 mA of power.
- The network ports on the servers can connect to external non-Sun servers using Internet small computer system interface (iSCSI) or network file system (NFS) protocols. However, the Fibre Channel over Ethernet (FCoE) protocol is not supported.
- Only switches specified for use in Oracle Big Data Appliance, Oracle Exadata Database Machine, and Oracle Exalogic Elastic Cloud can be connected to the InfiniBand network. Connecting third-party switches and other switches not used in these engineered systems is not supported.

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

See Also: *Oracle Big Data Appliance Site Checklists*

General Environmental Requirements

Table 2–1 shows the general environmental requirements for Oracle Big Data Appliance. The other sections in this chapter provide detailed information.

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

| Environmental Component | Oracle Big Data Appliance X3-2 | Oracle Big Data Appliance ¹ |
|---|--|--|
| Net Weight See Also: " Flooring Requirements " on page 2-4 | 816.5 kg (1800 lb) | 816.5 kg (1800 lb) |
| Acoustic levels | 8.3 Bel | 8.3 Bel |
| Power See Also: " Electrical Power Requirements " on page 2-4 | Maximum: 10 kW (10.2 kVA) Typical: 7 kW (7.15 kVA) | Maximum: 12.0 kW (12.25 kVA) Typical: 8.4 kW (8.6 kVA) ² |
| Cooling See Also: " Temperature and Humidity Requirements " on page 2-9 and " Ventilation and Cooling Requirements " on page 2-10 | Maximum: 34,142 BTU/hour (34,838 kJ/hour) Typical: 23,940 BTU/hour (24,420 kJ/hour) | Maximum: 40,971 BTU/hour (41,807 kJ/hour) Typical: 28,680 BTU/hour (29,265 kJ/hour) |

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

| Environmental Component | Oracle Big Data Appliance X3-2 | Oracle Big Data Appliance ¹ |
|---|--|--|
| Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements" on page 2-9, and "Ventilation and Cooling Requirements" on page 2-10 | Maximum: 1,573 CFM Typical: 1,103 CFM | Maximum: Approximately 1,886 CFM Typical: Approximately 1,340 CFM |
| IP addresses See Also: Chapter 3, "Understanding the Network Requirements." | 24 for Ethernet network 18 for InfiniBand network | 24 for Ethernet network 18 for InfiniBand network |
| Network drops See Also: Chapter 3, "Understanding the Network Requirements." | One network connection | Minimum of two network connections |
| External connectivity See Also: Chapter 3, "Understanding the Network Requirements." | 1 x 1 Gbps Ethernet ports 18 x 10 Gbps Ethernet ports | 2 x 1 Gbps Ethernet ports 18 x 10 Gbps Ethernet ports |

¹ Sun Fire X4270 M2-based rack

² Varies based on application load.

Table 2–2 shows the general environmental requirements for the starter rack and in-rack expansion kit.

Table 2–2 Environmental Requirements for Oracle Big Data Appliance Starter Rack

| Environmental Component | Starter Rack | Starter Rack with One In-Rack Expansion Kit |
|---|---|--|
| Net Weight See Also: "Flooring Requirements" on page 2-4 | 470.4 kg (1037 lb) | 635 kg (1400 lb) |
| Acoustic levels | 76 Bel | 79 Bel |
| Power See Also: "Electrical Power Requirements" on page 2-4 | Maximum: 4.2 kW Typical: 3 kW | Maximum: 7.7 kW) Typical: 5.4 kW ¹ |
| Cooling See Also: "Temperature and Humidity Requirements" on page 2-9 and "Ventilation and Cooling Requirements" on page 2-10 | Maximum: 14,052 BTU/hour (14,825.6 kJ/hour) Typical: 9,836 BTU/hour (10,377.5 kJ/hour) | Maximum: 26,411 BTU/hour (27,865.1 kJ/hour) Typical: 18,487 BTU/hour (19,504.8 kJ/hour) |
| Air flow front-to-back (subject to actual data center environment) See Also: "Temperature and Humidity Requirements" on page 2-9, and "Ventilation and Cooling Requirements" on page 2-10 | Maximum: 676 CFM Typical: 473 CFM | Maximum: 1223 CFM Typical: 856 CFM |

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

| Environmental Component | Starter Rack | Starter Rack with One In-Rack Expansion Kit |
|--|--|---|
| IP addresses See Also: Chapter 3, "Understanding the Network Requirements." | 12 for Ethernet network 6 for InfiniBand network | 18 for Ethernet network 12 for InfiniBand network |
| Network drops See Also: Chapter 3, "Understanding the Network Requirements." | One network connection | One network connection |
| External connectivity See Also: Chapter 3, "Understanding the Network Requirements." | 1 x 1 Gbps Ethernet port 6 x 10 Gbps Ethernet ports | 1 x 1 Gbps Ethernet port 12 x 10 Gbps Ethernet ports |

¹ Varies based on application load.

Space Requirements

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

- Height: 200 cm (79 inches)
- Width: 60 cm with side panels (24 inches)
- Depth: 120 cm (47.5 inches)

The minimum ceiling height for the cabinet is 230 cm (90 inches), measured from the true floor or raised floor, whichever is higher. An additional 92 cm (36 inches) is for top clearance. The space above the cabinet and its surroundings must not restrict the movement of cool air between the air conditioner and the cabinet, or the movement of hot air coming out of the top of the cabinet.

Receiving and Unpacking Requirements

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

- Shipping height: 216 cm (85 inches)
- Shipping width: 122 cm (48 inches)
- Shipping depth: 157.5 cm (62 inches)
- Shipping weight: 993 kg (2189 lb)

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

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

Use a conditioned space when removing the packaging material to reduce particles before entering the data center. Allow enough space for unpacking Oracle Big Data Appliance from its shipping cartons. [Table 2–3](#) lists the access route requirements.

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

| Access Route Item | With Shipping Pallet | Without Shipping Pallet |
|---|----------------------|-------------------------|
| Minimum door height | 218.4 cm (86 inches) | 204 cm (80.5 inches) |
| Minimum door width | 127 cm (50 inches) | 64 cm (25.5 inches) |
| Minimum elevator depth | 162.6 cm (64 inches) | 124 cm (49 inches) |
| Maximum incline | 6 degrees | 6 degrees |
| Minimum elevator, pallet jack, and floor loading capacity | 1134 kg (2500 lb) | 1134 kg (2500 lb) |

Maintenance Access Requirements

The maintenance area must be large enough for Oracle Big Data Appliance and have the required access space. For example, the required space to remove the side panels is 67.6 cm (26.6 inches). [Table 2–4](#) lists the maintenance access requirements.

Open tiles are required for electrical access.

Table 2–4 Maintenance Access Requirements for Oracle Big Data Appliance

| Location | Maintenance Access Requirement |
|-------------------|--------------------------------|
| Rear maintenance | 91.4 cm (36 inches) |
| Front maintenance | 91.4 cm (36 inches) |
| Top maintenance | 91.4 cm (36 inches) |

Flooring Requirements

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

[Table 2–5](#) lists the floor load requirements.

Table 2–5 Floor Load Requirements for Oracle Big Data Appliance

| Description | Requirement |
|---|-------------------|
| Maximum allowable weight of installed rack equipment | 998 kg (2200 lb) |
| Maximum allowable weight of installed power distribution units | 52 kg (115 lb) |
| Maximum dynamic load (maximum allowable weight of installed equipment including PDUs) | 1050 kg (2315 lb) |

Electrical Power Requirements

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

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

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

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

PDUs for Oracle Big Data Appliance

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

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

Refer to the appropriate sections for your location:

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

Low-Voltage 15 kVA Single-Phase PDUs

[Table 2–6](#) lists the requirements for low-voltage single-phase PDUs for North America, South America, Japan, and Taiwan. There are two PDUs per rack.

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

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

¹ UL489 2-pole circuit breaker

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

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

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

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



Low-Voltage 15 kVA Three-Phase PDUs

Table 2–7 lists the requirements for low-voltage three-phase PDUs for North America, South America, Japan, and Taiwan. There are two PDUs per rack.

Table 2–7 Low-Voltage 15 kVA Three-Phase PDUs

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

¹ UL489 2-pole circuit breaker

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

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

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

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

High-Voltage 15 kVA Single-Phase PDUs

Table 2–8 lists the requirements for 15 kVA high-voltage single-phase PDUs for Europe, the Middle East, Africa, and Asia Pacific, except Japan and Taiwan. There are two PDUs per rack.

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

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

¹ UL489 1-pole circuit breaker

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

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

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

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

High-Voltage 15 kVA Three-Phase PDUs

Table 2-9 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 2-9 High-Voltage 15 kVA Three-Phase PDUs

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

¹ UL489 1-pole circuit breaker

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

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

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

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



Facility Power Requirements

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

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

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

Circuit Breaker Requirements

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

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

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

Grounding Guidelines

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

Temperature and Humidity Requirements

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

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

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

[Table 2-10](#) lists the temperature, humidity, and altitude requirements for operating and nonoperating systems.

Table 2–10 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.

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

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

Network Connection Requirements

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

- Management network connection requirements
 - One 1 Gbps Ethernet connection for the management switch in the rack
 - One 1 Gbps Ethernet connection for the KVM switch in the rack (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: [Chapter 3, "Understanding the Network Requirements"](#)

Ensuring That the Site Is Ready

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

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

- [Task 4, "Prepare the Site Based on Requirements"](#)

Task 1 Review Site Requirements

Review the site requirements in this chapter and in the *Oracle Big Data Appliance Site Checklists* to understand the requirements of Oracle Big Data Appliance and ensure you are ready for delivery.

Task 2 Complete the Oracle Big Data Appliance Configuration Worksheets

Provide the network configuration details and select the optional software components to be installed. Complete the *Oracle Big Data Appliance Configuration Worksheets* for every set of appliances that will be configured as a single CDH cluster.

Give the completed *Oracle Big Data Appliance Configuration Worksheets* to your Oracle representative. An Oracle representative enters this information in the Oracle Big Data Appliance Configuration Utility and generates the Installation Template, the `preinstall-checkip.sh` script, and the configuration files.

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

Task 3 Configure the Network

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

See Also:

- [Chapter 3, "Understanding the Network Requirements"](#)
- [Chapter 7, "Configuring an Oracle Big Data Appliance Full Rack or Starter Rack"](#)

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"](#) on 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"](#) on page 2-3

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

See Also: ["Flooring Requirements"](#) on page 2-4

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

See Also:

- ["Electrical Power Requirements"](#) on page 2-4
- *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"](#) on page 2-10

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

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

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

See Also: ["Network Connection Requirements"](#) on page 2-12

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

Understanding the Network Requirements

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

- [Overview of Network Requirements](#)
- [Cabling the Client Network](#)
- [Factory Network Settings](#)
- [Port Assignments for Oracle Big Data Appliance Software](#)
- [Preparing the Network for Oracle Big Data Appliance](#)

Overview of Network Requirements

Oracle Big Data Appliance includes 6, 12, or 18 servers and the equipment to connect the servers to your network. The network connections allow the servers to be administered remotely and allow clients to connect to them. Use the information in this chapter to configure the environment for Oracle Big Data Appliance.

Each server has the following network components and interfaces:

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

Default Network Connections

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

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

Minimum Requirements for the Networks

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

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

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

There are two uplinks to the administrative network:

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

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

Network Diagram

The servers are configured on the network as follows:

- eth0: Provides access to the operating system using the administrative network.
- bondeth0: Provides access to the server using the client access network.
- ILOM: Provides access to Oracle Integrated Lights Out Manager (ILOM) using the administrative network.

Use the `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" on page 3-9.

Figure 3-1 shows the network diagram for a Sun Server X3-2L-based rack.

Figure 3-1 Network Diagram for Oracle Big Data Appliance X3-2

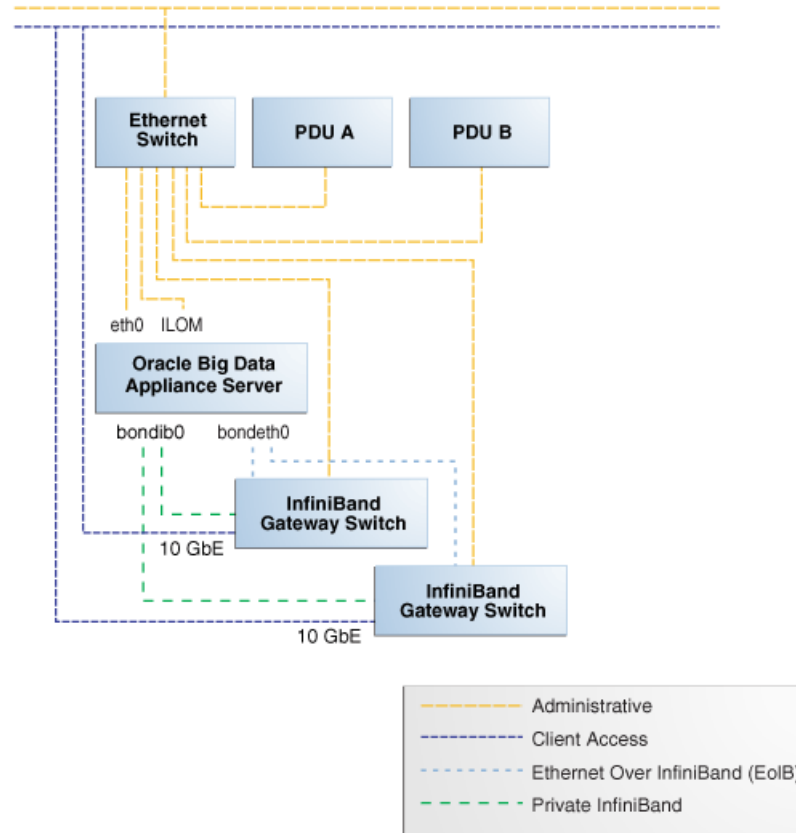
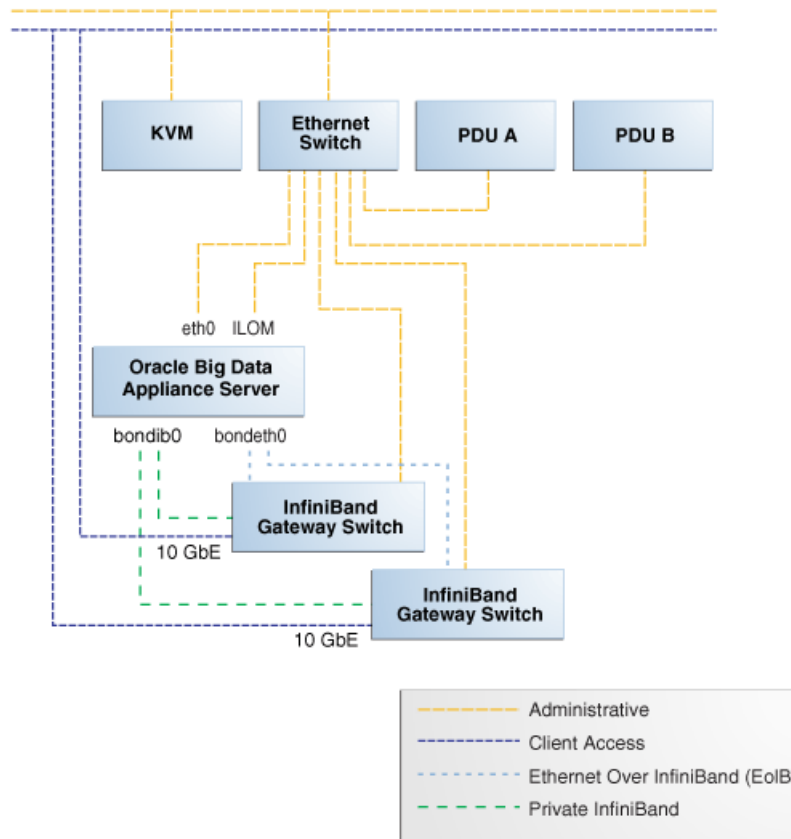


Figure 3-2 shows the network diagram for a Sun Fire X4270 M2 -based rack, which includes a KVM.

Figure 3–2 Network Diagram for Oracle Big Data Appliance

Cabling the Client Network

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

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

| Server Number | Active Link | Backup Link |
|---------------|--------------------|--------------------|
| 1 | GW Switch 1 Link 1 | GW Switch 2 Link 1 |
| 2 | GW Switch 1 Link 2 | GW Switch 2 Link 2 |
| 3 | GW Switch 1 Link 3 | GW Switch 2 Link 3 |
| 4 | GW Switch 1 Link 1 | GW Switch 2 Link 1 |
| ... | ... | ... |
| 17 | GW Switch 2 Link 2 | GW Switch 1 Link 2 |
| 18 | GW Switch 2 Link 3 | GW Switch 1 Link 3 |

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

Using Splitter Cables for Connecting to the 40 Gbps Physical Ports

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

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

What About Data Centers Without a 10 GbE Infrastructure?

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

Configuring Multiple Racks into a Single Hadoop Cluster

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

Factory Network Settings

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

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

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

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

| Host | Administrative IP Addresses | Oracle ILOM IP Addresses | InfiniBand Bonded IP Addresses |
|--------------------|-----------------------------|--------------------------|--------------------------------|
| bda18 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 | | |
| bdasw-ib2 | 192.168.1.202 | | |
| bda09 ² | 192.168.1.9 | 192.168.1.109 | 192.168.10.9 |
| bda08 ² | 192.168.1.8 | 192.168.1.108 | 192.168.10.8 |
| bda07 ² | 192.168.1.7 | 192.168.1.107 | 192.168.10.7 |
| bda06 | 192.168.1.6 | 192.168.1.106 | 192.168.10.6 |
| bda05 | 192.168.1.5 | 192.168.1.105 | 192.168.10.5 |
| bda04 | 192.168.1.4 | 192.168.1.104 | 192.168.10.4 |
| bda03 | 192.168.1.3 | 192.168.1.103 | 192.168.10.3 |
| bda02 | 192.168.1.2 | 192.168.1.102 | 192.168.10.2 |
| bda01 | 192.168.1.1 | 192.168.1.101 | 192.168.10.1 |
| bdasw-ib1 | 196.168.1.201 | | |

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

| Host | Administrative IP Addresses | Oracle ILOM IP Addresses | InfiniBand Bonded IP Addresses |
|-------|-----------------------------|--------------------------|--------------------------------|
| PDU A | 192.168.1.210 | | |
| PDU B | 192.168.1.211 | | |

¹ Full racks only² Full racks or starter racks with in-rack expansions only

Port Assignments for Oracle Big Data Appliance Software

[Table 3–2](#) identifies the port numbers used by Oracle Big Data Appliance software. Ensure that these ports are free before you configure the network.

Table 3–2 Oracle Big Data Appliance Port Numbers

| Port | Used by |
|--------------|---|
| 22 | ssh |
| 80 | yumrepos (only during installation) |
| 111 | portmap |
| 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 |
| 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) |

[Table 3–3](#) lists the ports used by Cloudera's Distribution including Apache Hadoop version 4 (CDH4). For additional details about these port assignments, go to the Cloudera website at

<http://ccp.cloudera.com/display/CDH4DOC/Configuring+Ports+for+CDH4>

Table 3–3 CDH Port Numbers

| Port | Component | Service | Access |
|------|-----------|-----------------|-----------|
| 0 | MapReduce | TaskTracker | Localhost |
| 1004 | HDFS | Secure DataNode | External |
| 1006 | HDFS | Secure DataNode | External |
| 2181 | HBase | HQuorumPeer | |
| 2181 | ZooKeeper | Server | External |

Table 3–3 (Cont.) CDH Port Numbers

| Port | Component | Service | Access |
|-------------|------------------|-----------------------------|---------------|
| 2888 | HBase | HQuorumPeer | |
| 2888 | ZooKeeper | Server | Internal |
| 3181 | ZooKeeper | Server | Internal |
| 3888 | HBase | HQuorumPeer | |
| 3888 | ZooKeeper | Server | Internal |
| 4181 | ZooKeeper | Server | Internal |
| 8002 | Hue | Beeswax Server | Internal |
| 8003 | Hue | Beeswax Metastore | Internal |
| 8020 | HDFS | NameNode | External |
| 8021 | MapReduce | JobTracker | External |
| 8030 | YARN | ResourceManager | |
| 8031 | YARN | ResourceManager | |
| 8032 | YARN | ResourceManager | |
| 8033 | YARN | ResourceManager | |
| 8040 | YARN | NodeManager | |
| 8041 | YARN | NodeManager | |
| 8042 | YARN | NodeManager | |
| 8080 | HBase | REST Service | External |
| 8088 | YARN | ResourceManager | External |
| 8888 | Hue | Server | External |
| 9083 | Hive | Metastore | External |
| 9090 | HBase | Thrift Server | External |
| 9090 | HBase | Avro Server | External |
| 9290 | MapReduce | JobTracker Thrift Plug-in | Internal |
| 10000 | Hive | HiveServer | External |
| 11000 | Oozie | Server | External |
| 11001 | Oozie | Server | Localhost |
| 16000 | Sqoop | Metastore | External |
| 19888 | YARN | MapReduce JobHistory Server | |
| 50010 | HDFS | DataNode | External |
| 50020 | HDFS | DataNode | External |
| 50030 | MapReduce | JobTracker | External |
| 50060 | MapReduce | TaskTracker | External |
| 50070 | HDFS | NameNode | External |
| 50075 | HDFS | DataNode | External |
| 50090 | HDFS | Secondary NameNode | Internal |
| 50470 | HDFS | Secure NameNode | External |

Table 3–3 (Cont.) CDH Port Numbers

| Port | Component | Service | Access |
|-------|-----------|---------------------------|----------|
| 50495 | HDFS | Secure Secondary NameNode | Internal |
| 60000 | HBase | Master | External |
| 60010 | HBase | Master | External |
| 60020 | HBase | RegionServer | External |
| 60030 | HBase | RegionServer | External |

Preparing the Network for Oracle Big Data Appliance

The Oracle Big Data Appliance Configuration Utility generates two files for you to use before installation:

- **bda-install-template.pdf:** This Installation Template lists the host names and custom IP addresses for Oracle Big Data Appliance. Verify that the Installation Template is correct, and add the host names and IP addresses to a name server on the network.
- **preinstall-checkip.sh:** Checks the network for name and address conflicts.

You can obtain the files from your Oracle representative, who typically runs the Oracle Big Data Appliance Configuration Utility. See "[About the Configuration Files](#)" on page 4-3.

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

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

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

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

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

[Example 3-1](#) shows success and error messages from the `preinstall-checkip.sh` script.

Example 3-1 Output from `preinstall-checkip.sh`

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

Using name server 198.51.100.21 for all DNS lookups

Domain is example.com

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

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

Using the Oracle Big Data Appliance Configuration Utility

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

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

Note: The network administrator must check the network configuration, using the output files from the configuration utility, before the installation of Oracle Big Data Appliance. Otherwise, network errors may cause extensive delays. See "[Preparing the Network for Oracle Big Data Appliance](#)" on page 3-9.

Overview of the Oracle Big Data Appliance Configuration Utility

The Oracle Big Data Appliance Configuration Utility is used to generate the installation and deployment files.

The utility is implemented as a spreadsheet containing these worksheets:

- Welcome
- Network Configuration
- Server Group Configuration (1 to 3 worksheets)
- Preview

An Oracle representative typically completes the Oracle Big Data Appliance Configuration Utility from information provided by the customer in the *Oracle Big Data Appliance Configuration Worksheets*. The customer and the Oracle representative work together to ensure that Oracle Big Data Appliance is configured correctly. For the sequence of steps involved in the configuration, see the *Oracle Big Data Appliance Configuration Worksheets*.

Using the Oracle Big Data Appliance Configuration Utility

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

- [Using OpenOffice Spreadsheet](#)
- [Generating the Configuration Files](#)
- [About the Configuration Files](#)

Using OpenOffice Spreadsheet

You must use Apache OpenOffice Spreadsheet 3.4 or Oracle Open Office Calc 3.3 to configure the spreadsheet. Do not use other spreadsheet programs like Microsoft Office Excel, which can modify the spreadsheet but not generate the files.

To obtain the Oracle Big Data Appliance Configuration Utility, refer to My Oracle Support Information Center ID 1445745.2.

To open the Oracle Big Data Appliance Configuration Utility:

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

AutoCorrect settings can make it difficult for you to enter passwords.
6. Save your changes and reload the spreadsheet.
7. If you see a security warning, select **Enable Macros**.

Generating the Configuration Files

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

To generate the configuration files:

1. Open the spreadsheet with OpenOffice as described in the previous procedure.
2. On the Welcome page, choose the type of installation, then the layout of the servers from the drop-down menus. The layout choices vary depending on your choice of installation type.
3. Click **Next** to edit the Network Configuration sheet.

4. Enter the network configuration settings as specified in the *Oracle Big Data Appliance Configuration Worksheets*. If the server IP addresses are nonsequential, then you can enter them on the Preview page.
5. Click **Next**. If your changes are invalid, then you see an error message. You must correct the errors before continuing to the Server Group 1 sheet.

To discard your changes, click **Reset**.

6. Enter the software configuration settings as specified in the *Oracle Big Data Appliance Configuration Worksheets*. If the group of servers is being added to an existing cluster, then all settings disappear from view.
7. Depending on your choices on the Welcome sheet, you must complete 1 to 3 Server Group sheets. Click **Next** to see the next Server Group sheet.
8. Click **Preview** to view the configuration settings defined by your entries. To make any changes, either click **Previous**, click the tab for a particular sheet, or enter the changes manually on the Preview page. You must click Preview on the Software Configuration sheet for any generated changes to appear on the Preview sheet.
9. When you are satisfied with the configuration, click **Save Configuration** on the Preview sheet to generate the configuration files.

This step does not save the worksheet. To save your changes, click the **Save** icon.
10. Choose a location for the configuration files in the Select Path dialog box. The files are created in a directory named `bda-customer_name/rack_name` in that location, such as `bda-Example Inc/bda1`.
11. If you are configuring a single rack or part of a rack, then you are done. For a multirack cluster, complete the Network Configuration in a separate set of worksheets for each extension that is outside the primary rack.
12. Send the Installation Template to the network administrator to verify that the configuration is correct.
13. Send the final Installation Template, the hosts file, and the `preinstall-checkip.sh` script to the network administrator to prepare the network for the installation of Oracle Big Data Appliance.

About the Configuration Files

The Oracle Big Data Appliance Configuration Utility generates the following files to use when you configure the system. They are stored in a directory named `bda-customer_name/rack_name`, such as `bda-Example Inc/bda1`. If you configure multiple clusters for the rack, then the utility generates a directory for each one.

BdaDeploy.json

Contains the network configuration for a full rack, a starter rack, or a starter rack with one in-rack expansion kit. It contains information about all the servers, switches, and PDUs. Transfer this file to a USB drive for copying to Oracle Big Data Appliance. See ["Configuring the Network"](#) on page 7-20.

BdaExpansion.json

Contains the network configuration for one or two in-rack expansion kits. It contains information about all the servers, but no information about the switches and PDUs. Transfer this file to a USB drive for copying to Oracle Big Data Appliance. See ["Configuring the Network"](#) on page 7-20.

mammoth-rack_name.params

Contains all information for a server group provided in the spreadsheet, including the network configuration, port numbers, default user names, and passwords. The spreadsheet creates a separate parameter file for each cluster. If several clusters are being configured, then each parameter file is located in a separate subdirectory named *config/cluster_name*. Transfer the files to a USB drive for copying to Oracle Big Data Appliance. See ["Installing the Software on a New Rack"](#) on page 9-2.

If an in-rack expansion kit is being configured as an addition to an existing cluster, then the spreadsheet does not generate a parameter file. The Mammoth utility generates it.

rack_name-bda-install-template.pdf

Reproduces the Installation Preview page of the spreadsheet in a printable format. The network administrator and others can use this Installation Template to verify the settings and make any last minute corrections.

rack_name-hosts

Lists the IP addresses and domains for all hosts on all networks:

- Servers on the public (client) network
- Servers on the administrative network
- Servers on the private (InfiniBand) network
- Oracle ILOMs on the administrative network
- Switches and PDUs on the administrative network

rack_name-preinstall-checkip.sh

Runs a series of tests to ensure the specified names and IP addresses for Oracle Big Data Appliance were added correctly to the name server, and they do not conflict with the existing network configuration. See ["Preparing the Network for Oracle Big Data Appliance"](#) on page 3-9 for instructions on running this script.

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

Figure 4–1 Installation Template

The screenshot shows the 'BDA Installation Preview' utility interface. At the top, there are navigation buttons '<< Previous' and 'Reset'. The main title is 'BDA Installation Preview' with the subtitle 'Big Data Appliance Configuration Utility'. Below this is a 'Network Configuration' section with a 'General' sub-section. The 'General' section contains two tables. The first table lists general properties: Cluster Name (bda), Rack Name (bda1), Region (America), Timezone (America/New_York), and Domain (example.com). The second table lists customer information: Customer Name (Example Inc) and Primary Appliance (Yes). Below these are fields for DNS Servers (172.16.100.45), NTP Servers (172.16.215.62), and Search Domains (example.com, us.example.com). The 'Server Network Info' section contains a table with columns 'Type', 'Netmask', and 'Gateway'. The rows are: Administrative – eth0 (255.255.255.0, 10.18.113.1), Private – bondib0 (255.255.255.0), and Client Access – bondeth0 (255.255.255.0, 10.18.114.1). A note on the right states 'BDA will be not be connected to a'.

Welcome Page

You must identify the type of equipment being delivered to the customer site and the planned configuration. [Table 4–1](#) describes these options.

Table 4–1 General Properties for the Installation

| Spreadsheet Field | Description |
|-------------------|--|
| Type of Install | Lists various configurations, including a full rack, a starter rack, and one or two in-rack expansion kits. Your selection in this field restricts the list in Layout of Servers to applicable choices only. |
| Layout of Servers | Lists the available options for configuring groups of servers into CDH clusters. The number of groups controls the number of Server Group pages in the spreadsheet: one, two, or three. |

Network Configuration

[Table 4–2](#) to [Table 4–6](#) briefly describe the fields in the Network Configuration sheet of the Oracle Big Data Appliance Configuration Utility. See the *Oracle Big Data Appliance Configuration Worksheets* for complete details.

[Table 4–2](#) describes the general rack properties.

Table 4–2 General Properties for the Rack

| Spreadsheet Field | Description |
|-------------------|---|
| Rack Name | The name of the Oracle Big Data Appliance rack. |

Table 4–2 (Cont.) General Properties for the Rack

| Spreadsheet Field | Description |
|--|---|
| Server Base Name | Base name for all servers. A two-digit suffix uniquely identifies each server. The rack name and server base name are used to generate the host names for all network interfaces: eth0, bondib0, bondeth0, and Oracle ILOM. For example, a cluster name of bda and a server base name of node results in host names of bda1node01, bda1node02, and so forth. |
| Admin Access Suffix | Suffix to the basic host name to form the bondeth0 host names |
| Private Name Suffix | Suffix to the basic host name to form the bondib0 host name |
| ILOM Name Suffix | Suffix to the basic host name to form the Oracle ILOM name |
| Switch Base Name | Suffix to the rack name to form the base name for all switches. For example, a rack name of bda and a switch base name of sw results in switch names of bda02sw-ip, bda02sw-ib1, and so forth. |
| Customer Name | Name of the enterprise. Required. |
| Region | Region in which the Oracle Big Data Appliance is located. Click the field, and then click the arrow for a list of regions. |
| Time Zone | Time zone for your installation. You must select a region before selecting the time zone. Click the field, and then click the arrow for a list of time zones for the region. |
| Domain Name | Name of the domain in which Oracle Big Data Appliance operates. Required. |
| Will this rack be connected via InfiniBand to other racks? | Choose Yes or No . |

[Table 4–3](#) describes the network properties for the entire rack. If the site has isolated client and administrative networks with different DNS servers, NTP servers, or search domains, then complete this section for the client network and [Table 4–7](#) for the administrative network.

Table 4–3 Network Properties for the Rack

| Spreadsheet Field | Description |
|-------------------|---|
| DNS Servers | Up to four IP addresses for the Domain Name System (DNS) servers. At least one DNS server must be accessible on the client network. DNS servers are not required on the administrative network, although Oracle Big Data Appliance uses them if they are available. |
| NTP Servers | Up to four IP addresses for the Network Time Protocol (NTP) servers. Both the client network and the administrative network must have access to at least one NTP server. The NTP server for the administrative network can be different from the NTP server for the client network. |
| Search Domains | Up to four domain names in which Oracle Big Data Appliance operates, such as example.com and us.example.com |

[Table 4–4](#) describes the network properties for individual servers in the rack. You specify the IP addresses assigned to node01. Sequential numbers are assigned automatically to the other servers. Ensure that an adequate range of addresses is free for each type of access. If sequential IP addresses are not available for the servers, then you can edit the Preview page.

Table 4–4 Network Properties for Servers

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

[Table 4–5](#) describes the network properties for all switches. You configure the switches manually, as described in [Chapter 7](#). If three sequential IP addresses are not available for the InfiniBand switches, then you can edit the Preview page.

Table 4–5 Network Properties for Switches

| Spreadsheet Field | Description |
|----------------------------|---|
| KVM Switch IP | IP address of the KVM switch. Not applicable to Oracle Big Data Appliance X3-2, which does not have a KVM switch. |
| Cisco Switch IP | IP address of the Cisco switch |
| First InfiniBand Switch IP | First IP address of three for the InfiniBand switches |

[Table 4–6](#) describes the network properties for the power distribution units (PDUs). You configure the PDUs manually, as described in [Chapter 7](#).

Table 4–6 Network Properties for Power Distribution Units

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

[Table 4–7](#) describes the network properties for the administrative network if they are different from the client network. Typically, these properties are the same for both networks, and you can leave these fields blank. They default to the values you entered for the general network properties. See [Table 4–3](#).

Table 4–7 Advanced Network Configuration Properties

| Spreadsheet Field | Description |
|--|--|
| Admin Domain | The administrative domain name, if it is different from the domain you identified in the general rack properties. |
| Are administrative host name entries in the DNS? | Choose Yes or No . |
| Admin DNS Servers | Up to four IP addresses for the administrative Domain Name System server, if the addresses are different from those you provided for the general network properties. |
| Admin NTP Servers | Up to four IP addresses for the administrative Network Time Protocol server, if the addresses are different from those you provided for the general network properties. |
| Admin Search Domains | Up to four domain names in which Oracle Big Data Appliance administrative network operates, if the domains are different from those you provided for the general network properties. |

Software Configuration

Table 4–8 to Table 4–12 briefly describe the fields in the Software Configuration sheet of the Oracle Big Data Appliance Configuration Utility.

Table 4–8 describes the optional software available for Oracle Big Data Appliance. If you have a license for Oracle Big Data Connectors, then you can install the software. For more information about these components, see the *Oracle Big Data Appliance Software User's Guide*.

Table 4–8 Installed Components

| Spreadsheet Field | Description |
|---|---|
| Are Big Data Connectors licensed? | Oracle Big Data Connectors facilitate data access between data stored in the CDH cluster and Oracle Database. The connectors require a separate license. If you have a license, choose Yes . |
| Install Oracle Data Integrator Agent? | The agent is a required component of Oracle Data Integrator, which is a graphical tool for performing extraction, transformation, and loading (ETL) of data into Oracle Database. To configure Oracle Data Integrator for use immediately, choose Yes . |
| Install Oracle NoSQL Database? | Oracle NoSQL Database is a distributed key-value database that provides a storage option to Hadoop Distributed File System (HDFS). To configure Oracle NoSQL Database for use, choose Community Edition or, if you have a license, Enterprise Edition . |
| Total disk space to allocate for Oracle NoSQL Database in Terabytes | If you install Oracle NoSQL Database, you must allocate space for its use. Disk space allocated to Oracle NoSQL Database is not available for HDFS. |

Table 4–9 describes the options for Auto Service Request (ASR). This service monitors the health of Oracle Big Data Appliance hardware. ASR Manager automatically submits a service request to Oracle Support Services when it detects a fault. See [Chapter 5](#) for more information about ASR.

Table 4–9 Oracle Auto Service Request

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

Table 4–10 identifies the users, groups, and passwords that can be set during the software installation. If you leave the passwords blank, then the installation prompts for them. The Oracle IDs must match those of a connected Oracle Exadata Database Machine to support the network file system (NFS) protocol.

Table 4–10 Users, Groups, and Passwords

| Spreadsheet Field | Description |
|---------------------------------|--|
| Cloudera Manager admin password | The password for the admin user for Cloudera Manager |

Table 4–10 (Cont.) Users, Groups, and Passwords

| Spreadsheet Field | Description |
|---|---|
| root operating system password | The root password on all servers in the cluster. |
| oracle operating system password | 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 user ID of a connected Oracle Exadata Database Machine. |
| oinstall group ID | The ID number of the Oracle Inventory Group (oinstall). It must match the oinstall group ID of a connected Oracle Exadata Database Machine. |
| dba group ID | The ID number of the dba group. It must match the dba group ID of a connected Oracle Exadata Database Machine. |
| MySQL administration password | The password for the MySQL Database administrative user |
| MySQL password for Oracle Data Integrator | The password for the Oracle Data Integrator user in MySQL Database |

Table 4–11 describes the configuration settings for the email server that Cloudera Manager uses to send alerts from the CDH cluster.

Table 4–11 Cloudera Manager Email Alerts

| Spreadsheet Field | Description |
|-----------------------------------|---|
| Email Server (SMTP) Host | The fully qualified name of the existing SMTP server that the company uses on its internal network |
| Email Server (SMTP) Port | The port number used by the email server |
| Email Server uses authentication? | Choose Yes or No . |
| Email Server (SMTP) User Name | User name for Cloudera Manager to use for authentication on the SMTP server. This field is hidden when authentication is set to No. |
| Email Server (SMTP) Password | Password for the user name. This field is hidden when authentication is set to No. |
| Email Server uses SSL? | Choose Yes or No . |
| Email Alert Recipients | Up to four email addresses. These users receive the alerts from Cloudera Manager. |

Table 4–12 describes the configuration settings for the Oracle Enterprise Manager system monitoring plugin for Oracle Big Data Appliance.

Table 4–12 Oracle Enterprise Manager Cloud Control

| Spreadsheet Field | Description |
|----------------------------|--|
| Enable Enterprise Manager? | To use the Oracle Enterprise Manager system monitoring plugin, choose Yes . |
| Cloud Control host | The fully qualified name of the server where Oracle Enterprise Manager is installed with the plugin for Oracle Big Data Appliance. |
| OMS port | The port number for the Oracle Enterprise Manager Cloud Control web interface. |

Table 4–12 (Cont.) Oracle Enterprise Manager Cloud Control

| Spreadsheet Field | Description |
|-------------------------------------|--|
| Cloud Control user name | A Cloud Control user with super-administrative privileges to perform administrative <code>emcli</code> commands. Typically, this user is <code>sysman</code> . |
| Cloud Control password | Password for the Cloud Control user name. |
| Cloud Control registration password | The Enterprise Manager agent registration password for validating the Oracle Management agents on Oracle Big Data Appliance. |

Setting Up Auto Service Request

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

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

Understanding Auto Service Request

Auto Service Request (ASR) is designed to automatically open service requests when specific Oracle Big Data Appliance hardware faults occur. 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>

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](#) 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://www.oracle.com/technetwork/systems/asr/overview/hardware-recommendations-330108.html>
6. Confirm that Java Development Kit 6 (JDK 1.6.0_04 or later) is running on the designated ASR Manager system:

```
java -version
```

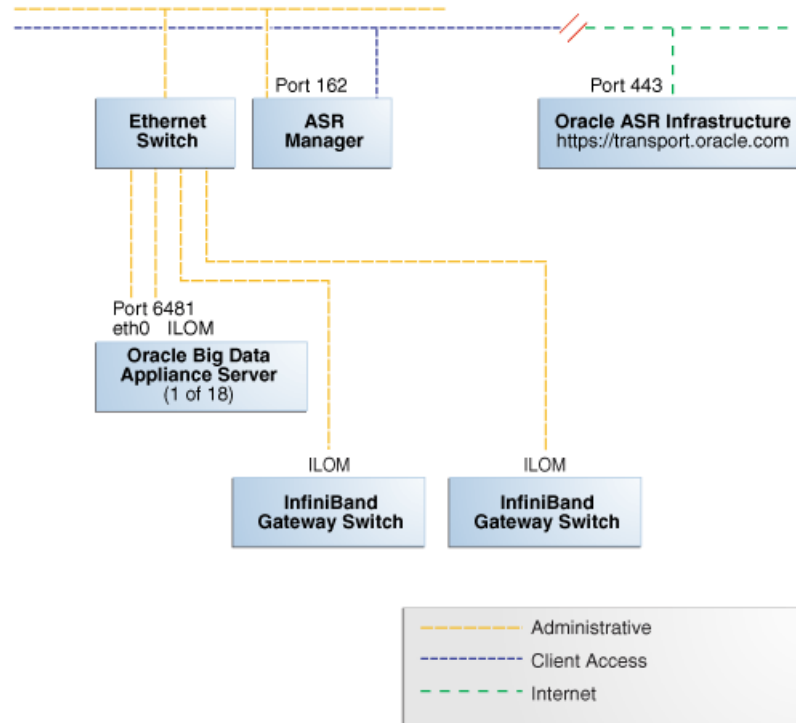
If necessary, download and install the latest version of JDK from the Java SE Downloads website:

<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 shows the network connections between ASR and Oracle Big Data Appliance.

Figure 5–1 Auto Service Request Network Connections



Installing ASR Manager

To install ASR Manager, download version 3.6 or higher from My Oracle Support ID 1185493.1. Then follow the instructions in the *Oracle Auto Service Request Installation and Operations Guide* at

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

Verifying the ASR Manager Installation

Perform these checks on ASR Manager to ensure that it is installed properly:

- As the `root` user, verify that the ASR installer created three CRON jobs:

```
crontab -l
```

The following are the recommended settings:

- `asr report`: Set once a month on Sunday
- `asr heartbeat`: Set twice daily or at least once daily
- `asr update_rules.sh`: Set once daily by default
- Verify that ASR Manager 3.5 or later is running:

```
asr show_rules_version
```

- Check the registration status:

```
asr show_reg_status
```

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

```
asr test_connection
```

Configuring the Trap Destinations on Oracle Big Data Appliance

ASR is an optional software component of Oracle Big Data Appliance. When planning the initial software configuration, you must complete the section about ASR in the *Oracle Big Data Appliance Configuration Worksheets*. Then during the software installation, Oracle Big Data Appliance Mammoth Utility configures the servers to trap the component fault indicators. See "[Mammoth Installation Steps](#)" on page 9-14.

You can reconfigure your appliance at any time to add or drop support for ASR. See "[Changing the Configuration of Optional Software](#)" on page 9-7.

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 | 1143FMM073 | Enabled | SUN FIRE X4270 M2 SERVER |
| 203.0.114.45 | bda1node16 | 1143FMM021 | Enabled | SUN FIRE X4270 M2 SERVER... |
| 203.0.115.140 | bda1node16-c | 1143FMM063 | Enabled | SUN FIRE X4270 M2 SERVER |
| 203.0.114.46 | bda1node17 | 1143FMM021 | Enabled | SUN FIRE X4270 M2 SERVER... |
| 203.0.115.141 | bda1node17-c | 1143FMM46B | Enabled | SUN FIRE X4270 M2 SERVER |

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/mammoth-rack_name.params | 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"](#) on page 9-7.

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

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. 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.

You received the emails from both of the previous tests (Oracle ILOM and operating system environment), then your validation is complete. If not, then proceed to "[Troubleshooting ASR](#)" on page 5-6.

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

Troubleshooting ASR

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

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

If you continue to have issues, contact ASR support. See My Oracle Support Note 1352349.1.

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:

- [Chapter 6, "Installing Oracle Big Data Appliance at the Site"](#)
- [Chapter 7, "Configuring an Oracle Big Data Appliance Full Rack or Starter Rack"](#)
- [Chapter 8, "Connecting Multiple Oracle Big Data Appliance Racks"](#)

Installing Oracle Big Data Appliance at the Site

This chapter describes how to move, install, and configure the hardware—from unpacking Oracle Big Data Appliance to powering on the system. This chapter contains these sections:

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

Reviewing Safety Guidelines

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

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

- Do not place anything on top of Oracle Big Data Appliance or perform any work directly above it.
- Do not let the room temperature rise sharply, especially in winter. Sudden temperature changes can cause condensation to form inside Oracle Big Data Appliance. Allow for a sufficient warm-up period 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/>

Unpacking Oracle Big Data Appliance

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

[Figure 6–1](#) 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>

Contents of the Shipping Kit

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

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

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

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

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

Removing Oracle Big Data Appliance from the Shipping Crate

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

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

To unpack Oracle Big Data Appliance:

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

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

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

5. Save the mounting brackets that secure the rack to the shipping pallet. You can use these mounting brackets to permanently secure Oracle Big Data Appliance to the installation site floor.

Note: Do not dispose of these brackets, because you cannot order replacement brackets.

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

Placing Oracle Big Data Appliance in Its Allocated Space

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

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

Moving Oracle Big Data Appliance


To move Oracle Big Data Appliance:

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

When moving Oracle Big Data Appliance to the installation site, the front casters do not turn; you must steer the unit by moving the rear casters, as shown in [Figure 6-2](#). You can safely maneuver Oracle Big Data Appliance by carefully pushing it.

[Figure 6-2](#) 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.

WARNINGS:

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

Stabilizing Oracle Big Data Appliance

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


Stabilize the Oracle Big Data Appliance Rack with Leveling Feet

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

To adjust the leveling feet:

1. Locate the four leveling feet located at the bottom corners of the rack.
2. Lower the leveling feet to the floor as shown in [Figure 6-3](#) using the SW 12 mm wrench. When lowered correctly, the four leveling feet should support the full weight of the rack.

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




Stabilize the Oracle Big Data Appliance Rack with Mounting Brackets

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

To install the mounting brackets:


1. Drill the mounting holes in the installation floor. [Figure 6-4](#) shows the location and dimensions of the mounting brackets.

Figure 6-4 Location of Mounting Brackets on the Rack



2. Obtain four bolts and washers to mount the Oracle Big Data Appliance rack to the floor. The bolt holes in the mounting brackets have a 10.0 mm diameter. Oracle does not provide mounting bolts because different floors require different bolt types and strengths. Select bolts that are appropriate for your location.
3. Position the rack over the predrilled holes. [Figure 6-5](#) shows the bottom view of the Oracle Big Data Appliance rack and the location for the mounting hole and floor cut dimensions.

Figure 6-5 Bottom View of the Oracle Big Data Appliance Rack




- 1: Distance from the center of the mounting bracket to the edge of the rack is 113 mm (4.45 inches).
- 2: Width between the centers of the mounting hole slots is 374 mm (14.72 inches).
- 3: Distance from the center of the mounting bracket to the edge of the rack is 113 mm (4.45 inches).
- 4: Distance between the centers of the front and rear mounting hole slots is 1120 mm (44.1 inches).
- 5: Depth of cable-routing floor cutout is 330 mm (13 inches).
- 6: Distance between the floor cutout and the edge of the rack is 160 mm (6.3 inches).
- 7: Width of cable-routing floor cutout is 280 mm (11 inches).

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


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

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

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



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

Figure 6-7 Securing the Oracle Big Data Appliance Rack Using Mounting Brackets

6. Using bolts and washers obtained in Step 2, permanently mount your system to the floor using the four mounting brackets shown in [Figure 6-7](#).
7. Firmly tighten all of the bolts that secure the mounting brackets to the rack and to the floor.

Attaching a Ground Cable (Optional)


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

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

To attach a ground cable:

1. Obtain a grounding cable. It is not shipped with the system.
2. Ensure that the installation site has properly grounded the power source in the data center. The facility PDU must have earth ground.
3. Ensure that all grounding points, such as raised floors and power receptacles, reference the facility ground.
4. Ensure that direct, metal-to-metal contact is made for this installation. During manufacturing, the ground cable attachment area might have been painted or coated.
5. Attach the ground cable to an attachment point located at the bottom rear of the system frame, as shown in [Figure 6-8](#). The attachment point is an adjustable bolt that is inside the rear of the cabinet.

Figure 6–8 Earth Ground Attachment Bolt Location



Powering On the System for the First Time

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

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

Inspecting the Oracle Big Data Appliance Rack After It Is in Place

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

To inspect the Oracle Big Data Appliance rack:

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

See Also: ["Ventilation and Cooling Requirements"](#) on page 2-10 for more information

Connecting Power Cords

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

To connect power cords to the rack:

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

Figure 6-9 Location of Power Cord Cable Ties




Figure 6–10 Power Cord Routing Out the Bottom





Figure 6–11 Power Cord Routing Out the Top




Powering On Oracle Big Data Appliance

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

To power 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 6–12](#). Press the ON (I) side of the toggle switch. When the breaker is ON, the switch is flush with the side of the PDU.

Figure 6–12 PDU Switch Locations

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

See Also: "[Reviewing Safety Guidelines](#)" on page 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.

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

Note: After powering on Oracle Big Data Appliance, configure the system as described in [Chapter 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](#)
- [Configuring the Cisco Ethernet Switch](#)
- [Configuring the InfiniBand Leaf and Spine Switches](#)
- [Configuring the Power Distribution Units](#)
- [Configuring the Oracle Big Data Appliance Servers](#)
- [Configuring the Network](#)

See Also:

- [Chapter 12, "Using the dcli Utility"](#)
- [Chapter 13, "Using Oracle Integrated Lights Out Manager"](#)

Note:

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

Configuring the KVM Switch

A KVM switch is provided only in Oracle Big Data Appliance racks with Sun Fire X4270 M2 servers. Oracle Big Data Appliance X3-2 does not have a KVM switch.

The KVM configuration consists of these procedures:

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

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

Starting the KVM Switch

To start the KVM switch:

1. Ensure that all connected components are powered off.
2. Pull the KVM tray out from the front of the rack, and open it using the handle.
3. Touch the touch pad.
4. Toggle between the host and KVM interface by pressing the Ctrl key on the left side twice, similar to double-clicking with a mouse. You see the main user interface page.
5. In the navigator on the left, select **Target Devices** under Unit View. In the main display area, verify that 18 target devices are listed with Action set to KVM Session.

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

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

Connecting the KVM Switch to the Management Network

To connect the KVM switch to the management network:

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

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

Checking the KVM Firmware Version

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

To check the KVM firmware version:

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

- Application: 1.10.2.17762
- Boot: 1.9.16473

If the application firmware version is earlier than 1.10.2, then you should upgrade it. To upgrade the firmware, continue with this procedure. Otherwise, you are done.

2. Download the firmware from this website to a USB flash drive:
<http://www.avocent.com/Pages/GenericTwoColumn.aspx?id=12541>
3. Plug the flash drive into the KVM USB port and open a browser session.
4. Log in to the KVM as Admin with password welcome1.
5. Under Appliance, select **Overview**.
6. From the Tools list, select **Upgrade Firmware**.
7. Select the connection method, such as **FTP** or **HTTP**.
8. Enter the file name of the downloaded firmware.
9. Click **Upgrade**.
The upgrade process takes 5 to 10 minutes, including an automatic restart.
10. Confirm the firmware version by selecting **Versions** under Appliance Settings.

Configuring the KVM Switch to Access the Servers

To configure the KVM switch to access the servers:

1. Under Unit View, select **Target Devices** to display the Target Devices page.
2. 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 through 5 for each server in the rack. Each server boots through BIOS, and boots the operating system with the default factory IP configuration.

Accessing a Server by Using the KVM Switch

To access a server by using the KVM switch:

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

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

Configuring the Cisco Ethernet Switch

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

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

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

Scope of the Configuration

This configuration disables IP routing and sets the following:

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

Prerequisites for Configuring the Ethernet Switch

To avoid disrupting the customer network, observe these prerequisites:

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

See Also:

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

Configuring the Ethernet Switch on the Customer Network

To configure the Ethernet switch on the customer network:

1. Connect a serial cable from the Cisco switch console to a laptop or similar device. An RJ45 to DB9 serial cable is included with the Cisco documentation package.

The status lights turn on and this message is displayed:

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

2. 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.

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

```
Switch> enable
Password:
```

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

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface vlan 1
Switch(config-if)# ip address 10.7.7.34 255.255.255.0
Switch(config-if)# end
Switch# *Jan 23 15:54:00.506: %SYS-5-CONFIG_I:Configured from console by console
Switch# write memory
Building configuration...
Compressed configuration from 2474 bytes to 1066 bytes [OK ]
```

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

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

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

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

```
Switch(config)# end
*Jan 23 15:55:02.506: %SYS-5-CONFIG_I:Configured from console by console
Switch# write memory
Building configuration...
Compressed configuration from 2502 bytes to 1085 bytes [OK ]
```

7. Set the host name of the switch provided in the *Oracle Big Data Appliance Configuration Worksheets*. This example uses the name `bdal-sw-ip`.

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

The system host name appears in the prompt.

8. Configure up to three DNS servers, using the IP addresses provided in the *Oracle Big Data Appliance Configuration Worksheets* 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.

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

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.

```
bdal-sw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdal-sw-ip(config)# enable password welcome1
bdal-sw-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.
bdal-sw-ip(config)# end
bdal-sw-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.
bdalsw-ip(config)# line vty 0 15
bdalsw-ip(config-line)# login
%Login disabled on line 1,until 'password' is set
%Login disabled on line 2,until 'password' is set
%Login disabled on line 3,until 'password' is set
...
bdalsw-ip(config-line)# password welcome1
bdalsw-ip(config-line)# login
bdalsw-ip(config-line)# end
bdalsw-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.
bdalsw-ip(config)# line vty 0 15
bdalsw-ip(config-line)# no password
bdalsw-ip(config-line)# login
%Login disabled on line 1, until 'password' is set
%Login disabled on line 2, until 'password' is set
%Login disabled on line 3, until 'password' is set
...
bdalsw-ip(config-line)# end
bdalsw-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]
```

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:

```
bdalsw-ip# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
bdalsw-ip(config)# crypto key generate rsa
% You already have RSA keys defined named bdalsw-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]
bdalsw-ip(config)# username admin password 0 welcome1
bdalsw-ip(config)# line vty 0 15
bdalsw-ip(config-line)# transport input ssh
bdalsw-ip(config-line)# exit
bdalsw-ip(config)# aaa new-model
```

```

bdal-sw-ip(config)# ip ssh time-out 60
bdal-sw-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
bdal-sw-ip(config)# ip ssh version 2
bdal-sw-ip(config)# end
*Sep 15 14:26:37.045: %SYS-5-CONFIG_I: Configured from console by console
bdal-sw-ip# write memory
Building configuration...
Compressed configuration from 2603 bytes to 1158 bytes[OK]

```

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:

```

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

```

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

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.

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

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.

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 prior to connecting to the customer network.

If spanning tree needs to 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:

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

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

2. To verify the disabling of the spanning tree:

```
bdal-sw-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
```

Verifying the Ethernet Configuration

To verify the Cisco Ethernet switch configuration:

1. Verify the configuration by entering the following command:

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

The following is an example of the output:

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

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

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

2. Save the current configuration by entering this command:

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

3. Exit from the session with this command:

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

Press RETURN to get started.

4. Disconnect the cable from the Cisco console.
5. To check the configuration, attach a laptop to port 48 and ping the IP address of the internal management network.

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

Configuring the InfiniBand Leaf and Spine Switches

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

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

Configuring an InfiniBand Switch

To configure an InfiniBand switch:

1. Connect to the switch using a serial or an Ethernet connection.
 - For an Ethernet connection to 192.168.1.201, 192.168.1.202, or 192.168.1.203, see ["Connecting to Oracle ILOM Using the Network"](#) on page 13-3.
 - For a serial connection, see ["Connecting to Oracle ILOM Using a Serial Connection"](#) on page 13-4.

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

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

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

```
cd /SP/network
```

4. Enter these commands to configure the switch:

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

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

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

```
-> show
```

```

/SP/network
  Targets:
    interconnect
    ipv6
    test

  Properties:
    commitpending = (Cannot show property)
    dhcp_ser_ip = none
    ipaddress = 10.135.42.24
    ipdiscovery = static
    ipgateway = 10.135.40.1
    ipnetmask = 255.255.255.0
    macaddress = 00:21:28:E7:B3:34
    managementport = SYS/SP/NET0
    outofbandmacaddress = 00:21:28:E7:B3:33
    pendingipaddress = 10.135.42.23
    pendingipdiscovery = static
    pendingipgateway = 10.135.42.1
    pendingipnetmask = 255.255.248.0
    pendingmanagementport = /SYS/SP/NET0
    sidebandmacaddress = 00:21:28:E7:B3:35
    state = enabled

  Commands:
    cd
    set
    show

->

```

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

```

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

```

```

/SP
  Properties:
    hostname = bdal-sw-ib2

```

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

```

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

```

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

8. Verify the settings:

```

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

  Properties:
    auto_dns = enabled
    nameserver = 10.196.23.245, 172.32.202.15
    retries = 1

```

```
searchpath = us.example.com
timeout = 5
```

```
Commands:
  cd
  set
  show
```

Setting the Time Zone and Clock on an InfiniBand Switch

To set the time zone on an InfiniBand switch:

1. Check the current time setting:

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

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

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

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

3. Check the current time setting:

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

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

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

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

5. Check the current time setting:

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

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

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

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

7. Enable the NTP servers:

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

Note: Properly synchronized clocks are required for the Mammoth Utility 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
```

Checking the Health of an InfiniBand Switch

To check the health of an InfiniBand leaf or spine switch:

1. Open the Fabric Management shell:

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

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

2. Check the firmware version, which should be 2.0.5-2 or later. Check My Oracle Support ID 1474875.1 for the current version.

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

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

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

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

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


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

```
FabMan@bdal-sw-ib2-> setpriority list
Current SM settings:
priority 5
controlled_handover TRUE
subnet_prefix 0xfe80000000000000
```

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

6. To correct the priority setting:

- a. Stop the InfiniBand Subnet Manager:

```
FabMan@bdal-sw-ib2-> disableism
```

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

```
FabMan@bdal-sw-ib2-> setpriority 5
```

- c. Restart the InfiniBand Subnet Manager:

```
FabMan@bdal-sw-ib2-> enableism
```

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

- a. Verify that the Exadata InfiniBand switches and the Exalogic spine switch are running firmware version 1.3.3_2 or later.
- b. Ensure that the subnet manager runs only on the switches with the highest firmware version.
- c. On systems running earlier firmware versions, disable the subnet manager. Log in to the switch as `root` and run the `disableism` command as described previously.

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

8. Exit the Fabric Management shell:

```
FabMan@bdal-sw-ib2-> exit
->
```

9. Exit the Oracle ILOM shell:

```
-> exit
```

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

```
reboot
```

11. Repeat these steps for the other InfiniBand switches.

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 Exadata Database Machine racks, then verify that the instance numbers are even and unique to each switch in the fabric.

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@bdalsw-ib2->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@bdalsw-ib2->setgwinstance 10
Stopping Bridge Manager... [ OK ]
Starting Bridge Manager. [ OK ]
FabMan@bdalsw-ib2->setgwinstance --list
BXM system name set to 10
```

5. Repeat Steps 1 to 4 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 Exadata Database Machine racks), then repeat Steps 1 to 5 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.

Configuring the Power Distribution Units

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

- [Connecting the PDUs to the Network](#)
- [Updating the PDU Firmware Version](#)

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 `admin`. Change this password after configuring the network.
 3. Click **Net Configuration** in the upper left of the page.
 4. Confirm that the DHCP Enabled option is not selected.
 5. Enter the following network settings for the PDU metering unit:
 - IP address
 - Subnet mask
 - Default gateway
 6. Click **Submit** to set the network settings and reset the PDU metering unit.
 7. Repeat Steps 5 and 6 for the second PDU.

Updating the PDU Firmware Version

To verify the PDU firmware version:

1. Select **Module Info**. If the output displays a firmware version of 1.06 or later, then you are done. Otherwise, continue this procedure to update the firmware version.
2. Download the latest firmware version from My Oracle Support:
 - a. Log in at <http://support.oracle.com>.

- b. Select the **Patches & Updates** tab.
 - c. For Patch Search, click **Product or Family (Advanced)**.
 - d. For Product, select Sun Rack II PDU.
 - e. For Release, select Sun Rack II PDU 1.0.6.
 - f. Click **Search** to see the Patch Search Results page.
 - g. Click the patch name, such as **16523441**.
 - h. Download the file.
3. 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.
 4. In the PDU web interface, select **Net Configuration** and log in as admin with password admin.
 5. Scroll down to Firmware Update.
 6. Click **Browse**, select the MKAPP_V1.06.DL file, and click **Submit**.
 7. Click **Browse**, select the HTML_V1.06.DL file, and click **Submit**.
 8. Click **Module Info** to verify the firmware version number.
 9. Click **Net Configuration**, and then click **Logout**.

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 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](#)" on page 12-1.

For starter racks, with or without an in-rack expansion kit, you must use the -j parameter. See [setup-root-ssh](#) on page 11-27.

```
# setup-root-ssh -j "etho0_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 again to verify the keys.

7. 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
```

8. 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)
```

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

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

The output should appear as follows for each server:

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

10. Power up each server and verify that the following files are in the `/root` directory:

```
BDA_IMAGING_SUCCEEDED
BDA_REBOOT_SUCCEEDED
```

11. 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
```

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

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

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

This command returns all lines that do not show `SUCCESS`, and 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
```
- 13.** Save the RAID configuration in a file, replacing *filename* with a file name of your choice:
- ```
dcli MegaCli64 -ldinfo -lall -a0 | grep "Virtual Drive\|State" > filename
```
- 14.** 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
```
- 15.** Save the software profile output from each system into a file for review, replacing *filename* with a file name of your choice:
- ```
dcli bdacheksw > filename
```
- 16.** 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
```
- 17.** Verify the system boots in this order: USB, RAID Slot 0, PXE:
- ```
dcli "ubiosconfig export all -x /tmp/bios.xml"
[# dcli "grep -m1 -A20 boot_order /tmp/bios.xml" | more
```
- 18.** 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.

## Configuring the Network

The Oracle Big Data Appliance Configuration Utility generates the BdaDeploy.json file, which is used to configure the administrative network and the private InfiniBand

network. See ["Generating the Configuration Files"](#) on page 4-2 if you do not have this file.

The network configuration consists of these procedures:

- [Copying the Configuration Files to Oracle Big Data Appliance](#)
- [Starting the Network Configuration](#)
- [Connecting to the Administrative Network](#)
- [Completing the Network Configuration](#)
- [Checking the Virtual Interface Cards \(VNICs\)](#)
- [Verifying Access to the Network Services](#)

## 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 D-1](#).
3. Log in as the `root` user on the first server. The initial password is `welcome1`.
4. Plug the USB drive into the USB port of the first server. The port is on the right front of the server. Information like the following is displayed on the console:

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

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

```
showusb
/dev/sdn1
```

6. Create a directory on the server:

```
mkdir /mnt/usb
```

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

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

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

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

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

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

---

---

**Note:** If mammoth-rack\_name.params is also on the drive, you can copy it to /opt/oracle/BDAMammoth for use in [Chapter 9](#).

---

---

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.

## 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 one or two in-rack expansion kits 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 BdaDeploy.json file. For one or two in-rack expansion kits, verify that /opt/oracle/bda contains both BdaDeploy.json from an existing node and the new BdaExpansion.json file.
3. Verify that the InfiniBand network is healthy:

```
cd /opt/oracle/bda
bdacheckib BdaShip.json

LINK bdasw-ib1.0B ... bdasw-ib3.8B UP
LINK bdasw-ib1.1B ... bdasw-ib2.8B UP
LINK bdasw-ib3.15A ... bda02.HCA-1.2 UP
LINK bdasw-ib3.15B ... bda01.HCA-1.2 UP
LINK bdasw-ib3.14A ... bda04.HCA-1.2 UP
LINK bdasw-ib3.14B ... bda03.HCA-1.2 UP
LINK bdasw-ib3.13A ... bda06.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](#)" on page 11-25.
5. Reconnect to the server using its ILOM instead of SSH:

```
ssh root@192.168.1.101
```



```

Password: welcome1
-> start -f /HOST/console
bda01 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](#) 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/BdaDeploy.json
networksetup-one: passed
networksetup-one: ping servers on ship admin network
networksetup-one: passed
networksetup-one: test ssh to servers on ship admin network
hello from node02
hello from node03
.
.
.
networksetup-one: passed
networksetup-one: copy /opt/oracle/bda/BdaDeploy.json to servers
BdaDeploy.json 0% 0 0.0KB/s --:-- ETABdaDeploy.json 100% 4304 4.2KB/s 00:00
BdaDeploy.json 0% 0 0.0KB/s --:-- ETABdaDeploy.json 100% 4304 4.2KB/s 00:00
.
.
.
networksetup-one: passed
networksetup-one: executing network settings on all servers
networksetup-one: wait a few seconds for the network to restart on 192.168.1.2
.
.
.
bda1node02.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
```

## 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:

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

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

## 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@bdal-sw-ib2-> 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)

```

- Repeat Steps 3 to 5 on the second gateway switch (-ib3).

## 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:

- 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"](#) on page 7-24.

---



---

- 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](#) 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/BdaDeploy.json
networksetup-two: passed
networksetup-two: ping server ips on admin network
networksetup-two: passed
networksetup-two: test ssh server ips on admin network
hello from bda1node02.example.com
hello from bda1node03.example.com
hello from bda1node04.example.com
.
.
.
networksetup-two: passed
networksetup-two: run connected network post script on each server
networksetup-two: post network setup for 10.133.42.253
networksetup-two: post network setup for 10.133.42.254
networksetup-two: post network setup for 10.133.43.1
.
.
.

```

```

networksetup-two: post network setup for this node
networksetup-two: ping admin servers by name on admin network
networksetup-two: passed
networksetup-two: verify infiniband topology
networksetup-two: passed
networksetup-two: start setup client network (10gigE over Infiniband)
networksetup-two: ping both gtw leaf switches
networksetup-two: passed
networksetup-two: verify existence of gateway ports
networksetup-two: passed
networksetup-two: ping server ips on admin network
networksetup-two: passed
networksetup-two: ping servers by name on admin network
networksetup-two: passed
networksetup-two: test ssh server ips on admin network
hello from bdalnode02.example.com
hello from bdalnode03.example.com
.
.
.
networksetup-two: passed
networksetup-two: check existence of default vlan for port 0A-ETH-1 on bdalsw-ib2
networksetup-two: no default vlan for port, create it
spawn ssh root@10.133.43.36 createvlan 0A-ETH-1 -vlan -1 -pkey default
networksetup-two: verify default vlan for port 0A-ETH-1 for bdalsw-ib2
.
.
.
networksetup-two: passed
networksetup-two: apply eoib on each server
networksetup-two: wait a few seconds for the network to restart on 10.133.42.253
networksetup-two: wait a few seconds for the network to restart on 10.133.42.254
.
.
.
check and delete vNIC for bdalnode02 eth9 on switch bdalsw-ib2
check and delete vNIC for bdalnode02 eth9 on switch bdalsw-ib3
create vNIC eth9 bdalnode02 using switch bdalsw-ib3
vNIC created
check and delete vNIC for bdalnode02 eth8 on switch bdalsw-ib2
.
.
.
networksetup-two: ping server ips on client network
networksetup-two: passed
networksetup-two: test ssh server ips on client network
hello from bdalnode02.example.com
hello from bdalnode03.example.com
.
.
.
networksetup-two: passed
networksetup-two: end setup client network

```

## 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@bdalsw-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.

## 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](#)" on page 11-27.
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
```

4. Verify network connectivity:

```
bdachecknet | tee -a /tmp/bdachecknet_exp.out
bdachecknet: check syntax and static semantics of
/opt/oracle/bda/BdaDeploy.json
bdachecknet: passed
bdachecknet: ping test private infiniband ips (bondib0 40gbs)
```

```
bdachecknet: passed
bdachecknet: ping test admin ips (eth0 1gbs)
bdachecknet: passed
bdachecknet: ping test client access ips (bondeth0 10gbs Eoib)
bdachecknet: passed
bdachecknet: test admin network resolve and reverse resolve
bdachecknet: passed
bdachecknet: test admin name array matches ip array
bdachecknet: passed
bdachecknet: test client network (eoib) resolve and reverse resolve
bdachecknet: passed
bdachecknet: test client name array matches ip array
bdachecknet: passed
bdachecknet: test ntp servers
bdachecknet: passed
bdachecknet: test arp -a
bdachecknet: passed
```

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.

## 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.

### 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](#)
  - [Configuring the Cisco Ethernet Switch](#)
  - [Configuring the InfiniBand Leaf and Spine Switches](#)
  - [Configuring the Power Distribution Units](#)
  - [Configuring the Oracle Big Data Appliance Servers](#)
  - [Configuring the Network](#)
2. Cable the Oracle Big Data Appliance racks together. See [Chapter 8](#).
3. Revalidate the InfiniBand network. See "[Validating a Multirack InfiniBand Network](#)" on page 7-30.

#### 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](#)
  - [Configuring the Cisco Ethernet Switch](#)
  - [Configuring the InfiniBand Leaf and Spine Switches](#)
  - [Configuring the Power Distribution Units](#)
  - [Configuring the Oracle Big Data Appliance Servers](#)
  - [Configuring the Network](#)
2. Unplug the power to the InfiniBand switches on the new racks.
  3. Cable the Oracle Big Data Appliance racks together. See [Chapter 8](#).
  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](#)" on page 7-30.

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

1. Complete the configuration steps for each Oracle Big Data Appliance rack individually:
  - [Configuring the KVM Switch](#)
  - [Configuring the Cisco Ethernet Switch](#)
  - [Configuring the InfiniBand Leaf and Spine Switches](#)
  - [Configuring the Power Distribution Units](#)
  - [Configuring the Oracle Big Data Appliance Servers](#)
  - [Configuring the Network](#)
2. Cable the Oracle Big Data Appliance racks together. See [Chapter 8](#).
3. Revalidate the InfiniBand network. See "[Validating a Multirack InfiniBand Network](#)" on page 7-30.

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

1. Complete the configuration steps for each new Oracle Big Data Appliance rack individually:
  - [Configuring the KVM Switch](#)
  - [Configuring the Cisco Ethernet Switch](#)
  - [Configuring the InfiniBand Leaf and Spine Switches](#)
  - [Configuring the Power Distribution Units](#)
  - [Configuring the Oracle Big Data Appliance Servers](#)
  - [Configuring the Network](#)
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 racks together. See [Chapter 8](#).
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](#)" on page 7-30.

## 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.

---



---

### 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](#) 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":
"bda02sw-ib3"}
{"SPINE_NAME": "exa1sw-ib1", "LEAF1_NAME": "exa1sw-ib2", "LEAF2_NAME":
"exa01sw-ib3"}]

```



---

---

# 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](#)
- [Cabling Two Racks Together](#)
- [Cabling Several Racks Together](#)

## Extending a Rack by Adding Another Rack

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

During the cabling procedure, note the following:

- There is some performance degradation while you are cabling the racks together. This degradation results from reduced network bandwidth, and the data retransmission due to packet loss when a cable is unplugged.
- The environment is not a high-availability environment because one leaf switch must be off. All traffic goes through the remaining leaf switch.
- Only the existing rack is operational, and any new rack is powered down.
- The software running on the systems must not have problems related to InfiniBand restarts.
- The new racks must be configured with the appropriate IP addresses to be migrated into the expanded system 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:**

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

## Cabling Two Racks Together

The following procedure describes how to cable two racks together. This procedure assumes that the racks are adjacent to each other. In the procedure, the existing rack is R1, and the new rack is R2.

**To cable two racks together:**

1. Set the priority of the current, active Subnet Manager Master to 10 on the spine switch, as follows:
  - a. Log in to any InfiniBand switch on the active system.
  - b. Use the `getmaster` command to verify that the Subnet Manager Master is running on the spine switch.
  - c. Log in to the spine switch.
  - d. Use the `disablesm` command to stop the Subnet Manager.
  - e. Use the `setsmpriority 10` command to set the priority to 10.
  - f. Use the `enablesm` command to restart the Subnet Manager.
  - g. Repeat Step b to ensure that the Subnet Manager Master is running on the spine switch.
2. Ensure that the new rack is near the existing rack. The InfiniBand cables must be able to reach the servers in each rack.
3. Completely shut down the new rack (R2).
4. Cable the leaf switch in the new rack according to [Table E-2](#).
5. Power off leaf switch R1 IB2. This causes all servers to fail over their InfiniBand traffic to R1 IB3.
6. Disconnect all interswitch links between R1 IB2 and R1 IB3.
7. Cable leaf switch R1 IB2 according to [Table E-1](#).
8. Power on leaf switch R1 IB2.
9. Wait for 3 minutes for R1 IB2 to become completely operational.

To check the switch, log in to it and run the `ibswitches` command. The output should show three switches: R1 IB1, R1 IB2, and R1 IB3.
10. Power off leaf switch R1 IB3. This causes all servers to fail over their InfiniBand traffic to R1 IB2.
11. Cable leaf switch R1 IB3 according to [Table E-1](#).
12. Power on leaf switch R1 IB3.
13. Wait for 3 minutes for R1 IB3 to become completely operational.

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

14. Power on all the InfiniBand switches in R2.
  15. Wait for 3 minutes for the switches to become completely operational.  
To check the switch, log in to it and run the `ibswitches` command. The output should show six switches: R1 IB1, R1 IB2, R1 IB3, R2 IB1, R2 IB2, and R2 IB3.
  16. Ensure that the Subnet Manager Master is running on R1 IB1 by running the `getmaster` command from any switch.
  17. Power on all servers in R2.
  18. Log in to spine switch R1 IB1, and lower its priority to 8 as follows:
    - a. Use the `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.

## 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.

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

- a. Power off leaf switch Rx IB2. This causes all servers to fail over their InfiniBand traffic to Rx IB3.
  - b. Cable leaf switch Rx IB2 according to [Appendix E](#).
  - 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,... Rn.
  - e. Power off leaf switch Rx IB3. This causes all servers to fail over their InfiniBand traffic to Rx IB2.
  - f. Cable leaf switch Rx IB3 according to [Appendix E](#).
  - 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,... Rn.  
All racks should now be rewired according to [Appendix E](#).
6. Power on all the InfiniBand switches in Rn+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 `disableesm` command to stop the Subnet Manager.
    - b. Enter the `setsmpriority 8` command to set the priority to 8.
    - c. Enter the `enableesm` 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 `enableesm` 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 `disableesm` 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:

- [Chapter 9, "Installing the Oracle Big Data Appliance Software"](#)





---

---

# Installing the Oracle Big Data Appliance Software

This chapter explains how to install, reinstall, or reconfigure the software on Oracle Big Data Appliance. This chapter contains these sections:

- [Installing the Software on a New Rack](#)
- [Adding Servers to a Cluster](#)
- [What If an Error Occurs During the Installation?](#)
- [Upgrading the Software on Oracle Big Data Appliance](#)
- [Changing the Configuration of Optional Software](#)
- [Reinstalling the Base Image](#)
- [Mammoth Utility Syntax](#)
- [Mammoth Reconfiguration Utility Syntax](#)
- [Oracle Big Data Appliance Base Imaging Utilities](#)

---

---

**Note:** Because Mammoth does not store passwords, you are prompted to enter each one. 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 reinstalling Oracle Big Data Connectors, then you also need the MySQL password for Oracle Data Integrator.

---

---

## About the Mammoth Utility

Mammoth is a command-line utility for installing and configuring the Oracle Big Data Appliance software. Using Mammoth, you can:

- Create a CDH cluster on one or more racks.
- Create multiple CDH clusters on a Oracle Big Data Appliance rack.
- Extend a CDH cluster on to new servers.
- Update a cluster with new software.
- Deploy Auto Service Request during or after the initial software installation.
- Deploy the Oracle Enterprise Manager system monitoring plugin for Oracle Big Data Appliance during or after the initial software installation.

## Installing the Software on a New Rack

The Mammoth Utility installs and configures the software on Oracle Big Data Appliance using the files generated by the Oracle Big Data Appliance Configuration Utility. At a minimum, Mammoth installs and configures Cloudera's Distribution including Apache Hadoop. This includes all the Hadoop software and Cloudera Manager, which is the tool for administering your Hadoop cluster. Mammoth optionally installs and configures Oracle NoSQL Database and, if you have a license, all components of Oracle Big Data Connectors.

In addition to installing the software across all servers in the rack, the Mammoth Utility creates the required user accounts, starts the correct services, and sets the appropriate configuration parameters. When it is done, you have a fully functional, highly tuned, up and running Hadoop cluster.

Complete the appropriate instructions for your installation:

- For an Oracle Big Data Appliance rack that forms one or more Hadoop clusters, follow the procedure in ["Installing the Software on a Single or Primary Rack"](#) on page 9-2.
- For multiple racks where each rack forms an independent Hadoop cluster, follow the procedure in ["Installing the Software on a Single or Primary Rack"](#) on page 9-2 for each rack.
- For multiple racks that form a single, multirack Hadoop cluster:
  - Identify the primary rack of the cluster, then follow the procedure in ["Installing the Software on a Single or Primary Rack"](#) on page 9-2.
  - For the other racks of the cluster, follow the procedure in ["Adding Servers to a Cluster"](#) on page 9-3.

### Installing the Software on a Single or Primary Rack

Follow this procedure to install and configure the software on a single Oracle Big Data Appliance rack or on the primary rack of a multiple-rack cluster.

**To install the software:**

1. Verify that the Oracle Big Data Appliance rack is configured according to the custom network settings described in `/opt/oracle/bda/BdaDeploy.json`. If the rack is still configured to the factory default IP addresses, first perform the network configuration steps described in ["Configuring the Network"](#) on page 7-20.
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 9.
3. Download the BDAMammoth zip file to any directory on node01 (such as `/tmp`). For the download location, see My Oracle Support Information Center ID 1445745.2.

4. Log in as `root` and extract all files from the downloaded zip file:

```
unzip p16694238_211_Linux-x86-64.zip
Archive: p16694238_211_Linux-x86-64.zip
 inflating: README.txt
 creating: BDAMammoth-2.1.1/
 inflating: BDAMammoth-2.1.1/bda-configurator-2.1.1.ods
 inflating: BDAMammoth-2.1.1/BDAMammoth-2.1.1.run
```

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

- ```
# cd BDAMammoth-2.1.1
```
6. Extract all files from BDAMammoth-*version*.run:


```
# ./BDAMammoth-2.1.1.run
```
 7. Change to the BDAMammoth directory.


```
# cd /opt/oracle/BDAMammoth
```
 8. Copy mammoth-*rack_name*.params to the current directory. See ["About the Configuration Files"](#) on page 4-3.
 9. Run the mammoth command with the appropriate options. See ["Mammoth Utility Syntax"](#) on page 9-13. This sample command runs all steps:


```
./mammoth -i rack_name
```
 10. To configure another Hadoop 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 4 to 9. Each cluster has its own mammoth-*rack_name*.params file. The file names are identical, but the Oracle Big Data Appliance Configuration Utility creates the files in separate directories named for the cluster.

The Mammoth Utility stores the current configuration in the `/opt/oracle/bda/install/state` directory. *Do not delete the files in this directory.* The Mammoth Utility fails without this information if you must use it again, such as adding a rack to the cluster.

Adding Servers to a Cluster

You can add servers to an existing cluster in groups of three servers, as defined in the *Oracle Big Data Appliance Configuration Worksheets*. You add a full rack the same way that you add a smaller number of servers. The spreadsheet does not generate the configuration file. Instead, you use Mammoth to generate the configuration file and then use it to configure the servers.

To install the software on additional servers in a cluster:

1. Ensure that all servers are running the same software version. See ["About Software Version Differences"](#) on page 9-4.
2. Ensure that all racks that form a single Hadoop cluster are cabled together. See [Chapter 8](#).
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 6
```

See the `-e` option in ["Mammoth Utility Syntax"](#) on page 9-13.

5. If you are prompted for an expansion number, enter an integer from 1 to 5.

6. Install the software. The following example begins installing software on the servers in bda3 with an expansion number of 1:

```
./mammoth -i bda3 1
```

If you have a license for Oracle Big Data Connectors, 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. You cannot add nodes to an Oracle NoSQL Database cluster after it is set up. However, a logical volume is created on the additional rack for future use when nodes can be added to an Oracle NoSQL Database cluster.

The Mammoth Utility obtains the current configuration from the files stored in `/opt/oracle/bda/install/state`. If those files are missing or if any of the services have been moved manually to run on other nodes, then the Mammoth Utility fails.

About Software Version Differences

All racks configured as one Hadoop cluster must have the same image. A new Oracle Big Data Appliance rack may be factory-installed with a newer base image than the previously installed racks. Use the `imageinfo` utility on any server to get the image version. When all racks of a single Hadoop cluster have the same image version, you can install the software on the new rack.

To synchronize the new rack with the rest of the Hadoop cluster, either upgrade the existing cluster to the latest image version or downgrade the image version of the new rack.

To upgrade the image version:

- Run the Mammoth Utility with the `-p` option to upgrade the software on the cluster to the latest version. See ["Upgrading the Software on Oracle Big Data Appliance"](#) on page 9-5.

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 older version of the Oracle Big Data Appliance Configuration Utility, which is on the first server of the existing cluster, to extend the cluster onto the new rack.

What If an Error Occurs During the Installation?

Each step generates a detailed log file listing the actions performed on each server and whether the step completed successfully. If an error occurs, the script stops. You can then check the log files in `/opt/oracle/BDAMammoth/bdaconfig/tmp`. The log files are named in this format:

```
STEP-i-yyyyymmddhhmmss.log
```

In this format, *i* is the step number and *yyyyymmddhhmmss* identifies the time that the file was created, as the year, month, day, hour, minute, and second.

After fixing the problem, you can rerun all steps or a range of steps. You cannot skip steps or run them out of order.

Upgrading the Software on Oracle Big Data Appliance

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

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

Software downgrades are not supported.

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

About the Changes from Version 1.1 to Version 2.1

Following is a summary of the changes an upgrade makes to the software on a cluster:

- CDH4 replaces CDH3, and Cloudera Manager 4.5 replaces Cloudera Manager 3.7.
- Cloudera Manager 4.5 manages Hive, so you use Cloudera Manager to start, stop, and configure Hive.
- Some services are located on different nodes:
 - The Cloudera Manager service runs on node03 instead of node02.
 - The primary and backup MySQL databases switch positions. The primary database runs on node03, and the backup database runs on node02.
 - The directory that is mounted on all nodes of a cluster is exported from node03 instead of node04.
- Namenode high availability and automatic failover are turned on:
 - There are two NameNodes, which run on node01 and node02. There is no secondary NameNode.
 - There are two failover controllers, which run on node01 and node02.
 - There are three journal nodes, which run on node01, node02, and node03.
 - The ZooKeeper service is configured with servers running on node01, node02, and node03.
 - The Oozie service is configured with a server running on node03.
- A mounted directory is not used to backup NameNode or secondary NameNode data.

See Also: *Oracle Big Data Appliance Software User's Guide* for more detailed information about the version 2.0 software.

Upgrading from Software Version 2.0 or 1.1

Follow these procedures to upgrade the software on an Oracle Big Data Appliance cluster from version 2.0 or 1.1 to version 2.1.

Note: This procedure upgrades version 1.1 software from CDH3 to CDH 4. You do not need to upgrade your mapreduce programs, but you do need to recompile them after this upgrade.

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))

To upgrade the software:

1. Download the BDAMammoth zip file to any directory (such as /tmp) on 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.

For the download location, see My Oracle Support Information Center ID 1445745.2. You upgrade to the version of the Mammoth Utility, and so you need Mammoth 2.1.1 to upgrade your appliance to 2.1.1 software.

2. Log in to the first node of the cluster as root.
3. Extract all files from the downloaded zip file:

```
# unzip p16694238_211_Linux-x86-64.zip
Archive:  p16694238_211_Linux-x86-64.zip
  inflating: README.txt
    creating: BDAMammoth-2.1.1/
  inflating: BDAMammoth-2.1.1/bda-configurator-2.1.1.ods
  inflating: BDAMammoth-2.1.1/BDAMammoth-2.1.1.run
```

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

```
# cd BDAMammoth-2.1.1
```

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

```
# ./BDAMammoth-2.1.1.run
```

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

6. Change to the BDAMammoth directory.

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

7. Run the mammoth command with the -p option:

```
./mammoth -p
```

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 Mammoth Reconfiguration Utility, you can reverse some of those decisions. You must provide the relevant server names, ports, user names, and passwords. See the `mammoth-reconfig` [add](#) and [remove](#) options.

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 [Chapter 5](#).
2. Log into the first NameNode (node01) of the primary rack and change to the BDAMammoth directory:

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

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

```
# cd /opt/oracle/BDAMammoth
# ./mammoth-reconfig add 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
INFO: Checking configuration file
/opt/oracle/bda/install/state/mammoth-saved.params...
INFO: Using saved configuration file
/opt/oracle/bda/install/state/mammoth-saved.params
INFO: To use the generated configuration file, remove
/opt/oracle/bda/install/state/mammoth-saved.params
INFO: Loading configuration file
/opt/oracle/bda/install/state/mammoth-saved.params...
INFO: Reading component versions from
/opt/oracle/BDAMammoth/bdaconfig/COMPONENTS
INFO: Creating nodelist files...
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 ...
.
.
.
```

The next procedure shows how to add support for Oracle Enterprise Manager Cloud Control:

To support Oracle Enterprise Manager Cloud Control:

1. Install the system monitoring plugin for Oracle Big Data Appliance in an Oracle Enterprise Manager Cloud Control installation on the same network. See the *Oracle Enterprise Manager System Monitoring Plug-in Installation Guide for Oracle Big Data Appliance*.
2. Log into the first NameNode (node01) of the primary rack and change to the BDAMammoth directory:

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

3. Add support for Oracle Enterprise Manager Cloud Control:

```
# cd /opt/oracle/BDAMammoth
# ./mammoth-reconfig add em
INFO: Logging all actions in
/opt/oracle/BDAMammoth/bdaconfig/tmp/bda1node01-20130205082218.log and traces
in /opt/oracle/BDAMammoth/bdaconfig/tmp/bda1node01-20130205082218.trc
INFO: Checking configuration file
/opt/oracle/bda/install/state/mammoth-saved.params...
INFO: Using saved configuration file
/opt/oracle/bda/install/state/mammoth-saved.params
INFO: To use the generated configuration file, remove
/opt/oracle/bda/install/state/mammoth-saved.params
INFO: Loading configuration file
/opt/oracle/bda/install/state/mammoth-saved.params...
INFO: Reading component versions from
/opt/oracle/BDAMammoth/bdaconfig/COMPONENTS
INFO: Creating nodelist files...
Enter the value for EM_HOST [Default:]: oem-host.example.com
Enter the value for EM_PORT [Default: 4473]:
Enter the value for EM_USER [Default: sysman]:

Please Check the values you entered for the EM parameters

EM_HOST = oem-host.example.com
EM_PORT = 4473
EM_USER = sysman

Are these values correct (y/n): y
Enter password for user sysman on machine oem-host.example.com
Enter password: password
Enter password again: password
Enter agent registration password for em setup on machine oem-host.example.com
Enter password: password
Enter password again: password
INFO: Checking passwordless ssh setup
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.
INFO: Creating environment.pp file ...
INFO: Making sure all puppet agents can be accessed.
INFO: Pinging puppet agents
INFO: Creating directories for em download This will take some time ...
.
.
```


Reinstalling the Base Image

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

You can reimage all of 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](#)
- [Reimaging an Oracle Big Data Appliance Rack](#)
- [Reimaging an Oracle Big Data Appliance Cluster](#)

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 latest Oracle Big Data Appliance Base Image tarball from My Oracle Support, and copy it to the server to be re-imaged.
2. If you are reimaging the server to the current customer settings, verify that `/opt/oracle/bda/BdaDeploy.json` reflects the intended network configuration. If it does not, then generate a new file using the Oracle Big Data Appliance Configuration Utility. See [Chapter 4](#).
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 base image zip file:

```
# unzip p16685583_210_Linux-x86-64.zip
Archive:  p16065021_201_Linux-x86-64.zip
  inflating: README.txt
   creating: BDABaseImage-2.1.0_RELEASE/
  inflating: BDABaseImage-2.1.0_RELEASE/reimagerack
  inflating: BDABaseImage-2.1.0_RELEASE/BDABaseImage-2.1.0_RELEASE.iso
  inflating: BDABaseImage-2.1.0_RELEASE/reimagecluster
```

```
extracting: BDABaseImage-2.1.0_RELEASE/BDABaseImage-2.1.0_RELEASE.md5sum
inflating: BDABaseImage-2.1.0_RELEASE/makebdaimage
inflating: BDABaseImage-2.1.0_RELEASE/ubiosconfig
inflating: BDABaseImage-2.1.0_RELEASE/biosconfig
```

5. Change to the BDABaseImage-2.1.0_RELEASE directory created in the previous step:

```
$ cd BDABaseImage-2.1.0_RELEASE
```

6. Reimage the server using the `makedbaimage` command. The following example reimages server 4, including the internal USB, from the 2.1.0 base image to the custom settings in `BdaDeploy.json`

```
./makebdaimage -usbint BDABaseImage-2.1.0_RELEASE.iso
/opt/oracle/bda/BdaDeploy.json 4
```

See [makebdaimage](#) on page 9-19 for the complete syntax of the command.

7. If the `makedbaimage` command succeeded without errors, reboot the server.

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.

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

1. If the Oracle Big Data Appliance software was installed previously on the rack, then save the `/opt/oracle/BDAMammoth/mammoth-rack_name.params` file to a safe place outside Oracle Big Data Appliance.
2. Download the most recent base image from My Oracle Support and copy it to the first (bottom) server of the rack being reimaged. For the download location, see My Oracle Support Information Center ID 1445745.2. You can copy the file to any directory, such as `/tmp`.
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/BdaDeploy.json` reflects the intended network configuration. If it does not, then generate a new file using the Oracle Big Data Appliance Configuration Utility. See [Chapter 4](#).
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"](#) on page 12-1. 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:

```
# 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% /
.
.
.
```

9. Extract all files from the base image zip file:

```
# unzip p16685583_210_Linux-x86-64.zip
Archive:  p16685583_210_Linux-x86-64.zip
  inflating: README.txt
   creating: BDABaseImage-2.1.0_RELEASE/
  inflating: BDABaseImage-2.1.0_RELEASE/reimagerack
  inflating: BDABaseImage-2.1.0_RELEASE/BDABaseImage-2.1.0_RELEASE.iso
  inflating: BDABaseImage-2.1.0_RELEASE/reimagecluster
  extracting: BDABaseImage-2.1.0_RELEASE/BDABaseImage-2.1.0_RELEASE.md5sum
  inflating: BDABaseImage-2.1.0_RELEASE/makebdaimage
  inflating: BDABaseImage-2.1.0_RELEASE/ubiosconfig
  inflating: BDABaseImage-2.1.0_RELEASE/biosconfig
```

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

```
# cd BDABaseImage-2.1.0_RELEASE
```

11. 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/BdaDeploy.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/BdaDeploy.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](#) on page 9-20 for the complete syntax of the command.

12. Run the Mammoth Utility. See "[Installing the Software on a Single or Primary Rack](#)" on page 9-2.

Reimaging an Oracle Big Data Appliance Cluster

Follow this procedure to reimage a group of servers that the Mammoth utility has deployed as a cluster. The existing network settings are automatically reapplied after reimaging.

Caution: If you reimage a cluster, then all files and data on the cluster are erased.

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

1. Save the `/opt/oracle/BDAMammoth/mammoth-rack_name.params` file to a safe place outside Oracle Big Data Appliance.
2. Download the most recent base image from My Oracle Support and copy it to the first server of the cluster being reimaged. You can copy the file to any directory, such as `/tmp`. For the download location, see My Oracle Support Information Center ID 1445745.2.

Depending on the configuration of clusters in the rack, the first server might be 1, 7, 10, or 13, counting from the bottom

3. Establish an SSH connection to the first server and log in as `root`.
4. Verify that `/opt/oracle/bda/BdaDeploy.json` reflects the intended network configuration. If it does not, then generate a new file using the Oracle Big Data Appliance Configuration Utility. See [Chapter 4](#).
5. Ensure that passwordless SSH is set up:

```
# dcli -C 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 in the cluster. If not, then follow the steps in "[Setting Up Passwordless SSH](#)" on page 12-1. Do not continue until the `dcli -C hostname` command runs successfully on all servers.

6. Check all Oracle Big Data Appliance servers for hardware issues:


```
# dcli -C 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:

```
# dcli -C 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% /
.
.
.
```

9. Extract all files from the base image zip file:

```
# unzip p16685583_210_Linux-x86-64.zip
Archive:  p16685583_210_Linux-x86-64.zip
  inflating: README.txt
   creating: BDABaseImage-2.1.0_RELEASE/
  inflating: BDABaseImage-2.1.0_RELEASE/reimagerack
  inflating: BDABaseImage-2.1.0_RELEASE/BDABaseImage-2.1.0_RELEASE.iso
  inflating: BDABaseImage-2.1.0_RELEASE/reimagecluster
 extracting: BDABaseImage-2.1.0_RELEASE/BDABaseImage-2.1.0_RELEASE.md5sum
  inflating: BDABaseImage-2.1.0_RELEASE/makebdaimage
  inflating: BDABaseImage-2.1.0_RELEASE/ubiosconfig
  inflating: BDABaseImage-2.1.0_RELEASE/biosconfig
```

10. Change to the `BDABaseImage-version_RELEASE` directory created in the previous step:

```
# cd BDABaseImage-2.1.0_RELEASE
```

11. Reimage the cluster:

```
./reimagecluster
```

See [reimagecluster](#) on page 9-20 for the complete syntax of the command.

12. Run the Mammoth Utility. See "[Installing the Software on a Single or Primary Rack](#)" on page 9-2.

Mammoth Utility Syntax

You must log in as `root` on the first server and change to the `/opt/oracle/BDAMammoth` directory to use the Mammoth Utility. It has this syntax:

```
./mammoth option [rack_name] [extension_number]
```

In this command, *rack_name* is the name of an Oracle Big Data Appliance rack. You must enter the rack name in the first command exactly as it appears in the configuration file name (`mammoth-rack_name[-extnumber].params`). Afterward, *rack_name* defaults to the rack specified in a previous `mammoth` command.

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

Example 9–1 Mammoth Utility Syntax Examples

This command displays Help for the Mammoth Utility:

```
./mammoth -h
```

This command does a complete install on rack `bda3`:

```
./mammoth -i bda3
```

This command generates a parameter file to add six servers in an in-rack expansion kit, beginning with `node07`, to an existing cluster:

```
./mammoth -e node07 6
```

The next command runs steps 2 through 6 on the rack being set up:

```
./mammoth -r 2-6
```

Mammoth Utility Options

The syntax of the `mammoth` command supports the configuration of new racks and in-rack expansion kits. You can also use the Mammoth 2.1.1 bundle to upgrade from 2.0 or 1.1.

-c

Run the Oracle Big Data Appliance cluster checks.

-e *start_node number_of_nodes*

Generates a parameter file for a group of servers being added to a cluster. The file is named `mammoth-rack_name.params` if the new servers are in a rack outside the cluster. Otherwise, Mammoth prompts for an in-rack expansion number from 1 to 5 that it uses in the name, such as `mammoth-bda1-1`.

start_node is the client network host name for the first server in the group being added.

number_of_nodes is the number of servers being added to the cluster. The number must be a multiple of three.

No passwords are included in the parameter file, so you must enter them when running Mammoth.

-h

Displays command Help including command usage and a list of steps.

-i [*rack_name*] [*expansion_number*]

Runs all mandatory steps, 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.

rack_name identifies the location of new servers being added to a cluster when they are outside the primary rack.

expansion_number identifies the `mammoth-rack-name.params` file when an expansion number is included in the name. See the `-e` option.

-l

List the steps of the Mammoth Utility.

-p

Upgrades the software on the cluster to the current version.

-r *n-N*

Run steps *n* through *N* of the Mammoth Utility while no errors occur

-s *n* [*rack_name*] [*expansion_number*]

Runs step *n* on the primary rack. Enter the *rack_name* to identify another rack in the same cluster, and an *expansion_number* if needed. See the `-e` option.

-v

Displays the version number of the Mammoth Utility.

Mammoth Installation Steps

Following are descriptions of the steps that the Mammoth Utility and the Mammoth Reconfiguration Utility perform when installing the software.

Step 1 SetupClusterAccess

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 Utility is run as the puppet master node. For example, if you run the Mammoth Utility from `bda1node01` with an IP address `192.168.41.1`, then a list of aliases for that IP address includes `bda1node01-master`. The Mammoth Utility uses Puppet for the software installation.
- Checks the network timing on all nodes. If the timing checks fail, then there are unresolved names and IP addresses that will prevent the installation from running correctly. Fix these issues before continuing with the installation.

Step 2 PreinstallChecks

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

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

Step 3 SetupPuppet

This step configures puppet agents on all nodes and start them, configures a puppet master on the node where the Mammoth Utility is being run, waits for the agents to submit their certificates, and automates their signing. After this step is completed, Puppet can deploy the software.

Puppet is a distributed configuration management tool that is commonly used for managing Hadoop clusters. The puppet master is a parent service and maintains a Puppet repository. A puppet agent operates on each Hadoop node.

A file named `/etc/puppet/puppet.conf` resides on every server and identifies the location of the puppet master.

Puppet operates in two modes:

- Periodic pull mode in which the puppet agents periodically contact the puppet master and asks for an update, or
- Kick mode in which the puppet master alerts the puppet agents that a configuration update is available, and the agents then ask for the update. Puppet operates in kick mode during the Mammoth Utility installation.

In both modes, the puppet master must trust the agent. To establish this trust, the agent sends a certificate to the puppet master node where the `sys admin` process signs it. When this transaction is complete, the puppet master sends the new configuration to the agent.

For subsequent steps, you can check the Puppet log files on each server, as described in ["What If an Error Occurs During the Installation?"](#) on page 9-4.

Step 4 PatchFactoryImage

Installs the most recent Oracle Big Data Appliance image and system parameter settings.

Step 5 CopyLicenseFiles

Copies third-party licenses to `/opt/oss/src/OSSLicenses.pdf` on every server, as required by the licensing agreements.

Step 6 CopySoftwareSource

Copies third-party software source code to `/opt/oss/src/` on every server, as required by the licensing agreements.

Step 7 CreateLogicalVolumes

Creates a logical volume if physical disks are allocated to Oracle NoSQL Database. This step varies depending on the amount of disk space allocated to Oracle NoSQL Database during configuration. Each cluster size (6, 9, or 12 servers) offers a choice in the number of terabytes allocated to Oracle NoSQL Database. For example, a full rack of 18 servers offers a choice of 54 or 108 terabytes.

- **0 terabytes:** This step does nothing.
- **Smaller number of terabytes:** The disk space is allocated across the cluster using one disk on each node. The disk mounted at `/u12` is used for the logical volume.
- **Larger number of terabytes:** The disk space is allocated across the cluster using two disks on each node. The disks mounted at `/u11` and `/u12` are used for the logical volume.

The logical volume is mounted at `/lv1` and corresponds to device `/dev/lvg1/lv1`.

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

Step 8 CreateUsers

Creates the `hdfs` and `mapred` users, and the `hadoop` group. It also creates the `oracle` user and the `dba` and `oinstall` groups.

The various packages installed in later steps also create users and groups during their installation.

See Also: *Oracle Big Data Appliance Software User's Guide* for more information about users and groups.

Step 9 SetupMountPoints

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 10 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`.

Step 11 InstallHadoop

Installs all packages in Cloudera's Distribution including Apache Hadoop (CDH) and Cloudera Manager. It then starts the Cloudera Manager server on node03 and configures the cluster.

Step 12 StartHadoopServices

Starts the agents on all nodes and starts all CDH services. After this step, you have a fully functional Hadoop installation.

Cloudera Manager runs on port 7180 of 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, which is changed in [Step 16](#).

Step 13 InstallBDASoftware

Installs Oracle NoSQL Database Community Edition and the server-side components of Oracle Big Data Connectors, if these options were selected in the Oracle Big Data Appliance Configuration Worksheets. Oracle NoSQL Database must be allocated disk space, and Oracle Big Data Connectors must be licensed separately.

Step 14 SetupEMAgent

Installs and configures the Oracle Enterprise Manager agent.

Step 15 SetupASR

Installs and configures Auto Service Request (ASR).

Note: For this step to run successfully, the ASR host system must be up with ASR Manager running and configured properly. See [Chapter 5](#).

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](#)" on page 5-4.

Step 16 CleanupInstall

Performs the following:

- Changes the root password on all nodes (optional).
- Changes the Cloudera Manager password if specified in the Installation Template.
- Deletes temporary files created during the installation.
- Copies log files from all nodes to subdirectories in /opt/oracle/bda/install/log.
- Runs cluster verification checks, including TeraSort, to ensure that everything is working properly. It also generates an install summary. All logs are stored in a subdirectory under /opt/oracle/bda/install/log on node01.

Step 17 CleanupSSHroot (Optional)

Removes passwordless SSH for `root` that was set up in [Step 1](#).

Mammoth Reconfiguration Utility Syntax

You must log in as `root` on the first server and change to the `/opt/oracle/BDAMammoth` directory to use the Mammoth Reconfiguration Utility. It has this syntax:

```
./mammoth-reconfig option parameter
```

Note:

- Where *parameter* is a node name in the syntax examples, `bda1` is the rack name, `node` is the server base name, and `-adm` is the administrative access suffix.
 - This utility uses the configuration settings stored in `/opt/oracle/bda/install/state/mammoth-saved.params`. When the utility makes a change, it modifies this file to reflect the new configuration.
-
-

Options

add

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

- `asr`: Turns on Auto Service Request monitoring on Oracle Big Data Appliance and activates assets on ASR Manager. The installation process prompts you for the ASR Manager host name, port number, user, and password. See [Chapter 5](#) for more information about Auto Service Request.
- `em`: Installs support for the Oracle Enterprise Manager system monitoring plug-in for Oracle Big Data Appliance. The installation process prompts you for the Oracle Management System (OMS) host name and port number, the Enterprise Manager super-administrative user and password, and the Enterprise Manager agent registration password.

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

```
# cd /opt/oracle/BDAMammoth
# ./mammoth-reconfig add asr
```

See ["Software Configuration"](#) on page 4-8 for detailed descriptions of the required information, and ["Changing the Configuration of Optional Software"](#) on page 9-7 for example output.

remove

Removes a service from the cluster. The *parameter* is a keyword that identifies the service.

- `em`: Removes support for the Oracle Enterprise Manager system monitoring plug-in. You must provide the Enterprise Manager super-administrative user password.

This example removes Oracle Enterprise Manager support from all servers in the cluster:

```
# cd /opt/oracle/BDAMammoth
# ./mammoth-reconfig remove em
```

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](#)
- [reimagecluster](#)
- [reimagerack](#)

makebdaimage

Reinstalls the base image on a single server, a cluster, or an entire rack. Both [reimagecluster](#) and [reimagerack](#) 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.

--noiloms

The reimaging process does not alter the Oracle ILOMs on the target servers.

target_servers

The 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 `BdaDeploy.json`. You can find them in `/opt/oracle/bda` on configured servers.
- **to.json**: Concatenation of JSON files similar to `BdaDeploy.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 taken as is.

reimagecluster

Reimages all servers in the cluster in parallel using `dcli` and `makebdaimage`.

The `reimagecluster` utility has this syntax:

```
reimagecluster [--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 `BdaDeploy.json`.

This option defaults to `/opt/oracle/bda/BdaDeploy.json`. If `BdaDeploy.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 `reimagecluster`.

to.json

The full path to the new network configuration in effect after reimaging. It can be either a `BdaDeploy.json` or a `BdaShip.json` file.

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.

current.json

The full path to the current configuration file, either BdaShip.json or BdaDeploy.json.

This option defaults to /opt/oracle/bda/BdaDeploy.json. If BdaDeploy.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.

new.json

The full path to the new network configuration in effect after reimaging. It can be either a BdaDeploy.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:

- [Chapter 10, "Maintaining Oracle Big Data Appliance"](#)
- [Chapter 11, "Monitoring the Health of Oracle Big Data Appliance"](#)
- [Chapter 12, "Using the dcli Utility"](#)
- [Chapter 13, "Using Oracle Integrated Lights Out Manager"](#)

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](#)
- [Powering On and Off Oracle Big Data Appliance](#)
- [Adding Memory to the Servers](#)
- [Maintaining the Physical Disks of Servers](#)
- [Replacing a Server Disk](#)
- [Maintaining the InfiniBand Network](#)
- [Changing the Number of Connections to a Gateway Switch](#)
- [Changing the NTP Servers](#)
- [Monitoring the PDU Current](#)

See Also: [Chapter 12, "Using the dcli Utility"](#)

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"](#) on page 12-1.
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 2-10](#).

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
.
.
.
```

Powering On and Off Oracle Big Data Appliance

This section includes the following topics:

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

Nonemergency Power Procedures

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

Powering On Oracle Big Data Appliance

Oracle Big Data Appliance is powered on by either pressing the power button on the front of the servers, or by logging in to the Oracle ILOM interface and applying power to the system.

To power on Oracle Big Data Appliance:

1. Turn on all 12 breakers on both PDUs.
Allow 4 to 5 minutes for Oracle ILOM to start.
2. Power up the servers.

Powering On Servers Remotely Using Oracle ILOM

You can power on the servers remotely using the Oracle ILOM interface. You can access Oracle ILOM using the web console, the command-line interface (CLI), the intelligent platform management interface (IPMI), or the simple network management protocol interface (SNMP). For example, to apply power to server bda1node01 using IPMI, run the following command as `root` from a server that has `ipmitool` installed:

```
ipmitool -H bda1node01-c -U root chassis power on
```

In this example, `bda1node01-c` is the host name of Oracle ILOM for the server to be powered on. You are prompted for the password.

See Also: [Chapter 13, "Using Oracle Integrated Lights Out Manager"](#)

Powering Off Oracle Big Data Appliance

To power off Oracle Big Data Appliance:

1. Power off the servers.
2. Turn off all 12 breakers on both PDUs.

Powering Off the Servers Use the Linux `shutdown` command to power off or restart the servers. Enter this command as `root` to shut down a server immediately:

```
# shutdown -hP now
```

The following command restarts a server immediately:

```
# shutdown -r now
```

See Also: Linux SHUTDOWN manual page for details

Powering Off Multiple Servers at the Same Time Use the `dcli` utility to run the `shutdown` command on multiple servers at the same time. Do not run the `dcli` utility from a server that will be shut down. Set up passwordless SSH for `root`, as described in ["Setting Up Passwordless SSH"](#) on page 12-1.

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: [Chapter 12, "Using the dcli Utility"](#)

Powering On and Off Network Switches

The network switches do not have power switches. They power off when power is removed by turning off a PDU or a breaker in the data center.

Emergency Power-Off Considerations

In an emergency, halt power to Oracle Big Data Appliance immediately. The following emergencies may require powering off Oracle Big Data Appliance:

- Natural disasters such as earthquake, flood, hurricane, tornado, or cyclone
- Abnormal noise, smell, or smoke coming from the system
- Threat to human safety

Emergency Power-Off Procedure

To perform an emergency power-off procedure for Oracle Big Data Appliance, turn off power at the circuit breaker or pull the emergency power-off switch in the computer room. After the emergency, contact Oracle Support Services to restore power to the system.

Emergency Power-Off Switch

Emergency power-off (EPO) switches are required when computer equipment contains batteries capable of supplying more than 750 volt-amperes for more than 5 minutes. Systems that have these batteries include internal EPO hardware for connection to a site EPO switch or relay. Use of the EPO switch removes power from Oracle Big Data Appliance.

Cautions and Warnings

The following cautions and warnings apply to Oracle Big Data Appliance:

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

Caution:

- Do not power off Oracle Big Data Appliance unless there is an emergency. In that case, follow the "[Emergency Power-Off Procedure](#)" on page 10-3.
 - Keep the front and rear cabinet doors closed. Failure to do so might cause system failure or result in damage to hardware components.
 - Keep the top, front, and back of the cabinets clear to allow proper airflow and prevent overheating of components.
 - Use only the supplied hardware.
-
-

Adding Memory to the Servers

You can add memory to all servers in the cluster or to specific servers that need more memory, such as the two NameNodes.

See Also: "[Parts for Oracle Big Data Appliance Servers](#)" on page B-5

Adding Memory to Sun Server X3-2L Servers

Oracle Big Data Appliance X3-2 ships from the factory with 64 GB of memory in each server. Eight of the 16 DIMM slots are populated with 8 GB DIMMs.

Oracle Big Data Appliance X3-2 supports 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 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 a Sun Server X3-2L server:

1. If you are mixing DIMM sizes, then review the DIMM population rules in the *Sun Server X3-2L Service Manual* at http://docs.oracle.com/cd/E23393_01/html/E27229/ceiehcdb.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 8 GB DIMMs last.

See the *Sun Server X3-2L Service Manual* at

http://docs.oracle.com/cd/E23393_01/html/E27229/ceicjagi.html#scrolltoc.

4. Power on the server.

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.

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](#)" on page B-5 for the repair procedures

Verifying the Server Configuration

The 12 disk drives in each Oracle Big Data Appliance server are controlled by an LSI MegaRAID SAS 92610-8i disk controller. Oracle recommends verifying the status of the RAID devices to avoid possible performance degradation or an outage. The effect on the server of validating the RAID devices is minimal. The corrective actions may affect operation of the server and can range from simple reconfiguration to an outage, depending on the specific issue uncovered.

Verifying Disk Controller Configuration

Enter this command to verify the disk controller configuration:

```
# MegaCli64 -AdpAllInfo -a0 | grep "Device Present" -A 8
```

The following is an example of the output from the command. There should be 12 virtual drives, no degraded or offline drives, and 14 physical devices. The 14 devices are the controllers and the 12 disk drives.

```

                                Device Present
                                =====
Virtual Drives      : 12
  Degraded          : 0
  Offline           : 0
Physical Devices   : 14
  Disks             : 12
  Critical Disks    : 0
  Failed Disks      : 0

```

If the output is different, then investigate and correct the problem.

Verifying Virtual Drive Configuration

Enter this command to verify the virtual drive configuration:

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

The following is an example of the output for Virtual Drive 0. Ensure that State is Optimal.

```
Adapter 0 -- Virtual Drive Information:
Virtual Drive: 0 (Target Id: 0)
Name                :
RAID Level           : Primary-0, Secondary-0, RAID Level Qualifier-0
Size                : 1.817 TB
Parity Size         : 0
State               : Optimal
Strip Size          : 64 KB
Number Of Drives    : 1
Span Depth          : 1
Default Cache Policy: WriteBack, ReadAheadNone, Cached, No Write Cache if Bad BBU
Current Cache Policy: WriteBack, ReadAheadNone, Cached, No Write Cache if Bad BBU
Access Policy       : Read/Write
Disk Cache Policy   : Disk's Default
Encryption Type     : None
```

Verifying Physical Drive Configuration

Use the following command to verify the physical drive configuration:

```
# MegaCli64 -PDList -a0 | grep Firmware
```

The following is an example of the output from the command. The 12 drives should be Online, Spun Up. If the output is different, then investigate and correct the problem.

```
Firmware state: Online, Spun Up
Device Firmware Level: 061A
Firmware state: Online, Spun Up
Device Firmware Level: 061A
Firmware state: Online, Spun Up
Device Firmware Level: 061A
.
.
.
```

Replacing a Server Disk

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

This section contains the following topics:

- [Overview of the Disk Replacement Process](#)
- [About Disk Drive Identifiers](#)
- [Prerequisites for Replacing a Working Disk](#)
- [What If a Server Fails to Restart?](#)
- [Replacing a Disk Drive](#)
- [Identifying the Function of a Disk Drive](#)

- [Configuring an Operating System Disk](#)
- [Configuring an Oracle NoSQL Database Disk](#)
- [Configuring an HDFS Disk](#)
- [Verifying the Disk Configuration](#)

Overview of the Disk Replacement Process

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

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

See Also:

"Servicing Storage Drives and Rear Drives" in the *Sun Server X3-2L Service Manual* at

http://docs.oracle.com/cd/E23393_01/html/E27229/z40000091011460.html

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

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

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 10-1](#).

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.

Standard Disk Drive Mappings

Table 10–1 shows the mappings between the RAID logical drives and the operating system identifiers, and the dedicated function of each drive in an Oracle Big Data Appliance server. Nonetheless, you must confirm that these mappings are correct on your system.

Table 10–1 Disk Drive Identifiers

| Symbolic Link to Physical Slot | Initial Operating System Location | 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 |
| /dev/disk/by-hba-slot/s3 | /dev/sdd | HDFS |
| /dev/disk/by-hba-slot/s4 | /dev/sde | HDFS |
| /dev/disk/by-hba-slot/s5 | /dev/sdf | HDFS |
| /dev/disk/by-hba-slot/s6 | /dev/sdg | HDFS |
| /dev/disk/by-hba-slot/s7 | /dev/sdh | HDFS |
| /dev/disk/by-hba-slot/s8 | /dev/sdi | HDFS |
| /dev/disk/by-hba-slot/s9 | /dev/sdj | HDFS |
| /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 |

Standard Mount Points

Table 10–2 show the mappings between HDFS partitions and mount points.

Table 10–2 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 |

Obtaining the Physical Slot Number of a Disk Drive

Use the following `MegaCli64` command to verify the mapping of virtual drive numbers to physical slot numbers. See ["Replacing a Disk Drive"](#) on page 10-11.


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

Prerequisites for Replacing a Working Disk

If you plan to replace an HDFS disk or an operating system disk before it fails, then you should first dismount the HDFS partitions. 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 ext3 (rw,noatime)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
/dev/md0 on /boot type ext3 (rw)
tmpfs on /dev/shm type tmpfs (rw)
/dev/sda4 on /u01 type ext4 (rw,nodev,noatime) [/u01]
/dev/sdb4 on /u02 type ext4 (rw,nodev,noatime) [/u02]
/dev/sdc1 on /u03 type ext4 (rw,nodev,noatime) [/u03]
/dev/sdd1 on /u04 type ext4 (rw,nodev,noatime) [/u04]
.
.
.
```

4. Check the list of mounted partitions for the failing disk. If the disk has no partitions listed, then proceed to ["Replacing a Disk Drive"](#) on page 10-11. 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"](#) on page 10-11. 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://bdainode03.example.com:7180`
7. Complete these steps in Cloudera Manager:
 - 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**.

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

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

```
# umount mountpoint
```
10. Complete the steps in "[Replacing a Disk Drive](#)" on page 10-11.

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://bdainode12-c.example.com`

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

2. Log in using your Oracle ILOM credentials.
3. Select the **Remote Control** tab.
4. Click the **Launch Remote Console** button.
5. Enter `Ctrl+d` to continue rebooting.

6. If the reboot fails, then enter the server root password at the prompt and attempt to fix the problem.
7. After the server restarts successfully, open the Redirection menu and choose **Quit** to close the console window.

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

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

Replacing a Disk Drive

Complete this procedure for all failed disk drives.

1. If you are replacing a working disk, then see "[Prerequisites for Replacing a Working Disk](#)" on page 10-9.
2. Replace the failed disk drive.
See "[Parts for Oracle Big Data Appliance Servers](#)" on 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 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.

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 fails because of cached data, then clear the cache:

```
# MegaCli64 discardpreservedcache l1 a0
```

11. Verify the disk is recognized by the operating system.

```
# lsscsi
```

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

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

```
[0:0:20:0]   enclosu SUN      HYDE12          0341  -
[0:2:0:0]    disk     LSI      MR9261-8i      2.12  /dev/sda
```

```

[0:2:1:0]   disk   LSI      MR9261-8i    2.12 /dev/sdb
[0:2:2:0]   disk   LSI      MR9261-8i    2.12 /dev/sdc
[0:2:3:0]   disk   LSI      MR9261-8i    2.12 /dev/sdd
[0:2:4:0]   disk   LSI      MR9261-8i    2.12 /dev/sde
[0:2:5:0]   disk   LSI      MR9261-8i    2.12 /dev/sdn
[0:2:6:0]   disk   LSI      MR9261-8i    2.12 /dev/sdg
[0:2:7:0]   disk   LSI      MR9261-8i    2.12 /dev/sdh
[0:2:8:0]   disk   LSI      MR9261-8i    2.12 /dev/sdi
[0:2:9:0]   disk   LSI      MR9261-8i    2.12 /dev/sdj
[0:2:10:0]  disk   LSI      MR9261-8i    2.12 /dev/sdo
[0:2:11:0]  disk   LSI      MR9261-8i    2.12 /dev/sdl
[7:0:0:0]   disk   Unigen  PSA4000      1100 /dev/sdm
[

```

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

```
# reboot
```

13. Reconnect to the server and reissue the `lsscsi` command to verify the original disk mappings. See [Table 10-1](#).

14. Check the hardware profile of the server:

```
# bdacheckhw
```

15. Identify the function of the drive, so you configure it properly. See ["Identifying the Function of a Disk Drive"](#) on page 10-13.

Identifying the Function of a Disk Drive

Most disks are used for HDFS, as shown in [Table 10-1](#). Nonetheless, you should verify that the failed disk was not used for either the operating system or Oracle NoSQL Database before configuring it for a particular function.

Checking for Use by the Operating System

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

To confirm that a failed disk supported the operating system:

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

```
# lsscsi
```

See the output from Step 11 of ["Replacing a Disk Drive"](#) on page 10-11.

2. Verify that `/dev/sda` and `/dev/sdb` are the operating system mirrored partitioned disks:

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

```

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

```

3. If the previous steps indicate that the failed disk is an operating system disk, then proceed to ["Configuring an Operating System Disk"](#) on page 10-14.
4. If the failed disk did not support the operating system, then proceed to ["Checking for Use by Oracle NoSQL Database"](#) on page 10-14.

Checking for Use by Oracle NoSQL Database

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

To discover whether a failed disk supported Oracle NoSQL Database:

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


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

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

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

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

5. If the previous steps indicate that the failed disk supported Oracle NoSQL Database, then proceed to ["Configuring an Oracle NoSQL Database Disk"](#) on page 10-20. Otherwise, proceed to ["Configuring an HDFS Disk"](#) on page 10-23.

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](#)
- [Repairing the RAID Arrays](#)
- [Formatting the HDFS Partition of an Operating System Disk](#)
- [Restoring the Swap Partition](#)
- [Restoring the GRUB Master Boot Records and HBA Boot Order](#)

Partitioning the Operating System Disk

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`.

To partition a logical drive:

1. Complete the steps in ["Replacing a Disk Drive"](#) on page 10-11.

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 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 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"](#) on page 10-16.

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:

```
# 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"](#) on page 10-18.

Formatting the HDFS Partition of an Operating System Disk

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

To format the HDFS partition:

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

```
# mke4fs -t ext4 /dev/disk/by-hba-slot/smp4

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

Writing inode tables:
```

Note: If this command fails because the device is mounted, then dismount the drive now and skip step 3. See ["Prerequisites for Replacing a Working Disk"](#) on page 10-9 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:

```
# tune4fs -c -1 -i 0 -m 0.2 -L /u0n /dev/disk/by-hba-slot/smp4
```

5. Mount the HDFS partition:

```
# mount /u0n
```

6. Complete the steps in ["Restoring the Swap Partition"](#) on page 10-19.

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/smp3
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. Complete the steps in ["Restoring the GRUB Master Boot Records and HBA Boot Order"](#) on page 10-19.

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.

To restore the GRUB boot record:

1. Open GRUB:

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

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

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

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

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

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

```
grub> root (hdn,0)

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

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

```
grub> setup (hdn)
```

```

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

```

4. Close the GRUB command-line interface:

```
grub> quit
```

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

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

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

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

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

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

8. Check the configuration. See ["Verifying the Disk Configuration"](#) on page 10-24.

Configuring an Oracle NoSQL Database Disk

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

Removing the Logical Volume for Oracle NoSQL Database

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

To remove the logical volume:

1. Complete the steps in ["Replacing a Disk Drive"](#) on page 10-11.
2. Open /etc/fstab in a text editor and make these changes:
 - a. Delete the mount entry for /dev/lvg1/lv1.
 - b. Add mount points for either /u12, or both /u11 and /u12.
3. Verify that the logical volume exists:

```
# lvdisplay
--- Logical volume ---
LV Name                /dev/lvg1/lv1
VG Name                lvg1
LV UUID                BaWmSp-Ha8J-Qm18-9qEr-i1qH-mKJm-eEj1wD
LV Write Access        read/write
LV Status              available
# open                 1
LV Size                3.45 TB
```

```

Current LE          905536
Segments           1
Allocation          inherit
Read ahead sectors auto
- currently set to 256
Block device       253:0

```

4. Verify that the physical volume exists:

```

# pvdisplay
--- Physical volume ---
PV Name            /dev/sd11
VG Name            lvg1
PV Size            1.82 TB / not usable 1.97 MB
Allocatable        yes
PE Size (KByte)    4096
Total PE           476598
Free PE            23830
Allocated PE       452768
PV UUID            kHtDSR-gmvv-uLIm-69Y1-fp4Q-7bgf-tMhzZU

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

```

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

6. Dismount the logical volume.

```
# umount /lv1
```

Note: If you are unable to dismount the volume, then complete the procedure in ["Stopping the Oracle NoSQL Database Service"](#) on page 10-22..

7. Deactivate the logical volume:

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

8. Remove the logical volume:

```
# lvremove -f /dev/lvg1/lv1
Logical volume "lv1" successfully removed
```

9. Deactivate the volume group:

```
# vgchange -a n /dev/lvg1
0 logical volume(s) in volume group "lvg1" now active
```

10. Remove the volume group.

```
# vgremove -f lvg1
Volume group "lvg1" successfully removed
```

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

```
# pvdisplay

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

"/dev/sdl1" is a new physical volume of "1.82 TB"
--- NEW Physical volume ---
PV Name                /dev/sdl1
VG Name
PV Size                1.82 TB
Allocatable            NO
PE Size (KByte)       0
Total PE              0
Free PE               0
Allocated PE          0
PV UUID                kHtDSR-gmvv-uLIIm-69Y1-fp4Q-7bgf-tMhzZU
```

12. Remove the physical volumes:

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

This example removes two allocated partitions:

```
# pvremove -ff /dev/sdk1 /dev/sdl1
Labels on physical volume "/dev/sdk1" successfully wiped
Labels on physical volume "/dev/sdl1" successfully wiped
```

13. Configure the drive or drives for HDFS. See ["Configuring an HDFS Disk"](#) on page 10-23. If you have two disks allocated to Oracle NoSQL Database, then you must configure both drives (/dev/sdk and /dev/sdl).**Stopping the Oracle NoSQL Database Service**

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

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

```
# fuser -m /dev/mapper/lvg1-lv1
/dev/mapper/lvg1-lv1: 11154 11430 11443
```

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

```
# ps -fp process_id
```

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

```
# ps -fp 11154
UID PID PPID C STIME TTY TIME CMD
```

```
oracle 11154 1 0 09:29 ? 00:00:00 java -Xms64m -Xmx64m
oracle.kv.impl.sna.StorageNodeAgentImpl -root /lv1/kvroot
```

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


```
# service nsdbservice stop
```
4. Restart the configuration beginning with Step 6 in ["Removing the Logical Volume for Oracle NoSQL Database"](#) on page 10-20, which dismounts the device.

Configuring an HDFS Disk

An HDFS disk is not used by the operating system nor by Oracle NoSQL Database. See ["Identifying the Function of a Disk Drive"](#) on page 10-13.

Partitioning a Disk for HDFS

To configure an HDFS 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:

1. Complete the steps in ["Replacing a Disk Drive"](#) on page 10-11, if you have not done so already.
2. Partition the drive:


```
# parted /dev/disk/by-hba-slot/sn -s mklabel gpt mkpart primary ext3 0% 100%
```
3. Format the HDFS partition as an ext4 file system:

```
# mke4fs -t ext4 /dev/disk/by-hba-slot/snp1

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

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

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

4. Reset the appropriate partition label to the missing device. See [Table 10–2](#).

```
# tune4fs -c -1 -i 0 -m 0.2 -L /unn /dev/disk/by-hba-slot/smp1
```

For example, this command resets the label for /dev/disk/by-hba-slot/s2p1 to /u03:

```
# tune4fs -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. Mount the HDFS partition, entering the appropriate mount point:

```
# mount /unn
```

For example, mount /u03.

6. Restart the server.
7. If you are configuring multiple drives, then repeat the previous steps.
8. If you previously removed a mount point in Cloudera Manager, 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**.
9. Check the configuration. See "[Verifying the Disk Configuration](#)" on page 10-24.

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:


```
# bdachcksw
```
2. If there are errors, then redo the configuration steps as necessary to correct the problem.
3. When bdachcksw returns no errors, restart the server.
4. Check the /root directory for a file named BDA_REBOOT_SUCCEEDED.

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

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

Maintaining the InfiniBand Network

The InfiniBand network connects the servers through the bondib0 interface to the InfiniBand switches in the rack. This section describes how to perform maintenance on the InfiniBand switches.

This section contains the following topics:

- [Backing Up and Restoring Oracle ILOM Settings](#)
- [Replacing a Failed InfiniBand Switch](#)
- [Verifying InfiniBand Network Operation](#)
- [Understanding the Network Subnet Manager Master](#)

Backing Up and Restoring Oracle ILOM Settings

Oracle ILOM supports remote administration of the Oracle Big Data Appliance servers. This section explains how to back up and restore the Oracle ILOM configuration settings, which are set by the Mammoth Utility.

See Also: Oracle Integrated Lights Out Manager 3.0 documentation at

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

Backing Up the Oracle ILOM Configuration Settings

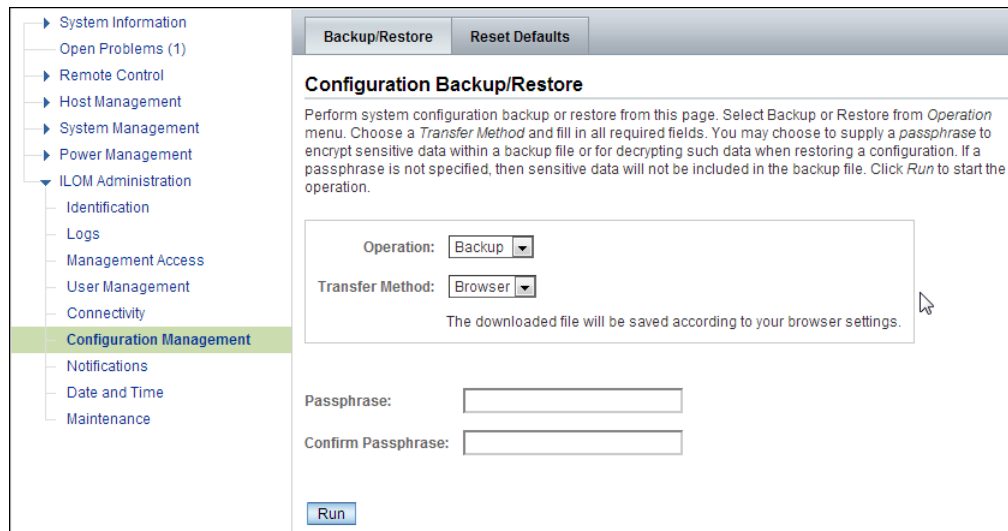
To back up the Oracle ILOM configuration settings:

1. Open your browser on any system that is on the same network as Oracle Big Data Appliance, and enter the Oracle ILOM address of a server. This example uses the Oracle ILOM address of node08:


```
http://bda1node08-c.example.com
```
2. Log in as the `root` user. The initial password is `welcome1`.
3. In the navigation tree, expand the ILOM Administration folder and select **Configuration Management**.
4. Select the Backup operation and the Browser transfer method, as shown in [Figure 10-1](#).
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 10–1 Oracle ILOM Configuration Backup



Restoring Oracle ILOM Configuration Settings

To restore the Oracle ILOM configuration settings:

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

```
http://bda1node08-c.us.example.com
```
2. Log in as the `ilom-admin` user. The default password is `welcome1`.
3. Select the **Maintenance** tab.
4. Select the **Backup/Restore** tab.
5. Select the Restore operation and the Browser transfer method.
6. Click **Choose File** and select the `config_backup.xml` file saved previously in a backup operation.
7. Enter the pass phrase that was set during the backup operation.
8. Click **Run** to restore the configuration.

Replacing a Failed InfiniBand Switch

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

See Also:

- "In-Rack InfiniBand Switch-to-Server Cable Connections" on page D-7 for information about cabling
- *Sun Network QDR InfiniBand Gateway Switch Installation Guide* at <http://docs.oracle.com/cd/E19671-01/821-1186-11/821-1186-11.pdf>
- *Sun Datacenter InfiniBand Switch 36 User's Guide* at <http://docs.oracle.com/cd/E19197-01/index.html>

To replace a failed InfiniBand switch:

1. Disconnect the cables from the switch. All InfiniBand cables have labels at both ends indicating their locations. If any cables do not have labels, then label them.
2. Power off both power supplies on the switch by removing the power plugs.
3. Remove the switch from the rack.
4. Install the new switch in the rack.
5. Restore the switch settings using the backup file, as described in "Backing Up and Restoring Oracle ILOM Settings" on page 10-25.

6. Connect to the switch as `ilom_admin` and open the Fabric Management shell:

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

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

7. Disable the Subnet Manager:

```
FabMan@bdalsw-02-> disablesm
```

8. Connect the cables to the new switch, being careful to connect each cable to the correct port.

9. Verify that there are no errors on any links in the fabric:

```
FabMan@bdalsw-02-> ibdiagnet -c 1000 -r
```

10. Enable the Subnet Manager:

```
FabMan@bdalsw-02-> enablesm
```

Note: If the replaced switch was the Sun Datacenter InfiniBand Switch 36 spine switch, then manually fail the master Subnet Manager back to the switch by disabling the Subnet Managers on the other switches until the spine switch becomes the master. Then reenabling the Subnet Manager on all the other switches.

Verifying InfiniBand Network Operation

If any component in the InfiniBand network has required maintenance, including replacing an InfiniBand Host Channel Adapter (HCA) on a server, an InfiniBand switch, or an InfiniBand cable, or if operation of the InfiniBand network is suspected to be substandard, then verify that the InfiniBand network is operating properly. The following procedure describes how to verify network operation:

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

To verify InfiniBand network operation:

1. Enter the `ibdiagnet` command to verify InfiniBand network quality:

```
# ibdiagnet -c 1000
```

Investigate all errors reported by this command. It generates a small amount of network traffic and can run during a normal workload.

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

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

2. Report switch port error counters and port configuration information. The `LinkDowned`, `RcvSwRelayErrors`, `XmtDiscards`, and `XmtWait` errors are ignored by this command:

```
# ibqueryerrors.pl -rR -s LinkDowned,RcvSwRelayErrors,XmtDiscards,XmtWait
```

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

3. Check the status of the hardware:

```
# bdacheckhw
```

The following is an example of the output:

```
[SUCCESS: Correct system model : SUN FIRE X4270 M2 SERVER
[SUCCESS: Correct processor info : Intel(R) Xeon(R) CPU X5675 @ 3.07GHz
[SUCCESS: Correct number of types of CPU : 1
[SUCCESS: Correct number of CPU cores : 24
[SUCCESS: Sufficient GB of memory (>=48): 48
[SUCCESS: Correct GB of swap space : 24
[SUCCESS: Correct BIOS vendor : American Megatrends Inc.
[SUCCESS: Sufficient BIOS version (>=08080102): 08080102
[SUCCESS: Recent enough BIOS release date (>=05/23/2011) : 05/23/2011
[SUCCESS: Correct ILOM version : 3.0.16.10.a r68533
[SUCCESS: Correct number of fans : 6
[SUCCESS: Correct fan 0 status : ok
[SUCCESS: Correct fan 1 status : ok
[SUCCESS: Correct fan 2 status : ok
[SUCCESS: Correct fan 3 status : ok
[SUCCESS: Correct fan 4 status : ok
[SUCCESS: Correct fan 5 status : ok
[SUCCESS: Correct number of power supplies : 2
[Im[34mINFO: Detected Santa Clara Factory, skipping power supply checks
[SUCCESS: Correct disk controller model : LSI MegaRAID SAS 9261-8i
[SUCCESS: Correct disk controller firmware version : 12.12.0-0048
[SUCCESS: Correct disk controller PCI address : 13:00.0
[SUCCESS: Correct disk controller PCI info : 0104: 1000:0079
[SUCCESS: Correct disk controller PCIe slot width : x8
[SUCCESS: Correct disk controller battery type : iBBU08
[SUCCESS: Correct disk controller battery state : Operational
[SUCCESS: Correct number of disks : 12
```

```
[SUCCESS: Correct disk 0 model : SEAGATE ST32000SSSUN2.0
[SUCCESS: Sufficient disk 0 firmware (>=61A): 61A
[SUCCESS: Correct disk 1 model : SEAGATE ST32000SSSUN2.0
[SUCCESS: Sufficient disk 1 firmware (>=61A): 61A
.
.
.
[SUCCESS: Correct disk 10 status : Online, Spun Up No alert
[SUCCESS: Correct disk 11 status : Online, Spun Up No alert
[SUCCESS: Correct Host Channel Adapter model : Mellanox Technologies MT26428
ConnectX VPI PCIe 2.0
[SUCCESS: Correct Host Channel Adapter firmware version : 2.9.1000
[SUCCESS: Correct Host Channel Adapter PCI address : 0d:00.0
[SUCCESS: Correct Host Channel Adapter PCI info : 0c06:15b3:673c
[SUCCESS: Correct Host Channel Adapter PCIe slot width : x8
[SUCCESS: Big Data Appliance hardware validation checks succeeded
```

4. Check the status of the software:

```
# 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 : 5.5.17
[SUCCESS: All required programs are accessible in $PATH
[SUCCESS: All required RPMs are installed and valid
[SUCCESS: Big Data Appliance software validation checks succeeded
```

Understanding the Network Subnet Manager Master

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

- Discover the network topology
- Assign a local identifier to all ports connected to the network
- Calculate and program switch forwarding tables

- Monitor changes in the fabric

The InfiniBand network can have multiple Subnet Managers, but only one Subnet Manager is active at a time. The active Subnet Manager is the Master Subnet Manager. The other Subnet Managers are the Standby Subnet Managers. If a Master Subnet Manager is shut down or fails, then a Standby Subnet Manager automatically becomes the Master Subnet Manager.

Each Subnet Manager has a configurable priority. When multiple Subnet Managers are on the InfiniBand network, the Subnet Manager with the highest priority becomes the master Subnet Manager. On Oracle Big Data Appliance, the Subnet Managers on the leaf switches are configured as priority 5, and the Subnet Managers on the spine switches are configured as priority 8.

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

- Run the Subnet Managers only on the switches in Oracle Big Data Appliance. Running a Subnet Manager on any other device is not supported.
- When the InfiniBand network consists of one, two, or three racks cabled together, all switches must run a Subnet Manager. The master Subnet Manager runs on a spine switch.
- When the InfiniBand network consists of four or more racks cabled together, then only the spine switches run a Subnet Manager. The leaf switches must disable the Subnet Manager.

See Also:

- Sun Network QDR InfiniBand Gateway library at <http://docs.oracle.com/cd/E19671-01/>
- Sun Datacenter InfiniBand Switch 36 library at <http://docs.oracle.com/cd/E19197-01/index.html>

Changing the 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](#)" on page 11-14.

To re-create the VNICs in a rack:

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

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

2. Connect to node01 (bottom of rack) using 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](#)" on page 12-1 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](#).

6. Restore passwordless SSH (optional):

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

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](#)" on page 7-13.

To update the Cisco Ethernet switch:

1. Use telnet to connect to the Cisco Ethernet switch.
2. Delete the current setting:

```
# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
(config)# no ntp server current_IPaddress
```

3. Enter the new IP address:

```
# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
(config)# ntp server new_IPaddress
```

4. Save the current configuration:

```
# copy running-config startup-config
```

5. Exit from the session:

```
# exit
```

See Also: "[Configuring the Cisco Ethernet Switch](#)" on page 7-4

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

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/index.html>

Table 10–3 lists the threshold values for the Oracle Big Data Appliance rack using a single-phase, low-voltage PDU.

Table 10–3 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 10–4 lists the threshold values for the Oracle Big Data Appliance rack using a three-phase, low-voltage PDU.

Table 10–4 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 10–5 lists the threshold values for the Oracle Big Data Appliance rack using a single-phase, high-voltage PDU.

Table 10–5 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 10–6 lists the threshold values for the Oracle Big Data Appliance rack using a three-phase, high-voltage PDU.

Table 10–6 *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 |

Monitoring the Health of Oracle Big Data Appliance

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`
- `bdcheckhw`
- `bdcheckib`
- `bdchecknet`
- `bdchecksw`
- `bdadiag`
- `bdagetclientconfig`
- `bdaid`
- `bdaimagevalidate`
- `bdaredoclientnet`
- `bdaserials`
- `bdaswapoff`
- `bdaswapon`
- `bdaupdatefw`
- `iblinkinfo`
- `imagehistory`
- `imageinfo`
- `listlinkup`
- `remove-root-ssh`
- `setup-root-ssh`
- `showvlan`
- `showvnics`

bdcheckcluster

Checks the health of a CDH cluster, including the software, hardware, and network, and logs the results in a file in the /tmp directory.

Syntax

```
bdcheckcluster
```

Example

This example shows the output from the utility:

```
# bdcheckcluster
INFO: Logging results to /tmp/bdcheckcluster_1358517388/
SUCCESS: Mammoth configuration file is valid.
SUCCESS: All cluster host names are pingable
SUCCESS: All cluster hosts passed checks on last reboot
INFO: Starting cluster host hardware checks
SUCCESS: All cluster hosts pass hardware checks
INFO: Starting cluster host software checks
SUCCESS: All cluster hosts pass software checks
SUCCESS: All ILOM hosts are pingable
SUCCESS: All client interface IPs are pingable
SUCCESS: All admin eth0 interface IPs are pingable
SUCCESS: All private Infiniband interface IPs are pingable
SUCCESS: Puppet master is running on bda1node01-master
SUCCESS: Puppet running on all cluster hosts
SUCCESS: Cloudera SCM server is running on bda1node03
SUCCESS: Cloudera SCM agent running on all cluster hosts
.
.
.
SUCCESS: Big Data Appliance cluster health checks succeeded
```

bdacheckhw

Checks the hardware profile of the server.

Syntax

```
bdacheckhw
```

Usage Notes

See "[Configuring the Oracle Big Data Appliance Servers](#)" on page 7-18 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
```

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/BdaDeploy.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: bdalsw-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: bdalsw-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: bdalsw-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
.
.
.
```

bdchecknet

Checks whether the network configuration is working properly.

Syntax

```
bdchecknet
```

Usage Notes

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

Example

This example shows the output from the utility:

```
[root@node01 network]# bdchecknet
bdchecknet: check syntax and static semantics of /opt/oracle/bda/BdaDeploy.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
```

bdachecksw

Checks the software profile of a server.

Syntax

```
bdachecksw
```

Usage Notes

See ["Configuring the Oracle Big Data Appliance Servers"](#) on page 7-18 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 : Linux 2.6.32-200.21.1.el5uek
SUCCESS: Correct Java Virtual Machine version : HotSpot(TM) 64-Bit Server 1.6.0_37
SUCCESS: Correct puppet version : 2.6.11
SUCCESS: Correct MySQL version : 5.5.17
SUCCESS: All required programs are accessible in $PATH
SUCCESS: All required RPMs are installed and valid
SUCCESS: Correct bda-monitor status : bda monitor is running
SUCCESS: Big Data Appliance software validation checks succeeded
```

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 v2.0.1

Gathering Linux information

Skipping ILOM collection. Use the ilom or snapshot options, or login to
ILOMover 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_2013_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 v2.0.1

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" : "2013-01-18T07:25", "endTime" : "2013-01-17T07:25" }
Succeeded. Output in : /opt/oracle/BDAMammoth/bdaconfig/tmp/cm_commands_
collectDiagnosticData.out
Command ID is 159
....
Command 159 finished after 25 seconds
Operation completed successfully
Diagnostic data successfully collected
Can be downloaded from URL
http://bda1node03.us.oracle.com:7180/cmf/command/159/download
Downloading diagnostic data ...
3609df48-4930-11e1-9006-b8ac6f8061c1.scaj.20131801-15-30.support-bundle.zip
Data successfully downloaded and placed in /tmp/bdadiag_bda1node01_1216FM5497_
2013_01_18_07_
24/3609df48-4930-11e1-9006-b8ac6f8061c1.scaj.20131801-15-30.support-bundle.zip
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.

.
.
.
```

The next example shows the additional output from the snapshot option:

```
# bdadiag snapshot
```

```
Big Data Appliance Diagnostics Collection Tool v2.0.1
```

```
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: Fri Jan 18 07:35:23 PST 2013
```

```
snapshot running: Fri Jan 18 07:36:26 PST 2013
```

```
snapshot running: Fri Jan 18 07:37:30 PST 2013
```

```
Snapshot Collection completed.
```

```
Generating diagnostics tarball and removing temp directory
```

```
=====  
Done. The report files are bzip2 compressed in /tmp/bdadiag_bda1node01_1216FM5497_  
2013_01_18_07_33.tar.bz2  
=====
```

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 !
```

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
Cluster Name         : Cluster 1
Appliance Name       : bda1
```

bdaimagevalidate

Validates the hardware and software by running [bdacheckhw](#), and then [bdachecksw](#).

Syntax

```
bdaimagevalidate
```

Example

```
# bdaimagevalidate
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: Correct puppet version : 2.6.11
SUCCESS: Correct MySQL version : 5.5.17
SUCCESS: All required programs are accessible in $PATH
SUCCESS: All required RPMs are installed and valid
SUCCESS: Correct bda-monitor status : bda monitor is running
SUCCESS: Big Data Appliance software validation checks succeeded
SUCCESS: All Big Data Appliance validation checks succeeded
```

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

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"](#) on page 10-30.

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: check syntax and static semantics of
/opt/oracle/bda/BdaDeploy.json
bdaredoclientnet: passed
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 bdalnode02.us.oracle.com
hello from bdalnode03.us.oracle.com
.
.
.
bdaredoclientnet: passed
bdaredoclientnet: check existence of default vlan for port 0A-ETH-1 on bdalsw-ib2
bdaredoclientnet: use existing default vlan for port 0A-ETH-1 on bdalsw-ib2
bdaredoclientnet: check existence of default vlan for port 0A-ETH-1 on bdalsw-ib3
bdaredoclientnet: use existing default vlan for port 0A-ETH-1 on bdalsw-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 bdalnode02 on switch bdalsw-ib2
vNIC ID 757 deleted
IO Adapter for vNIC deleted
check and delete vNIC for bdalnode02 on switch bdalsw-ib3
check and delete vNIC for bdalnode02 on switch bdalsw-ib2
check and delete vNIC for bdalnode02 on switch bdalsw-ib3
vNIC ID 707 deleted
IO Adapter for vNIC deleted
create vNIC eth9 bdalnode02 on switch bdalsw-ib3
vNIC created
create vNIC eth8 bdalnode02 on switch bdalsw-ib2
vNIC created
.
.
.
bdaredoclientnet: ping server ips on client network
bdaredoclientnet: passed
bdaredoclientnet: test ssh server ips on client network
hello from bdalnode02.us.oracle.com
hello from bdalnode03.us.oracle.com
.
.
.
bdaredoclientnet: passed
bdaredoclientnet: end setup client network
```

bdserials

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

```
bdserials
```

Example

This example shows the output from the utility:

```
# bdserials
Rack serial number :
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
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
ib1 InfiniBand : 80:00:00:4B:FE:80:00:00:00:00:00:00:00:00:00:00:00:00:00:00
```

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"](#) on page 11-18

Example

This example turns off swapping:

```
# bdaswapoff  
#
```

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"](#) on page 11-17

Example

This example turns on swapping:

```
# bdaswapon
Filename                               Type      Size    Used    Priority
/dev/sda3                               partition 12287992 0      1
/dev/sdb3                               partition 12287992 0      1
```

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, the Mammoth Utility 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. See "[Oracle Big Data Appliance Restrictions on Use](#)" on page 1-4.

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/or
```

```
acle/bda/firmware/ILOM-3_2_0_r74388-Sun_Fire_X4270_M3.pkg
[INFO:GENERAL] Updating ILOM firmware with the firmware file /opt/oracle/bda/fir
mware/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 afterwa
rds.
[INFO:GENERAL]
[INFO:GENERAL] ipmiflash -v -I lanplus -H 10.133.46.218 -U root write /opt/oracl
e/bda/firmware/ILOM-3_1_2_12_r74388-Sun_Fire_X4270_M3.pkg
[INFO:GENERAL]
```

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@bdalnode01 network]# iblinkinfo
Switch 0x002128df348ac0a0 SUN IB QDR GW switch bdalsw-ib2 10.133.43.36:
 149 1[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 130 2[ ] "SUN IB QDR GW switch bdalsw-ib2 10.133...
 149 2[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 127 1[ ] "SUN IB QDR GW switch bdalsw-ib2 10.133...
 149 3[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 111 2[ ] "SUN IB QDR GW switch bdalsw-ib2 10.133...
 149 4[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 109 1[ ] "SUN IB QDR GW switch bdalsw-ib2 10.133...
 149 5[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 143 1[ ] "bdalnode02 BDA 192.168.41.20 HCA-1" ( )
 149 6[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 137 1[ ] "bdalnode01 BDA 192.168.41.19 HCA-1" ( )
 149 7[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 141 1[ ] "bdalnode04 BDA 192.168.41.22 HCA-1" ( )
 149 8[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 123 1[ ] "bdalnode03 BDA 192.168.41.21 HCA-1" ( )
 149 9[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 151 1[ ] "bdalnode06 BDA 192.168.41.24 HCA-1" ( )
 149 10[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 112 1[ ] "bdalnode05 BDA 192.168.41.23 HCA-1" ( )
 149 11[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 139 1[ ] "bdalnode07 BDA 192.168.41.25 HCA-1" ( )
 149 12[ ] == (                               Down/Disabled) ==> [ ] "" ( )
 149 13[ ] == (                               Down/Disabled) ==> [ ] "" ( )
 149 14[ ] == ( 4X 10.0 Gbps Active/ LinkUp) ==> 85 9[ ] "SUN DCS 36P QDR dm01sw-ib1 10.133.40.203" ( )
 149 15[ ] == (                               Down/Disabled) ==> [ ] "" ( )
.
.
.
```

imagehistory

Displays a history of operating system upgrades.

Syntax

```
imagehistory
```

Example

This example shows that the appliance was imaged with version 2.1.0 with no upgrades:

```
$ imagehistory
Big Data Appliance Image History

IMAGE_VERSION           : 2.1.0
IMAGE_CREATION_DATE     : Sat Mar 23 20:49:49 UTC 2013
IMAGING_START_DATE     : Mon Apr 22 19:42:08 UTC 2013
IMAGING_END_DATE       : Mon Apr 22 13:31:53 PDT 2013
DEPLOYMENT_VERSION     : 2.1.1
DEPLOYMENT_START_DATE  : Mon Apr 22 17:56:41 PDT 2013
DEPLOYMENT_END_DATE    : Mon Apr 22 18:58:34 PDT 2013
```

imageinfo

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

Syntax

```
imageinfo
```

Example

This example identifies the 21.0 image:

```
$ imageinfo
Big Data Appliance Image Info

IMAGE_CREATION_DATE      : Thu Apr 18 19:48:35 PDT 2013
IMAGE_LABEL               : BDA_2.1.1_LINUX.X64_RELEASE
IMAGE_VERSION             : 2.1.1
KERNEL_VERSION            : 2.6.32-200.21.1.el5uek
BDA_RPM_VERSION           : bda-2.1.1-1
OFA_RPM_VERSION           : ofa-2.6.32-200.21.1.el5uek-1.5.5-4.0.55.4
JDK_VERSION               : jdk-1.6.0_43-fcs
HADOOP_VERSION            : 2.0.0-cdh4.2.0
```

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 `bdal1sw-ib3`:

```
[root@bdal1sw-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)
```

remove-root-ssh

Removes passwordless SSH previously established by the [setup-root-ssh](#) 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 and expansion kit [1-12]. This parameter is required in the 2.1.1 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, as listed in `/opt/oracle/bda/rack-hosts-infiniband`.

See Also

["setup-root-ssh"](#) on page 11-27, ["Setting Up Passwordless SSH"](#) on page 12-1

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
```

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 and expansion kit [1-12]. This parameter is required in the 2.1.1 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, as listed in `/opt/oracle/bda/rack-hosts-infiniband`.

See Also

["remove-root-ssh"](#) on page 11-25, ["Setting Up Passwordless SSH"](#) on page 12-1

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
```

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
```

showvnic

Lists the virtual network interface cards (VNICs) created for the switch.

Syntax

```
showvnic
```

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`:

```
# showvnic
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
.
.
.
```

Using the dcli Utility

The `dcli` utility executes commands across a group of servers on Oracle Big Data Appliance and returns the output. You use `dcli` in [Chapter 7](#).

This chapter contains the following sections:

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

Overview of the dcli Utility

The `dcli` utility executes commands on multiple Oracle Big Data Appliance servers using the administrative (`eth0`) interface to make the connections. You can run the utility from any server.

Setting Up Passwordless SSH

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

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

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

See [Table 12-1](#) for details about these options.

To set up passwordless SSH for root:

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

```
setup-root-ssh
```

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](#) on page 11-27 and [remove-root-ssh](#) on page 11-25.

Basic Use of dcli

This topic identifies some basic options to the dcli command. For a complete list of options, see [Table 12-1](#).

Getting Help

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

Identifying the Target Servers

You can identify the servers where you want the commands to run either in the command line or in a file. For a list of default target servers, use the -t option. To change the target servers for the current command, use the -c or -g options described in [Table 12-1](#).

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.

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 12-1](#).

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

Controlling the Output Levels

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

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

Note: The output from only one server (node07) is shown. The syntax in these examples executes the date command on all 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: ['bdalnode07.example.com']
```

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

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

dcli Syntax

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

Parameters

options

The options described in [Table 12-1](#). 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 12–1 dcli Options

| Option | Description |
|-----------------------------|--|
| -c <i>nodes</i> | Specifies a comma-separated list of Oracle Big Data Appliance servers where the command is executed |
| -C | Uses the list of servers in <code>/opt/oracle/bda/cluster-rack-infiniband</code> as the target. See "Identifying the Target Servers" on page 12-2. |
| -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 <code>/root/.ssh/authorized_keys</code> file. See "Setting Up Passwordless SSH" on page 12-1 for an easier alternative. |
| -l <i>userid</i> | Identifies the user ID for logging in to another server. The default ID is <code>root</code> . |
| --maxlines= <i>maxlines</i> | Identifies the maximum lines of output displayed from a command executed on multiple servers. The default is 10,000 lines. |
| -n | Abbreviates the output for non-error messages. Only the server name is displayed when a server returns normal output (return code 0). You cannot use the -n and -r options together. |
| -r <i>regexp</i> | Replaces the output with the server name for lines that match the specified regular expression |
| -s <i>sshoptions</i> | Specifies a string of options that are passed to SSH |
| --scp= <i>scpoptions</i> | Specifies a string of options that are passed to Secure Copy (SCP), when these options are different from <i>sshoptions</i> |
| --serial | Serializes execution over the servers. The default is parallel execution. |
| -t | Lists the target servers |
| --unkey | Drops the keys from the <code>authorized_key</code> files of the target servers |
| -v | Displays the verbose version of all messages |
| --version | Displays the dcli version number |
| --vmstat=VMSTATOPS | Displays the syntax of the Linux Virtual Memory Statistics utility (<code>vmstat</code>). This command returns process, virtual memory, disk, trap, and CPU activity information. To enter a <code>vmstat</code> command, enclose its options in quotation marks. For example: <code>--vmstat="-a 3 5"</code> See your Linux documentation for more information about <code>vmstat</code> . |
| -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. |

dcli Return Values

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

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

dcli Examples

Following are examples of the `dcli` utility.

This example returns the default list of target servers:

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

The next example checks the temperature of all servers:

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

bda1node01-adm.example.com: value = 22.000 degree C
bda1node02-adm.example.com: value = 22.000 degree C
bda1node03-adm.example.com: value = 22.000 degree C
bda1node04-adm.example.com: value = 23.000 degree C
.
.
.
```

Using Oracle Integrated Lights Out Manager

Oracle Integrated Lights Out Manager (Oracle ILOM) is a preinstalled component of Oracle Big Data Appliance that you can use to monitor the servers and switches. You use Oracle ILOM in [Chapter 7](#).

This chapter contains the following sections:

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

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

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

Oracle ILOM Overview

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

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

With Oracle ILOM, you can proactively:

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

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

(Oracle Linux). Using Oracle ILOM, you can remotely manage Oracle Big Data Appliance as if you were using a local KVM.

Oracle ILOM Interfaces

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

Web Interface

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

Command-Line Interface

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

Oracle ILOM Users

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.

Administrative Network Diagram

Figure 13–1 illustrates the administrative Oracle ILOM network. It shows two of the 18 servers and the two Sun Network QDR InfiniBand Gateway Switches. The Cisco Ethernet management switch is connected to the servers and the InfiniBand switches.

Figure 13–1 Administrative Network in Oracle Big Data Appliance



Oracle ILOM IP Addresses for Oracle Big Data Appliance Components

You require at least one static IP address for service processor (Oracle ILOM) access. For the list of default Oracle ILOM IP addresses assigned to Oracle Big Data Appliance components at the time of manufacturing, see "[Factory Network Settings](#)" on page 3-6.

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

Connecting to Oracle ILOM Using the Network

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

Note: You can use this connection method when Oracle ILOM IP addresses can be accessed over the network. Oracle recommends that you use this connection method. See [Chapter 7](#)

This section discusses the following topics:

- [Connecting to the Command-Line Interface](#)
- [Connecting to the Web GUI](#)

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/>

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" on page 13-2 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 13–2.

Figure 13–2 Oracle ILOM Web Interface on an Oracle Big Data Appliance X3-2 Server



See Also: *Oracle ILOM 3.0 Daily Management -- Web Procedures Guide* at

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

Connecting to Oracle ILOM Using a Serial Connection

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

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

Connecting to the Oracle ILOM of a Server

To connect to Oracle ILOM of a server using a serial connection:

1. Attach a serial cable from a terminal or a PC running terminal emulation software to the SER MGT port of an Oracle Big Data Appliance server. The cable should be 15 feet or shorter.
2. Verify that your terminal or laptop is operational.
3. Configure the terminal device or the terminal emulation software to use the following settings:
 - 8N1: eight data bits, no parity, one stop bit
 - 9600 baud (default, but can be set to any standard rate up to 115200)
 - Disable software flow control (XON/XOFF)
 - Disable hardware control
4. Verify that power is supplied to either PSU.
If power is supplied to either PSU, then Oracle ILOM is functional regardless of the power state of the servers.
5. Press Enter on the terminal device. A connection between the terminal device and Oracle ILOM is established.
The Oracle ILOM login prompt is displayed.
6. Log in to the CLI as root using the default password of welcome1.
The Oracle ILOM default command prompt is displayed.

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.

Part V

Appendixes

Part V provides reference information. It contains the following chapters:

- [Appendix A, "Error Messages for Oracle Big Data Appliance"](#)
- [Appendix B, "Replacement Units"](#)
- [Appendix C, "In-Rack Cabling Tables for Oracle Big Data Appliance X3-2"](#)
- [Appendix D, "In-Rack Cabling Tables for Oracle Big Data Appliance"](#)
- [Appendix E, "Multitrack Cabling Tables"](#)

Error Messages for Oracle Big Data Appliance

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

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

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

Information is also written to the system logs.

DISKALRT-02001 System OS hard disk failure

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

Action: Replace the physical disk as soon as possible. See "[Replacing a Server Disk](#)" on page 10-6 and follow the instructions for configuring an operating system disk.

DISKALRT-02002 System OS hard disk predictive failure

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

Action: Replace the physical disk as soon as possible. See "[Replacing a Server Disk](#)" on page 10-6 and follow the instructions for configuring an operating system disk.

DISKALRT-02003 System OS hard disk poor performance

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

Action: Plan to replace the disk. If multiple disk disks in a server are identified as poor performers, then check for problems with the cable connections, the controllers, and the disk backplane before replacing the disks. To replace a disk, see "[Replacing a Server Disk](#)" on page 10-6 and follow the instructions for configuring an operating system disk.

DISKALRT-02004 Data hard disk failure

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

Action: Replace the physical disk as soon as possible. See "[Replacing a Server Disk](#)" on page 10-6 and follow the instructions for configuring a disk for either HDFS or Oracle NoSQL Database.

DISKALRT-02005 Data hard disk predictive failure

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

Action: Replace the physical disk as soon as possible. See ["Replacing a Server Disk"](#) on page 10-6 and follow the instructions for configuring a disk for either HDFS or Oracle NoSQL Database.

DISKALRT-02006 Data hard disk poor performance

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

Action: Replace the physical disk as soon as possible. See ["Replacing a Server Disk"](#) on page 10-6 and follow the instructions 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.

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](#)
- [Disk Controller Batteries](#)
- [LED Status Descriptions](#)
- [Parts for Oracle Big Data Appliance Servers](#)
- [Parts for Sun Datacenter InfiniBand Switch 36](#)
- [Parts for the Cisco Ethernet Switch](#)
- [Parts for Power Distribution Units](#)
- [Cables and Cable Accessories for an Oracle Big Data Appliance Rack](#)

See Also: ["Spare Parts Kit Components"](#) on page 1-3

Understanding Repair Categories

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

Table B-1 Oracle Big Data Appliance Repair Categories

| Category | Description | Scenario Examples |
|----------------------------|--|---|
| Hot Swap (HS) | The repair part is hot swappable and can be replaced without shutting down the host system. Procedures may be needed before and after replacement to protect the data. | <ul style="list-style-type: none"> ■ Disks ■ Fans ■ Power supplies |
| Infrastructure Repair (IR) | Repair of a connectivity component within an Oracle Big Data Appliance rack. No downtime of the rack is required; however, individual components may require downtime. | <ul style="list-style-type: none"> ■ External cables ■ InfiniBand switch ■ Ethernet switch |

Table B-1 (Cont.) Oracle Big Data Appliance Repair Categories

| Category | Description | Scenario Examples |
|----------------------------|---|---|
| BDA Server Offline (BDA-O) | <p>Repair of the part requires shutting down one Big Data Appliance server. No downtime of the rack is required; however, individual servers may require downtime and temporary removal from the cluster. Rebalancing the workload and data duplication may temporarily affect performance.</p> <p>If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 10-2 for additional information.</p> <p>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.</p> | <ul style="list-style-type: none"> ■ System boards ■ PCIe cards ■ Memory ■ Processors |

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.

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](#)
- [Sun InfiniBand Switch LEDs](#)

Sun Server LEDs

Table B–2 identifies some LED status indicators on a Sun Server X3-2L that can indicate a fault.

See Also: "About Server and Component LEDs" in the *Sun Server X3-2L Service Manual* at

http://docs.oracle.com/cd/E23393_01/html/E27229/z40003ec1405942.html

Table B–2 LED Status Descriptions

| Component | LED Status |
|----------------|---|
| General Status | <p>The following LEDs are located on the front and back panels of each server:</p> <p>Service Required (amber)</p> <ul style="list-style-type: none"> ■ Off indicates normal operation. ■ Steady on indicates a fault in the server. <p>Power/OK (green)</p> <ul style="list-style-type: none"> ■ 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. ■ Off indicates no AC power. |
| Drives | <p>The following LEDs are located on the storage and boot drives:</p> <p>Service Required (amber)</p> <ul style="list-style-type: none"> ■ Off indicates normal operation. ■ Steady on indicates a fault in the drive. <p>Activity (green)</p> <ul style="list-style-type: none"> ■ 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. |
| Fans | <p>Each fan module has one bicolored LED.</p> <p>Fan Status (amber/green)</p> <ul style="list-style-type: none"> ■ 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. |

Table B–2 (Cont.) LED Status Descriptions

| Component | LED Status |
|--------------|---|
| Power supply | <p>The following LEDs are on each power supply. You can see them from the back of the server.</p> <p>Service Required (amber)</p> <ul style="list-style-type: none"> ■ 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. <p>AC OK/DC OK (green)</p> <ul style="list-style-type: none"> ■ 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. <p>AC OK (green)</p> <ul style="list-style-type: none"> ■ Steady on indicates normal operation. ■ Off indicates no AC power. |

Sun InfiniBand Switch LEDs

Table B–3 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 Datacenter InfiniBand Switch 36 LED Status Descriptions

| Component | LED Status |
|-----------|---|
| Chassis | <p>The following LEDs are located on the left side of the rear panel.</p> <p>Attention (amber)</p> <ul style="list-style-type: none"> ■ Steady on indicates a fault. ■ Off indicates normal operation. ■ Flashing indicates no function. <p>OK (green)</p> <ul style="list-style-type: none"> ■ Steady on indicates normal operation. ■ Off indicates that the switch is off or initializing. ■ Flashing indicates no function. |

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

| Component | LED Status |
|--------------------------|---|
| Link status | <p>The link status LEDs are located by the data cable connectors on the rear panel.</p> <p>Link (green)</p> <ul style="list-style-type: none"> ■ 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 | <p>The following LEDs are located on the network management connector on the rear panel.</p> <p>Link Speed (amber/green):</p> <ul style="list-style-type: none"> ■ Green steady on indicates a 1000BASE-T link. ■ Amber steady on indicates a 100BASE-T link. ■ Off indicates no link. ■ Flashing indicates no function. <p>Activity (green):</p> <ul style="list-style-type: none"> ■ Flashing indicates packet activity. ■ On indicates no function. ■ Off indicates no activity. |
| Power supply | <p>The following LEDs are located on the power supply at the front of the chassis.</p> <p>Attention (amber)</p> <ul style="list-style-type: none"> ■ Steady on indicates a fault, and 12 VDC is shut down. ■ Off indicates normal operation. ■ Flashing indicates no function. <p>OK (green)</p> <ul style="list-style-type: none"> ■ Steady on indicates 12 VDC. ■ Off indicates no DC voltage. ■ Flashing indicates no function. <p>AC LED is green</p> <ul style="list-style-type: none"> ■ Steady on indicates AC power is supplied. ■ Off indicates that AC power is not present. ■ Flashing indicates no function. |
| Fan | <p>The following LED is located on the air inlet on the front of the chassis.</p> <p>FT</p> <ul style="list-style-type: none"> ■ On indicates a problem with the fan. ■ Off indicates normal operation |

Parts for Oracle Big Data Appliance Servers

An Oracle Big Data Appliance rack contains either Sun Server X3-2L or Sun Fire X4270 M2 servers.

Note: Oracle Big Data Appliance ships with two spare drives. If a drive is faulty, then the customer can replace the drive before an Oracle service representative arrives if circumstances warrant immediate action.

Table B-4 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-4 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 |
| 371-5002 | 4 GB USB 2.0 flash drive, RoHS:Y | BDA-O | See "Servicing the Internal USB Flash Drives" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceiicfbd.html#scrolltoc |
| 371-4985 | 6 Gigabit SAS RAID PCI battery module, RoHS:Y (LION), BBU-08 | BDA-O | <ol style="list-style-type: none"> 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 | <ol style="list-style-type: none"> 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 |

Table B-4 (Cont.) Replacement Parts for a Sun Server X3-2L Server

| Part Number | Description | Repair Category | Repair Procedure |
|-------------|---|-----------------|--|
| 375-3696 | InfiniBand (CX2) dual port CX2 4x QDR PCI-e | BDA-O | <ol style="list-style-type: none"> Power down the server gracefully. See "Powering Down the Server" at http://docs.oracle.com/cd/E23393_01/html/E27229/z400061b1011735.html See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000f91037409.html |
| 375-3701 | 6 Gbps SAS RAID disk PCIe HBA, RoHS:Y | BDA-O | <ol style="list-style-type: none"> Power down the server gracefully. See "Powering Down the Server" at http://docs.oracle.com/cd/E23393_01/html/E27229/z400061b1011735.html See "Servicing PCIe Cards" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000f91037409.html <p>Note: Do not run any controller commands when replacing the disk.</p> |
| 7010036 | DR 3 TB, 3.5-inch, SAS/7200, CORAL | HS | <p>See "Servicing Storage Drives and Rear Drives" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40000091011460.html</p> <p>Do not run any controller commands when replacing the disk.</p> <p>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</p> |
| 7014391 | SAS expander assembly | BDA-O | <p>See "Servicing the SAS Expander Module" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceieidja.html</p> |
| 7020486 | DIMM, 8 GB, DDR3, 1600 MHz, 1.3V | BDA-O | <p>See "Servicing the DIMMs" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceicjagi.html#scrolltoc</p> |
| 7018704 | DIMM, 16 GB, DDR3, 1600 MHz, 1.35V | BDA-O | <p>See "Servicing the DIMMs" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceicjagi.html#scrolltoc</p> |
| 7041013 | DIMM, 32 GB, DDR3, RDIMM-1066, QUAD-RANK | BDA-O | <p>See "Servicing the DIMMs" at http://docs.oracle.com/cd/E23393_01/html/E27229/ceicjagi.html#scrolltoc</p> |
| 7020774 | Pregreased CPU heatsink | BDA-O | <p>See "Servicing Processors" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40001d31037480.html#scrolltoc</p> |

Table B–4 (Cont.) Replacement Parts for a Sun Server X3-2L Server

| Part Number | Description | Repair Category | Repair Procedure |
|-------------|--|-----------------|---|
| 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: <ul style="list-style-type: none"> ■ Disk drive flex cable ■ Dual backplane power cable ■ Single backplane power cable ■ Disk data cable, iPass+ HD to iPass+ HD, 600 mm | BDA-O | <ol style="list-style-type: none"> 1. Shut down the rack. See "Powering On and Off Oracle Big Data Appliance" on page 10-2. 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 | <ol style="list-style-type: none"> 1. Shut down the rack. See "Powering On and Off Oracle Big Data Appliance" on page 10-2. 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 | <ol style="list-style-type: none"> 1. Shut down the rack. See "Powering On and Off Oracle Big Data Appliance" on page 10-2. 2. "Servicing the Motherboard Assembly" at http://docs.oracle.com/cd/E23393_01/html/E27229/z40001d31037512.html |

Table B–5 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-5 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 | <ol style="list-style-type: none"> 1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 10-2. 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 |
| 371-4743 | 4 GB USB 2.0 flash drive, RoHS:Y | BDA-O | Not applicable |
| 371-4966 | LV DIMM, 8 GB, DDR3, 1RX4, 1333 MHz | BDA-O | See Section 4.2, "Servicing Memory Modules (DIMMs)" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_71526 |
| 371-4982 | 6 Gigabit SAS RAID PCI battery module, RoHS:Y (LION), BBU-08 | BDA-O | <ol style="list-style-type: none"> 1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 10-2. 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 | <ol style="list-style-type: none"> 1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 10-2. 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 |

Table B–5 (Cont.) Replacement Parts for Sun Fire X4270 M2 Servers

| Part Number | Description | Repair Category | Repair Procedure |
|-------------|--|-----------------|---|
| 375-3701 | 6 Gbps SAS RAID disk PCIe HBA, RoHS:Y | BDA-O | <ol style="list-style-type: none"> 1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 10-2. 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 <p>Note: Do not run any controller commands in the service manual when replacing the disk.</p> |
| 530-3927 | Cable, PDB, MB, 1U+2U, ribbon | BDA-O | <p>See Section 5.10, "Servicing Cables " in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_58217</p> |
| 530-4228 | Cable, mini-SAS, 36 POS LONG, LY, RoHS:Y | BDA-O | <p>See Section 5.10, "Servicing Cables" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_78806</p> |
| 541-2075 | Bus_bar, power, 1U+2U, RoHS:Y | BDA-O | <p>See Section 5.6, "Servicing the Power Distribution Board" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_85936</p> |
| 541-2884 | Assembly, riser, X8, X8, 2U, PCI-E, RoHS:Y | BDA-O | <p>See Section 4.4, "Servicing PCIe Risers" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_53252</p> |
| 541-3310 | Xeon Heatsink with grease, RoHS:Y, 2U | BDA-O | <p>See Section 4.7, "Servicing 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</p> |
| 541-3356 | <p>x8/x8 switched PCIe riser assembly, RoHS:Y with the following:</p> <ul style="list-style-type: none"> ■ x8/x8 switch PCIe riser ■ bracket | BDA-O | <ol style="list-style-type: none"> 1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 10-2. 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-5 (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 | <ol style="list-style-type: none"> 1. If the system is currently running, then shut it down gracefully. See "Powering On and Off Oracle Big Data Appliance" on page 10-2. 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 |
| 541-4306 | Assembly, connector board bracket, Hyde | BDA-O | See Section 5.7, "Servicing the Connector Board" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_pgfid-1296581 |
| 541-4363 | 12-slot disk backplane, RoHS:Y | BDA-O | See Section 5.4, "Servicing the Disk Backplane" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_71557 |

Table B-5 (Cont.) Replacement Parts for Sun Fire X4270 M2 Servers

| Part Number | Description | Repair Category | Repair Procedure |
|-------------|--|-----------------|---|
| 541-4371 | Indicator, left, Hyde 12/24, RoHS:Y | BDA-O | See Section 5.5, "Servicing the Front Panel LED Modules" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/infrastructure_x4270_M2.html#50438167_71563 |
| 560-2952 | Cable kit, Hyde, RoHS:Y with the following: <ul style="list-style-type: none"> ■ Disk backplane power cable ■ 24-slot DBP power cable ■ 12/24 slot disk backplane signal cable ■ Fan signal cable ■ Fan power cable | BDA-O | Refer to the entries for the individual items in the kit. |
| 7010036 | FRU DR 3 TB 3.5-inch SAS/7200, CORAL | -- | See Section 3.2, "Servicing Storage Drives and Boot Drives" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/hotswap.html#50438178_61628 Note: Do not run any controller commands in the service manual when replacing the disk. |
| 7012433 | 3.06 GHz processor | BDA-O | See Section 4.7, "Servicing CPUs" in <i>Sun Fire X4270 M2 Server Service Manual</i> at http://docs.oracle.com/cd/E19245-01/821-0488-12/motherboard.html#50438179_23425 |

Parts for Sun Datacenter InfiniBand Switch 36

Table B-6 lists the replaceable parts for the Sun Datacenter InfiniBand Switch 36, which are covered under the Oracle Premier Support for Systems warranty. Unless otherwise noted, the repair procedures reference the *Sun Datacenter InfiniBand Switch 36 Service Manual* at

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

See "Cables and Cable Accessories for an Oracle Big Data Appliance Rack" on page B-15 for information about InfiniBand cables.

Table B–6 Replacement Parts for the Sun Datacenter InfiniBand Switch 36

| Part Number | Description | Repair Category | Repair Procedure |
|-------------|--|-----------------|--|
| 300-2299 | A247A 760-watt power supply, RoHS:Y | HS | See "Servicing Power Supplies" at http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f49112.html |
| 350-1312 | Fan module for Sun Datacenter InfiniBand Switch 36, RoHS:Y | HS | See "Servicing Fans" at http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f59112.html |
| 371-2210 | CR2032 3V battery, RoHS:Y | IR | See "Servicing the Battery" at http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f79112.html |
| 541-3495 | Sun Datacenter InfiniBand Switch 36, RoHS:YL | IR | <ol style="list-style-type: none"> 1. See "Replacing a Failed InfiniBand Switch" on page 10-26. 2. See "Installing the Switch" in the <i>Sun Datacenter InfiniBand Switch 36 Installation Guide</i> at http://docs.oracle.com/cd/E26698_01/html/E26434/cggbhbf.html |
| 594-6603 | 10 Gbps QSFP short wavelength transceiver | IR | See "Servicing Data Cables" at http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f69112.html |

Parts for Sun Network QDR InfiniBand Gateway Switches

Table B–7 lists the replaceable parts for the Sun Network QDR InfiniBand Gateway switch, which are covered under the Oracle Premier Support for Systems warranty. Unless otherwise noted, the repair procedures refer to the *Sun Network QDR InfiniBand Gateway Switch Service Manual* at

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

See "Cables and Cable Accessories for an Oracle Big Data Appliance Rack" on page B-15 for information about InfiniBand cables.

Table B–7 Replacement Parts for the Sun Network QDR InfiniBand Gateway Switch

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

Table B-7 (Cont.) Replacement Parts for the Sun Network QDR InfiniBand Gateway Switch

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

Parts for the Cisco Ethernet Switch

Table B-8 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-8 Replacement Parts for Cisco Catalyst Ethernet Switch

| 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> . |
| 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> . |

Parts for Power Distribution Units

Table B-9 lists the replacement parts for the power distribution units (PDUs), which are covered under the Oracle Premier Support for Systems warranty. The repair procedures references the *Sun Rack II Power Distribution Units User's Guide* at

<http://docs.oracle.com/cd/E19844-01/E23443-01/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–9 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/E23443-01/z40000121011643.html |
| 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/E23443-01/z40000121011643.html |
| 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/E23443-01/z40000121011643.html |
| 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/E23443-01/z40000121011643.html |

Cables and Cable Accessories for an Oracle Big Data Appliance Rack

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

See Also:

- *Sun Network QDR InfiniBand Gateway Switch Service Manual* at
http://docs.oracle.com/cd/E26699_01/html/E26706/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–10 Replacement Cables for an Oracle Big Data Appliance Rack

| Part Number | Description | Repair Category | Repair |
|-------------|--|-----------------|--|
| 350-1287 | Slim rail, CMA, 1U-2U for Sun servers | BDA-O | <ol style="list-style-type: none"> Power down the server gracefully. See "Powering Down the Server" at http://docs.oracle.com/cd/E23393_01/html/E27229/z400061b1011735.html 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: <ul style="list-style-type: none"> USB to DB 9-pin M serial cable DB 9-pin F to DB 9-pin F null modem cable | HS | See "Understanding Cabling" in the <i>Sun Datacenter InfiniBand Switch 36 Installation Guide</i> at http://docs.oracle.com/cd/E26698_01/html/E26434/cggfhjcc.html |
| 530-4567 | 4x QSFP copper QDR InfiniBand cable, 2 m, RoHS:Y | HS | <ol style="list-style-type: none"> See "Maintaining the InfiniBand Network" on page 10-25. See "Understanding Cabling" in the <i>Sun Datacenter InfiniBand Switch 36 Installation Guide</i> at http://docs.oracle.com/cd/E26698_01/html/E26434/cggfhjcc.html <p>Note: Redundancy is lost during cable replacement.</p> |
| 530-4445 | 4x QSFP copper QDR InfiniBand cable, 3 m, RoHS:Y | HS | <ol style="list-style-type: none"> See "Maintaining the InfiniBand Network" on page 10-25. 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 | <ol style="list-style-type: none"> See "Maintaining the InfiniBand Network" on page 10-25. 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 |
| 530-4432 | 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> . |

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

The tables in this appendix show the in-rack cable connections for Sun Server X3-2L-based Oracle Big Data Appliance racks. This appendix contains the following sections:

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

Cable Color Coding

The cables used in Oracle Big Data Appliance are color-coded as follows:

- **Black:** InfiniBand cables or AC power jumper cables
- **Red:** Integrated Lights Out Management (Oracle ILOM) Ethernet management cables
- **Blue:** Gigabit Ethernet management (eth0) cables

Oracle Big Data Appliance Rack Layout

The following figures show the front and rear layouts for three configurations of Oracle Big Data Appliance X3-2: full rack, starter rack, and starter rack with one in-rack expansion kit.

[Figure C-1](#) shows the layout of a full rack for Oracle Big Data Appliance X3-2.

Figure C-1 Oracle Big Data Appliance X3-2 Full Rack Layout

| Rear View | | | | | Front View | | | | |
|-----------|---------------|--------------------------------------|-------|---------------|------------|-------|---------------------------|-------|----|
| U# | PDU B | Big Data Appliance | PDU A | U# | U# | PDU B | Big Data Appliance | PDU A | U# |
| 42 | INP0 | 1U Solid Filler | INP2 | 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 | BDA Server 10 | 32 | 32 | BDA Server 10 | 32 | | | | |
| 31 | BDA Server 9 | 31 | 31 | BDA Server 9 | 31 | | | | |
| 30 | BDA Server 8 | 30 | 30 | BDA Server 8 | 30 | | | | |
| 29 | BDA Server 7 | 29 | 29 | BDA Server 7 | 29 | | | | |
| 28 | BDA Server 6 | 28 | 28 | BDA Server 6 | 28 | | | | |
| 27 | BDA Server 5 | 27 | 27 | BDA Server 5 | 27 | | | | |
| 26 | BDA Server 4 | 26 | 26 | BDA Server 4 | 26 | | | | |
| 25 | BDA Server 3 | 25 | 25 | BDA Server 3 | 25 | | | | |
| 24 | BDA Server 2 | 24 | 24 | BDA Server 2 | 24 | | | | |
| 23 | BDA Server 1 | 23 | 23 | BDA Server 1 | 23 | | | | |
| 22 | INP1 | Sun Network QDR Infiniband GW Switch | INP1 | 22 | 22 | INP1 | 1U Vented Filler | INP1 | 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 | BDA Server 9 | 19 | 19 | BDA Server 9 | 19 | | | | |
| 18 | BDA Server 8 | 18 | 18 | BDA Server 8 | 18 | | | | |
| 17 | BDA Server 7 | 17 | 17 | BDA Server 7 | 17 | | | | |
| 16 | BDA Server 6 | 16 | 16 | BDA Server 6 | 16 | | | | |
| 15 | BDA Server 5 | 15 | 15 | BDA Server 5 | 15 | | | | |
| 14 | BDA Server 4 | 14 | 14 | BDA Server 4 | 14 | | | | |
| 13 | BDA Server 3 | 13 | 13 | BDA Server 3 | 13 | | | | |
| 12 | BDA Server 2 | 12 | 12 | BDA Server 2 | 12 | | | | |
| 11 | BDA Server 1 | 11 | 11 | BDA Server 1 | 11 | | | | |
| 10 | INP2 | BDA Server 4 | INP0 | 10 | 10 | INP2 | BDA Server 3 | INP0 | 10 |
| 9 | | BDA Server 3 | | 9 | 9 | | BDA Server 2 | | 9 |
| 8 | | BDA Server 2 | | 8 | 8 | | BDA Server 1 | | 8 |
| 7 | | BDA Server 1 | | 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 | | | | |

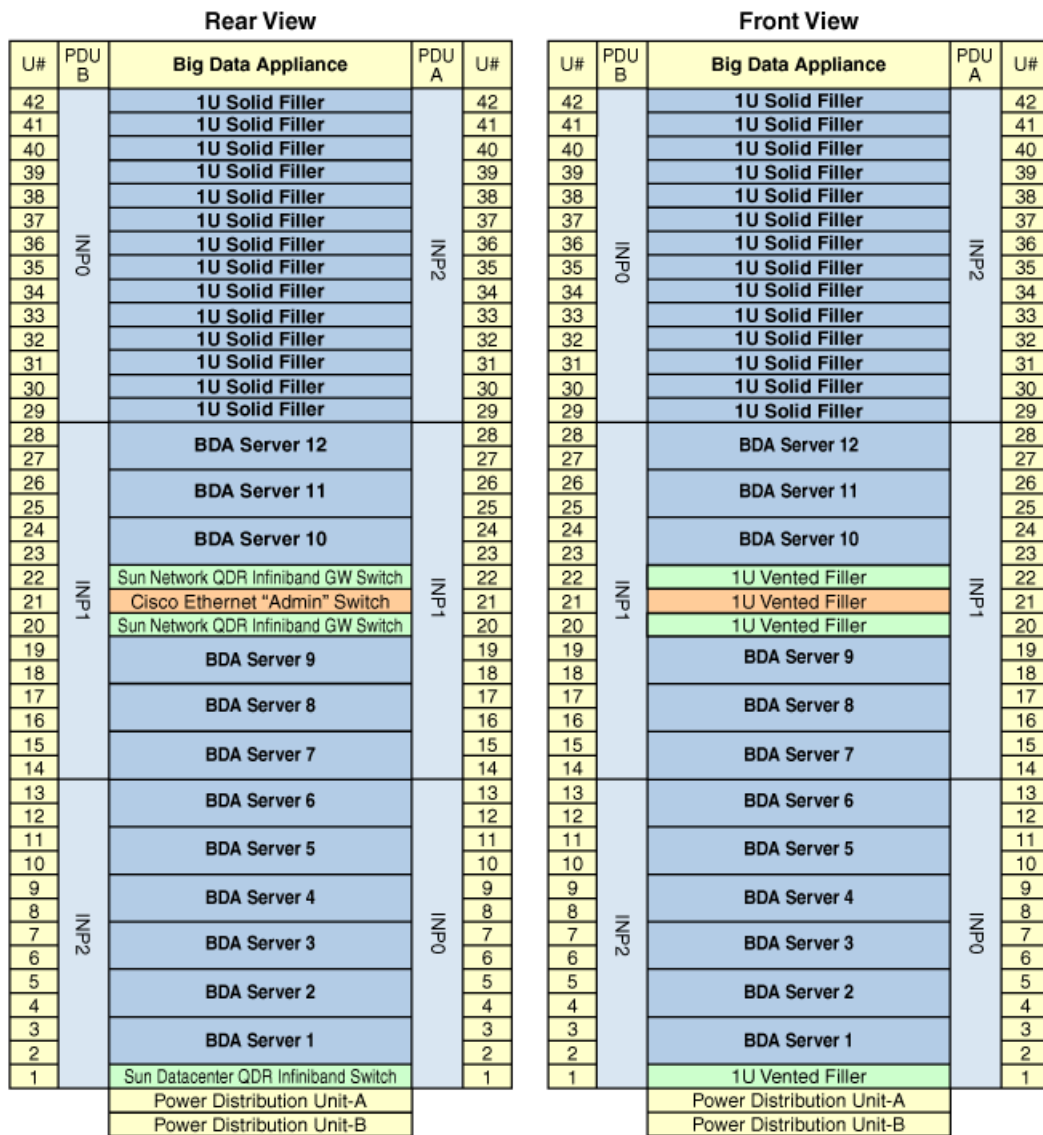
Figure C-2 shows the layout of a starter rack for Oracle Big Data Appliance X3-2.

Figure C-2 Oracle Big Data Appliance X3-2 Starter Rack Layout

| Rear View | | | | | Front View | | | | | | | | |
|-----------|---------------------------|--------------------------------------|-------|-----------------|------------|------------------|---------------------------|-------|----|------------------|----|----|----|
| U# | PDU B | Big Data Appliance | PDU A | U# | U# | PDU B | Big Data Appliance | PDU A | U# | | | | |
| 42 | INP0 | 1U Solid Filler | INP2 | 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 | 1U Solid Filler | 30 | 30 | 1U Solid Filler | 30 | | | | | | | | |
| 29 | 1U Solid Filler | 29 | 29 | 29 | 29 | 1U Solid Filler | 29 | 29 | 29 | | | | |
| 28 | INP1 | 1U Solid Filler | INP1 | 28 | 28 | INP1 | 1U Solid Filler | INP1 | 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 | | 22 | | 22 | 1U Vented Filler | 22 | 22 | 22 |
| 21 | | Cisco Ethernet "Admin" Switch | | 21 | 21 | | 21 | | 21 | 1U Vented Filler | 21 | 21 | 21 |
| 20 | | Sun Network QDR Infiniband GW Switch | | 20 | 20 | | 20 | | 20 | 1U Vented Filler | 20 | 20 | 20 |
| 19 | | 1U Solid Filler | | 19 | 19 | | 19 | | 19 | 1U Solid Filler | 19 | 19 | 19 |
| 18 | 1U Solid Filler | 18 | 18 | 18 | 18 | 1U Solid Filler | 18 | 18 | 18 | | | | |
| 17 | 1U Solid Filler | 17 | 17 | 17 | 17 | 1U Solid Filler | 17 | 17 | 17 | | | | |
| 16 | 1U Solid Filler | 16 | 16 | 16 | 16 | 1U Solid Filler | 16 | 16 | 16 | | | | |
| 15 | 1U Solid Filler | 15 | 15 | 15 | 15 | 1U Solid Filler | 15 | 15 | 15 | | | | |
| 14 | 1U Solid Filler | 14 | 14 | 14 | 14 | 1U Solid Filler | 14 | 14 | 14 | | | | |
| 13 | INP2 | BDA Server 6 | INP0 | 13 | 13 | INP2 | BDA Server 6 | INP0 | 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 | INP0 | BDA Server 3 | INP0 | 7 | 7 | INP0 | BDA Server 3 | INP0 | 7 | | | | |
| 6 | | BDA Server 2 | | 6 | 6 | | BDA Server 2 | | 6 | | | | |
| 5 | | BDA Server 1 | | 5 | 5 | | BDA Server 1 | | 5 | | | | |
| 4 | | 1U Vented Filler | | 4 | 4 | | 1U Vented Filler | | 4 | | | | |
| 3 | | Power Distribution Unit-A | | 3 | 3 | | Power Distribution Unit-A | | 3 | | | | |
| 2 | | Power Distribution Unit-B | | 2 | 2 | | Power Distribution Unit-B | | 2 | | | | |
| 1 | Power Distribution Unit-B | 1 | 1 | 1 | 1 | 1U Vented Filler | 1 | 1 | 1 | | | | |

Figure C-3 shows the layout of an Oracle Big Data Appliance X3-2 starter rack with one in-rack expansion kit.

Figure C-3 Oracle Big Data Appliance X3-2 Starter Rack With In-Rack Expansion



Administrative Gigabit Ethernet Cable Connections

Table C-1 shows the cable connections for the administrative Gigabit Ethernet network (Net 0). The Ethernet switch is located in rack unit 21. The cables are 10 feet long and black.

Table C-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 |

Table C-1 (Cont.) Administrative Gigabit Ethernet Cabling

| From Rack Unit ¹ | Component | To Ethernet Port |
|-----------------------------|---------------|------------------|
| 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 U*n* is the unit location in the rack, where *n* is the number.

Oracle Integrated Lights Out Manager Cable Connections

Table C-2 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 | 4 | 10 feet | Red |
| U37 | Server 17 | 6 | 10 feet | Red |
| U35 | Server 16 | 8 | 10 feet | Red |
| U33 | Server 15 | 10 | 10 feet | Red |
| U31 | Server 14 | 12 | 10 feet | Red |
| U29 | Server 13 | 14 | 10 feet | Red |
| U27 | Server 12 | 18 | 10 feet | Red |
| U25 | Server 11 | 22 | 10 feet | Red |
| U23 | Server 10 | 26 | 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 |

Table C-2 (Cont.) Oracle ILOM Cabling

| From Rack Unit ¹ | Component | Ethernet Port | Cable Length | Cable Color |
|-----------------------------|-----------|---------------|--------------|-------------|
| U10 | Server 5 | 36 | 10 feet | Red |
| U8 | Server 4 | 38 | 10 feet | Red |
| U6 | Server 3 | 40 | 10 feet | Red |
| U4 | Server 2 | 42 | 10 feet | Red |
| U2 | Server 1 | 44 | 10 feet | Red |
| PDU-A | PDU-A | 15 | 1 meter | White |
| PDU-B | PDU-B | 20 | 1 meter | White |
| Not applicable | Service | 48 | 10 feet | Blue |

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

Single-Phase Power Distribution Unit Cable Connections

Table C-3 shows the connections for single-phase cabling from each power distribution unit (PDU) to the power supplies in the rack. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table 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 |
| 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 |

Table C-3 (Cont.) Single-Phase PDU Cabling

| Rack Unit ¹ | PDU-A/PS-00 | PDU-B/PS-01 | Cable Length |
|------------------------|-------------|-------------|--------------|
| 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 U n is the unit location in the rack, where n is the number.

Three-Phase Power Distribution Unit Cable Connections

Table C-4 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 |
| U24 | G3-1 | G0-5 | 2 meters |
| U23 | Not applicable | G5-0 | Included |
| U22 | G2-5 | G5-1 | 1 meter |
| U21 | G3-0 | G0-6 | 2 meters |
| U20 | G2-4 | G5-2 | 2 meters |
| U18 | G2-2 | G5-4 | 2 meters |
| U16 | G1-6 | G4-0 | 2 meters |
| U14 | G2-0 | G5-6 | 2 meters |
| U12 | G1-4 | G4-2 | 2 meters |
| U10 | G1-2 | G4-4 | 2 meters |
| U8 | G1-0 | G4-6 | 2 meters |
| U6 | G0-6 | G3-0 | 2 meters |
| U4 | G0-4 | G3-2 | 2 meters |
| U2 | G0-2 | G3-4 | 2 meters |
| U1 | G0-0 | G3-6 | 2 meters |

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

In-Rack InfiniBand Switch-to-Switch Cable Connections

Table C-5 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 | 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 C-4 identifies the locations of the ports on a Sun Network QDR InfiniBand Gateway switch.

Figure C-4 Sun Network QDR InfiniBand Gateway Ports



In-Rack InfiniBand Switch-to-Server Cable Connections

Table C-6 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 |

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

| From InfiniBand Switch Rack Unit ¹ | Port | To Rack Unit ¹ | Port ² | Cable Length |
|---|------|---------------------------|-------------------|--------------|
| 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 | 2 meters |
| U22 | 4A | U23 | PCIe 3 P1 | 2 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 | 2 meters |
| U20 | 4A | U23 | PCIe 3 P2 | 2 meters |
| U20 | 4B | U18 | PCIe 3 P1 | 2 meters |
| U20 | 12A | U16 | PCIe 3 P1 | 2 meters |
| U20 | 12B | U14 | PCIe 3 P1 | 3 meters |
| U20 | 13A | U12 | PCIe 3 P1 | 3 meters |
| U20 | 13B | U10 | PCIe 3 P1 | 3 meters |
| U20 | 14A | U8 | PCIe 3 P1 | 3 meters |
| U20 | 14B | U6 | PCIe 3 P1 | 3 meters |
| U20 | 15A | U4 | PCIe 3 P1 | 3 meters |
| U20 | 15B | U2 | PCIe 3 P1 | 3 meters |
| U22 | 4B | U18 | PCIe 3 P2 | 2 meters |
| U22 | 12A | U16 | PCIe 3 P2 | 2 meters |
| U22 | 12B | U14 | PCIe 3 P2 | 3 meters |
| U22 | 13A | U12 | PCIe 3 P2 | 3 meters |
| U22 | 13B | U10 | PCIe 3 P2 | 3 meters |
| U22 | 14A | U8 | PCIe 3 P2 | 3 meters |
| U22 | 14B | U6 | PCIe 3 P2 | 3 meters |
| U22 | 15A | U4 | PCIe 3 P2 | 3 meters |
| U22 | 15B | U2 | PCIe 3 P2 | 3 meters |

¹ Un is the unit location in the rack, where n is the number.² Pn is the InfiniBand port, where n is the port number.

In-Rack Cabling Tables for Oracle Big Data Appliance

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

Cable Color Coding

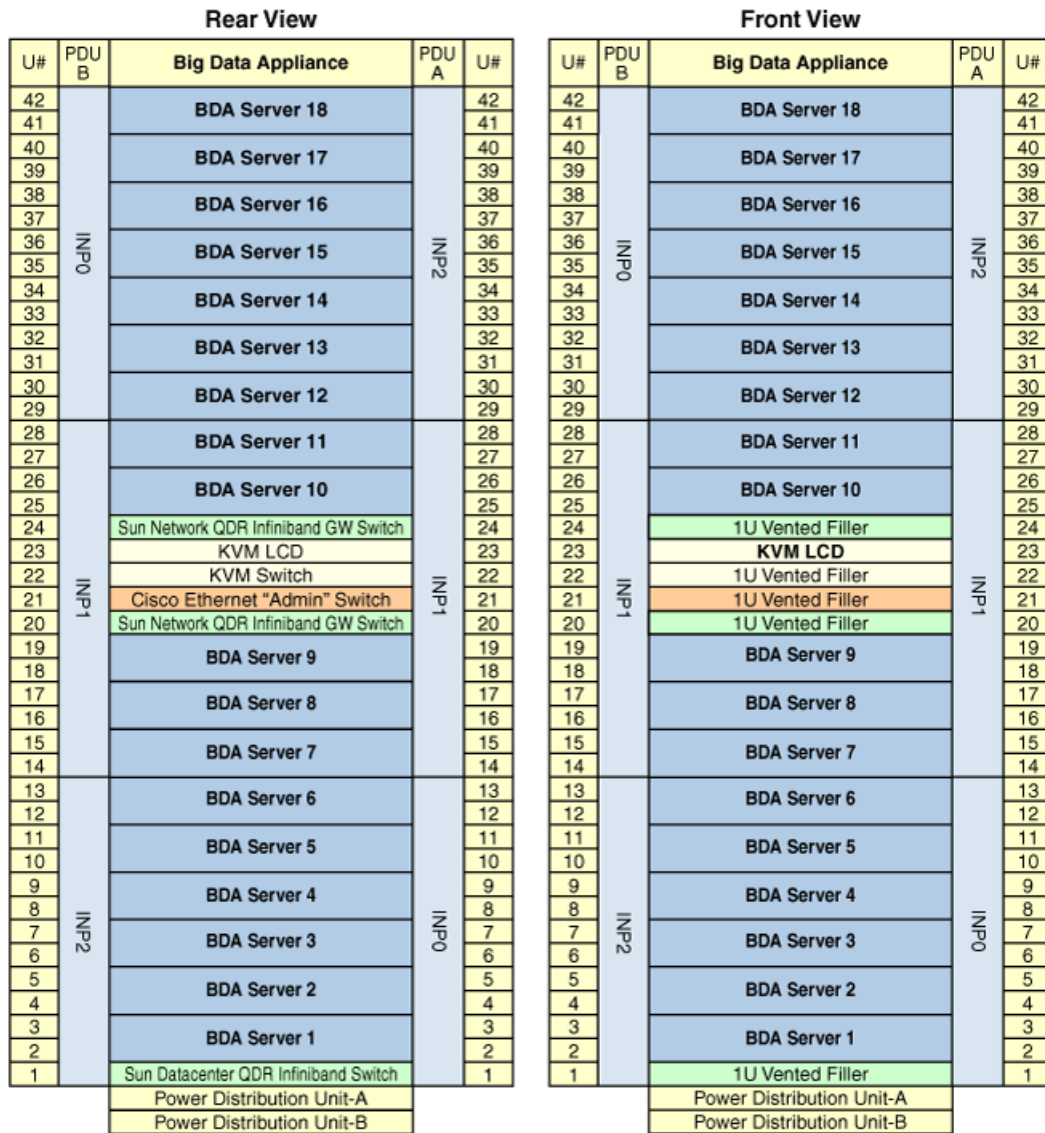
The cables used in Oracle Big Data Appliance are color-coded as follows:

- Black: InfiniBand cables or AC power jumper cables
- Red: Integrated Lights Out Management (Oracle ILOM) Ethernet management cables
- Blue: Gigabit Ethernet management (eth0) cables
- Orange: KVM switch to dongle cables

Oracle Big Data Appliance Rack Layout

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

Figure D-1 Rack Layout of Oracle Big Data Appliance



KVM Network Cable Connections

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

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

Table D-1 KVM Cabling

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

Table D-1 (Cont.) KVM Cabling

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

Administrative Gigabit Ethernet Cable Connections

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

Table D-2 Administrative Gigabit Ethernet Cabling

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

Table D–2 (Cont.) Administrative Gigabit Ethernet Cabling

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

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

Oracle Integrated Lights Out Manager Cable Connections

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

Table D–3 Oracle ILOM Cabling

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

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

Single-Phase Power Distribution Unit Cable Connections

Table D-4 shows the connections for single-phase cabling from each power distribution unit (PDU) to the power supplies in the rack. The cables terminate at PDU-A on the left and are routed to the right to enter the cable management arm (CMA). The cables are bundled in groups of four.

Table D-4 Single-Phase PDU Cabling

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

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

Three-Phase Power Distribution Unit Cable Connections

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

Table D-5 Three-Phase PDU Cabling

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

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

In-Rack InfiniBand Switch-to-Switch Cable Connections

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

Table D-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 |

Table D-6 (Cont.) InfiniBand Switch-to-Switch Cabling

| From InfiniBand Switch Rack Unit ¹ | Port | To InfiniBand Switch Rack Unit ¹ | Port | Cable Length | Description |
|---|------|---|------|--------------|---------------|
| U20 | 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 U n is the unit location in the rack, where n is the number.

Figure D-2 Sun Network QDR InfiniBand Gateway Ports



In-Rack InfiniBand Switch-to-Server Cable Connections

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

Table D-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 |

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

| From InfiniBand Switch Rack Unit¹ | Port | To Rack Unit¹ | Port² | Cable Length |
|---|-------------|---------------------------------|-------------------------|---------------------|
| U20 | 2B | U31 | PCIe 3 P2 | 3 meters |
| U20 | 3A | U29 | PCIe 3 P2 | 3 meters |
| U20 | 3B | U27 | PCIe 3 P2 | 2 meters |
| U20 | 4A | U25 | PCIe 3 P2 | 2 meters |
| U20 | 4B | U18 | PCIe 3 P1 | 2 meters |
| U20 | 12A | U16 | PCIe 3 P1 | 2 meters |
| U20 | 12B | U14 | PCIe 3 P1 | 3 meters |
| U20 | 13A | U12 | PCIe 3 P1 | 3 meters |
| U20 | 13B | U10 | PCIe 3 P1 | 3 meters |
| U20 | 14A | U8 | PCIe 3 P1 | 3 meters |
| U20 | 14B | U6 | PCIe 3 P1 | 3 meters |
| U20 | 15A | U4 | PCIe 3 P1 | 3 meters |
| U20 | 15B | U2 | PCIe 3 P1 | 3 meters |
| U24 | 4B | U18 | PCIe 3 P2 | 2 meters |
| U24 | 12A | U16 | PCIe 3 P2 | 2 meters |
| U24 | 12B | U14 | PCIe 3 P2 | 3 meters |
| U24 | 13A | U12 | PCIe 3 P2 | 3 meters |
| U24 | 13B | U10 | PCIe 3 P2 | 3 meters |
| U24 | 14A | U8 | PCIe 3 P2 | 3 meters |
| U24 | 14B | U6 | PCIe 3 P2 | 3 meters |
| U24 | 15A | U4 | PCIe 3 P2 | 3 meters |
| U24 | 15B | U2 | PCIe 3 P2 | 3 meters |

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

² Pn is the InfiniBand port, where n is the port number.

Multirack Cabling Tables

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

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

Understanding Multirack Cabling

You can connect multiple Oracle Big Data Appliance racks, or a combination of Oracle Big Data Appliance racks and Oracle Exadata Database Machine racks. Exadata full racks, half racks, and quarter racks must have a spine switch. If your Exadata rack does not have a spine switch, then be sure to have an Exadata Expansion Switch Kit on hand for the installation. Exadata eighth racks do not require a spine switch to connect to Oracle Big Data Appliance.


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 D-1](#).

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

Cabling the Switches in One Rack

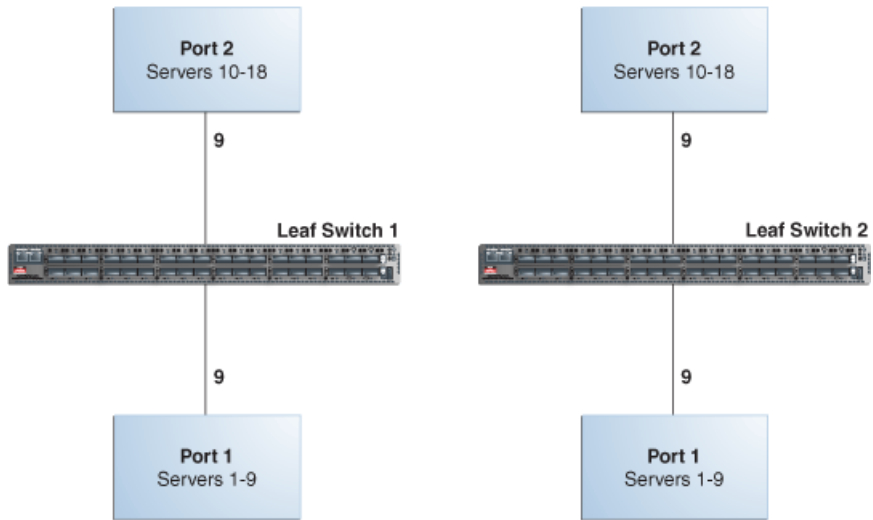
In a single rack, the two leaf switches are connected by seven cables. Each leaf switch has one cable to the spine switch. Nine servers connect to one leaf switch, and the other nine servers connect to the other leaf switch, as shown in [Figure E-1](#).

Figure E-1 Connections Between Spine Switch and Leaf Switches



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

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



Cabling the Switches in Two Racks


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

[Figure E-3](#) shows that each leaf switch in rack 1 connects to the following switches:

- Four connections to its internal spine switch
- Four connections to the spine switch in rack 2

The spine switch in rack 1 connects to the following switches:


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

Figure E-3 Connecting Switches Across Two Racks

Cabling the Switches in Up to Eight Racks

As the number of racks increases from two to eight, the pattern continues as shown in [Figure E-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 E-4 Connecting Switches Across Multiple Racks

Key to Cabling Table Abbreviations

The following abbreviations are used in the tables:

| Abbreviation | Description |
|--------------|--|
| R_n | Rack n , where n is the number of the rack, such as R1. |
| IB_n | Unit location in rack, where n is the number, such as IB3. |
| P_n | InfiniBand port n , where n is port number, such as P8A. |

Two-Rack Cabling

Table E-1 shows the cable connections for the first spine switch (R1-U1) when two racks are cabled together.

Note:

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

Table E-1 Leaf Switch Connections for the First Rack in a Two-Rack System

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

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

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

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

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

| Leaf Switch | Connection | Cable Length |
|----------------------|---|--------------|
| R2 IB2 within Rack 2 | R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B R2-IB2-P9A to R2-U1-P5B R2-IB2-P9B to R2-U1-P6B | 3 meters |
| R2 IB2 to Rack 1 | R2-IB2-P10A to R1-U1-P7B R2-IB2-P10B to R1-U1-P8B R2-IB2-P11A to R1-U1-P9B R2-IB2-P11B to R1-U1-P10B | 5 meters |

Three-Rack Cabling

Table E-3 shows the cable connections for the first spine switch (R1-U1) when three racks are cabled together.

Note:

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

Table E-3 Leaf Switch Connections for the First Rack in a Three-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|---|--------------|
| R1 IB3 within Rack 1 | R1-U24-P8A to R1-U1-P3A R1-U24-P8B to R1-U1-P4A R1-U24-P9A to R1-U1-P5A | 3 meters |
| R1 IB3 to Rack 2 | R1-U24-P9B to R2-U1-P6A R1-U24-P10A to R2-U1-P7A R1-U24-P10B to R2-U1-P8A | 5 meters |
| R1 U24 to Rack 3 | R1-U24-P11A to R3-U1-P9A R1-U24-P11B to R3-U1-P10A | 5 meters |
| R1 IB2 within Rack 1 | R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B R1-IB2-P9A to R1-U1-P5B | 3 meters |
| R1 IB2 to Rack 2 | R1-IB2-P9B to R2-U1-P6B R1-IB2-P10A to R2-U1-P7B R1-IB2-P10B to R2-U1-P8B | 5 meters |
| R1 IB2 to Rack 3 | R1-IB2-P11A to R3-U1-P9B R1-IB2-P11B to R3-U1-P10B | 5 meters |

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

Table E-4 Leaf Switch Connections for the Second Rack in a Three-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|---|--------------|
| R2 IB3 within Rack 2 | R2-U24-P8A to R2-U1-P3A R2-U24-P8B to R2-U1-P4A R2-U24-P9A to R2-U1-P5A | 3 meters |
| R2 IB3 to Rack 1 | R2-U24-P11A to R1-U1-P9A R2-U24-P11B to R1-U1-P10A | 5 meters |
| R2 IB3 to Rack 3 | R2-U24-P9B to R3-U1-P6A R2-U24-P10A to R3-U1-P7A R2-U24-P10B to R3-U1-P8A | 5 meters |
| R2 IB2 within Rack 2 | R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B R2-IB2-P9A to R2-U1-P5B | 3 meters |
| R2 IB2 to Rack 1 | R2-IB2-P11A to R1-U1-P9B R2-IB2-P11B to R1-U1-P10B | 5 meters |
| R2 IB2 to Rack 3 | R2-IB2-P9B to R3-U1-P6B R2-IB2-P10A to R3-U1-P7B R2-IB2-P10B to R3-U1-P8B | 5 meters |

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

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

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

Four-Rack Cabling

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

Note:

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

Table E-6 Leaf Switch Connections for the First Rack in a Four-Rack System

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

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

Table E-7 Leaf Switch Connections for the Second Rack in a Four-Rack System

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

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

| Leaf Switch | Connection | Cable Length |
|----------------------|---|--------------|
| R2 IB2 within Rack 2 | R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B | 3 meters |
| R2 IB2 to Rack 1 | R2-IB2-P11A to R1-U1-P9B R2-IB2-P11B to R1-U1-P10B | 5 meters |
| R2 IB2 to Rack 3 | R2-IB2-P9A to R3-U1-P5B R2-IB2-P9B to R3-U1-P6B | 5 meters |
| R2 IB2 to Rack 4 | R2-IB2-P10A to R4-U1-P7B R2-IB2-P10B to R4-U1-P8B | 5 meters |

Table E-8 shows the cable connections for the third spine switch (R3-U1) when four racks are cabled together.

Table E-8 Leaf Switch Connections for the Third Rack in a Four-Rack System

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

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

Table E-9 Leaf Switch Connections for the Fourth Rack in a Four-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|--|--------------|
| R4 IB3 within Rack 4 | R4-IB3-P8A to R4-U1-P3A R4-IB3-P8B to R4-U1-P4A | 3 meters |
| R4 IB3 to Rack 1 | R4-IB3-P9A to R1-U1-P5A R4-IB3-P9B to R1-U1-P6A | 10 meters |

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

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

Five-Rack Cabling

Table E–10 shows the cable connections for the first spine switch (R1-U1) when five racks are cabled together.

Note:

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

Table E–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 |

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

| Leaf Switch | Connection | Cable Length |
|------------------|---------------------------|--------------|
| R1 IB2 to Rack 5 | R1-IB2-P11B to R5-U1-P10B | 10 meters |

Table E–11 shows the cable connections for the second spine switch (R2-U1) when five racks are cabled together.

Table E–11 Leaf Switch Connections for the Second Rack in a Five-Rack System

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

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

Table E–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 |

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

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

Table E-13 Leaf Switch Connections for the Fourth Rack in a Five-Rack System

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

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

Table E-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 |

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

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

Six-Rack Cabling

Table E-15 shows the cable connections for the first spine switch (R1-U1) when cabling six racks.

Note:

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

Table E-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 E-16 shows the cable connections for the second spine switch (R2-U1) when six racks are cabled together.

Table E-16 Leaf Switch Connections for the Second Rack in a Six-Rack System

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

Table E-17 shows the cable connections for the third spine switch (R3-U1) when six racks are cabled together.

Table E-17 Leaf Switch Connections for the Third Rack in a Six-Rack System

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

Table E-18 shows the cable connections for the fourth spine switch (R4-U1) when six racks are cabled together.

Table E-18 Leaf Switch Connections for the Fourth Rack in a Six-Rack System

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

Table E-19 shows the cable connections for the fifth spine switch (R5-U1) when six racks are cabled together.

Table E-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-IB3-P10B to R2-U1-P8B | 10 meters |
| R5 IB2 to Rack 3 | R5-IB2-P11A to R3-U1-P9B | 5 meters |
| R5 IB2 to Rack 4 | R5-IB2-P11B to R4-U1-P10B | 5 meters |

Table E-19 (Cont.) Leaf Switch Connections for the Fifth Rack in a Six-Rack System

| Leaf Switch | Connection | Cable Length |
|------------------|--|--------------|
| R5 IB2 to Rack 6 | R5-IB2-P9A to R6-U1-P5B R5-IB2-P9B to R6-U1-P6B | 5 meters |

Table E-20 shows the cable connections for the sixth spine switch (R6-U1) when six racks are cabled together.

Table E-20 Leaf Switch Connections for the Sixth Rack in a Six-Rack System

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

Seven-Rack Cabling

Table E-21 shows the cable connections for the first spine switch (R1-U1) when seven racks are cabled together.

Note:

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

Table E-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 |

Table E–21 (Cont.) Leaf Switch Connections for the First Rack in a Seven-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|--|--------------|
| R1 IB3 to Rack 3 | R1-IB3-P9B to R3-U1-P6A | 5 meters |
| R1 IB3 to Rack 4 | R1-IB3-P10A to R4-U1-P7A | 10 meters |
| R1 IB3 to Rack 5 | R1-IB3-P10B to R5-U1-P8A | 10 meters |
| R1 IB3 to Rack 6 | R1-IB3-P11A to R6-U1-P9A | 10 meters |
| R1 IB3 to Rack 7 | R1-IB3-P11B to R7-U1-P10A | 10 meters |
| R1 IB2 within Rack 1 | R1-IB2-P8A to R1-U1-P3B R1-IB2-P8B to R1-U1-P4B | 3 meters |
| R1 IB2 to Rack 2 | R1-IB2-P9A to R2-U1-P5B | 5 meters |
| R1 IB2 to Rack 3 | R1-IB2-P9B to R3-U1-P6B | 5 meters |
| R1 IB2 to Rack 4 | R1-IB2-P10A to R4-U1-P7B | 10 meters |
| R1 IB2 to Rack 5 | R1-IB2-P10B to R5-U1-P8B | 10 meters |
| R1 IB2 to Rack 6 | R1-IB2-P11A to R6-U1-P9B | 10 meters |
| R1 IB2 to Rack 7 | R1-IB2-P11B to R7-U1-P10B | 10 meters |

Table E–22 shows the cable connections for the second spine switch (R2-U1) when cabling seven racks.

Table E–22 Leaf Switch Connections for the Second Rack in a Seven-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|--|--------------|
| R2 IB3 within Rack 2 | R2-IB3-P8A to R2-U1-P3A R2-IB3-P8B to R2-U1-P4A | 3 meters |
| R2 IB3 to Rack 1 | R2-IB3-P11B to R1-U1-P10A | 5 meters |
| R2 IB3 to Rack 3 | R2-IB3-P9A to R3-U1-P5A | 5 meters |
| R2 IB3 to Rack 4 | R2-IB3-P9B to R4-U1-P6A | 5 meters |
| R2 IB3 to Rack 5 | R2-IB3-P10A to R5-U1-P7A | 10 meters |
| R2 IB3 to Rack 6 | R2-IB3-P10B to R6-U1-P8A | 10 meters |
| R2 IB3 to Rack 7 | R2-IB3-P11A to R7-U1-P9A | 10 meters |
| R2 IB2 within Rack 2 | R2-IB2-P8A to R2-U1-P3B R2-IB2-P8B to R2-U1-P4B | 3 meters |
| R2 IB2 to Rack 1 | R2-IB2-P11B to R1-U1-P10B | 5 meters |
| R2 IB2 to Rack 3 | R2-IB2-P9A to R3-U1-P5B | 5 meters |
| R2 IB2 to Rack 4 | R2-IB2-P9B to R4-U1-P6B | 5 meters |
| R2 IB2 to Rack 5 | R2-IB2-P10A to R5-U1-P7B | 10 meters |
| R2 IB2 to Rack 6 | R2-IB2-P10B to R6-U1-P8B | 10 meters |
| R2 IB2 to Rack 7 | R2-IB2-P11A to R7-U1-P9B | 10 meters |

Table E–23 shows the cable connections for the third spine switch (R3-U1) when seven racks are cabled together.

Table E-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 E-24 shows the cable connections for the fourth spine switch (R4-U1) when seven racks are cabled together.

Table E-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 |
| R4 IB2 to Rack 7 | R4-IB2-P10A to R7-U1-P7B | 10 meters |

Table E-25 shows the cable connections for the fifth spine switch (R5-U1) when seven racks are cabled together.

Table E-25 Leaf Switch Connections for the Fifth Rack in a Seven-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|--|--------------|
| R5 IB3 within Rack 5 | R5-IB3-P8A to R5-U1-P3A R5-IB3-P8B to R5-U1-P4A | 3 meters |
| R5 IB3 to Rack 1 | R5-IB3-P10A to R1-U1-P7A | 10 meters |
| R5 IB3 to Rack 2 | R5-IB3-P10B to R2-U1-P8A | 10 meters |
| R5 IB3 to Rack 3 | R5-IB3-P11A to R3-U1-P9A | 5 meters |
| R5 IB3 to Rack 4 | R5-IB3-P11B to R4-U1-P10A | 5 meters |
| R5 IB3 to Rack 6 | R5-IB3-P9A to R6-U1-P5A | 5 meters |
| R5 IB3 to Rack 7 | R5-IB3-P9B to R7-U1-P6A | 5 meters |
| R5 IB2 within Rack 5 | R5-IB2-P8A to R5-U1-P3B R5-IB2-P8B to R5-U1-P4B | 3 meters |
| R5 IB2 to Rack 1 | R5-IB2-P10A to R1-U1-P7B | 10 meters |
| R5 IB2 to Rack 2 | R5-IB3-P10B to R2-U1-P8B | 10 meters |
| R5 IB2 to Rack 3 | R5-IB2-P11A to R3-U1-P9B | 5 meters |
| R5 IB2 to Rack 4 | R5-IB2-P11B to R4-U1-P10B | 5 meters |
| R5 IB2 to Rack 6 | R5-IB2-P9A to R6-U1-P5B | 5 meters |
| R5 IB2 to Rack 7 | R5-IB2-P9B to R7-U1-P6B | 5 meters |

Table E-26 shows the cable connections for the sixth spine switch (R6-U1) when seven racks are cabled together.

Table E-26 Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|--|--------------|
| R6 IB3 within Rack 6 | R6-IB3-P8A to R6-U1-P3A R6-IB3-P8B to R6-U1-P4A | 3 meters |
| R6 IB3 to Rack 1 | R6-IB3-P9B to R1-U1-P6A | 10 meters |
| R6 IB3 to Rack 2 | R6-IB3-P10A to R2-U1-P7A | 10 meters |
| R6 IB3 to Rack 3 | R6-IB3-P10B to R3-U1-P8A | 5 meters |
| R6 IB3 to Rack 4 | R6-IB3-P11A to R4-U1-P9A | 5 meters |
| R6 IB3 to Rack 5 | R6-IB3-P11B to R5-U1-P10A | 5 meters |
| R6 IB3 to Rack 7 | R6-IB3-P9A to R7-U1-P5A | 5 meters |
| R6 IB2 within Rack 6 | R6-IB2-P8A to R6-U1-P3B R6-IB2-P8B to R6-U1-P4B | 3 meters |
| R6 IB2 to Rack 1 | R6-IB2-P9B to R1-U1-P6B | 10 meters |
| R6 IB2 to Rack 2 | R6-IB3-P10A to R2-U1-P7B | 10 meters |
| R6 IB2 to Rack 3 | R6-IB2-P10B to R3-U1-P8B | 5 meters |
| R6 IB2 to Rack 4 | R6-IB2-P11A to R4-U1-P9B | 5 meters |
| R6 IB2 to Rack 5 | R6-IB2-P11B to R5-U1-P10B | 5 meters |

Table E–26 (Cont.) Leaf Switch Connections for the Sixth Rack in a Seven-Rack System

| Leaf Switch | Connection | Cable Length |
|------------------|-------------------------|--------------|
| R6 IB2 to Rack 7 | R6-IB3-P9A to R7-U1-P5B | 5 meters |

Table E–27 shows the cable connections for the seventh spine switch (R7-U1) when seven racks are cabled together.

Table E–27 Leaf Switch Connections for the Seventh Rack in a Seven-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|--|--------------|
| R7 IB3 within Rack 7 | R7-IB3-P8A to R7-U1-P3A R7-IB3-P8B to R7-U1-P4A | 3 meters |
| R7 IB3 to Rack 1 | R7-IB3-P9A to R1-U1-P5A | 10 meters |
| R7 IB3 to Rack 2 | R7-IB3-P9B to R2-U1-P6A | 10 meters |
| R7 IB3 to Rack 3 | R7-IB3-P10A to R3-U1-P7A | 10 meters |
| R7 IB3 to Rack 4 | R7-IB3-P10B to R4-U1-P8A | 10 meters |
| R7 IB3 to Rack 5 | R7-IB3-P11A to R5-U1-P9A | 5 meters |
| R7 IB3 to Rack 6 | R7-IB3-P11B to R6-U1-P10A | 5 meters |
| R7 IB2 within Rack 7 | R7-IB2-P8A to R7-U1-P3B R7-IB2-P8B to R7-U1-P4B | 3 meters |
| R7 IB2 to Rack 1 | R7-IB2-P9A to R1-U1-P5B | 10 meters |
| R7 IB2 to Rack 2 | R7-IB3-P9B to R2-U1-P6B | 10 meters |
| R7 IB2 to Rack 3 | R7-IB2-P10A to R3-U1-P7B | 10 meters |
| R7 IB2 to Rack 4 | R7-IB2-P10B to R4-U1-P8B | 10 meters |
| R7 IB2 to Rack 5 | R7-IB2-P11A to R5-U1-P9B | 5 meters |
| R7 IB2 to Rack 6 | R7-IB3-P11B to R6-U1-P10B | 5 meters |

Eight-Rack Cabling

Table E–28 shows the cable connections for the first spine switch (R1-U1) when cabling eight racks.

Note:

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

Table E–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 |

Table E–28 (Cont.) Leaf Switch Connections for the First Rack in an Eight-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|---------------------------|--------------|
| R1 IB3 to Rack 5 | R1-IB3-P10A to R5-U1-P7A | 10 meters |
| R1 IB3 to Rack 6 | R1-IB3-P10B to R6-U1-P8A | 10 meters |
| R1 IB3 to Rack 7 | R1-IB3-P11A to R7-U1-P9A | 10 meters |
| R1 IB3 to Rack 8 | R1-IB3-P11B to R8-U1-P10A | 10 meters |
| R1 IB2 within Rack 1 | R1-IB2-P8A to R1-U1-P3B | 3 meters |
| R1 IB2 to Rack 2 | R1-IB2-P8B to R2-U1-P4B | 5 meters |
| R1 IB2 to Rack 3 | R1-IB2-P9A to R3-U1-P5B | 5 meters |
| R1 IB2 to Rack 4 | R1-IB2-P9B to R4-U1-P6B | 10 meters |
| R1 IB2 to Rack 5 | R1-IB2-P10A to R5-U1-P7B | 10 meters |
| R1 IB2 to Rack 6 | R1-IB2-P10B to R6-U1-P8B | 10 meters |
| R1 IB2 to Rack 7 | R1-IB2-P11A to R7-U1-P8B | 10 meters |
| R1 IB2 to Rack 8 | R1-IB2-P11B to R8-U1-P10B | 10 meters |

Table E–29 shows the cable connections for the second spine switch (R2-U1) when eight racks are cabled together.

Table E–29 Leaf Switch Connections for the Second Rack in an Eight-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|---------------------------|--------------|
| R2 IB3 within Rack 2 | R2-IB3-P8A to R2-U1-P3A | 3 meters |
| R2 IB3 to Rack 1 | R2-IB3-P11B to R1-U1-P10A | 5 meters |
| R2 IB3 to Rack 3 | R2-IB3-P8B to R3-U1-P4A | 5 meters |
| R2 IB3 to Rack 4 | R2-IB3-P9A to R4-U1-P5A | 5 meters |
| R2 IB3 to Rack 5 | R2-IB3-P9B to R5-U1-P6A | 10 meters |
| R2 IB3 to Rack 6 | R2-IB3-P10A to R6-U1-P7A | 10 meters |
| R2 IB3 to Rack 7 | R2-IB3-P10B to R7-U1-P8A | 10 meters |
| R2 IB3 to Rack 8 | R2-IB3-P11A to R8-U1-P9A | 10 meters |
| R2 IB2 within Rack 2 | R2-IB2-P8A to R2-U1-P3B | 3 meters |
| R2 IB2 to Rack 1 | R2-IB2-P11B to R1-U1-P10B | 5 meters |
| R2 IB2 to Rack 3 | R2-IB2-P8B to R3-U1-P4B | 5 meters |
| R2 IB2 to Rack 4 | R2-IB2-P9A to R4-U1-P5B | 5 meters |
| R2 IB2 to Rack 5 | R2-IB2-P9B to R5-U1-P6B | 10 meters |
| R2 IB2 to Rack 6 | R2-IB2-P10A to R6-U1-P7B | 10 meters |
| R2 IB2 to Rack 7 | R2-IB2-P10B to R7-U1-P8B | 10 meters |
| R2 IB2 to Rack 8 | R2-IB2-P11A to R8-U1-P9B | 10 meters |

Table E–30 shows the cable connections for the third spine switch (R3-U1) when eight racks are cabled together.

Table E-30 Leaf Switch Connections for the Third Rack in an Eight-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|---------------------------|--------------|
| R3 IB3 within Rack 3 | R3-IB3-P8A to R3-U1-P3A | 3 meters |
| R3 IB3 to Rack 1 | R3-IB3-P11A to R1-U1-P9A | 5 meters |
| R3 IB3 to Rack 2 | R3-IB3-P11B to R2-U1-P10A | 5 meters |
| R3 IB3 to Rack 4 | R3-IB3-P8B to R4-U1-P4A | 5 meters |
| R3 IB3 to Rack 5 | R3-IB3-P9A to R5-U1-P5A | 5 meters |
| R3 IB3 to Rack 6 | R3-IB3-P9B to R6-U1-P6A | 5 meters |
| R3 IB3 to Rack 7 | R3-IB3-P10A to R7-U1-P7A | 10 meters |
| R3 IB3 to Rack 8 | R3-IB3-P10B to R8-U1-P8A | 10 meters |
| R3 IB2 within Rack 3 | R3-IB2-P8A to R3-U1-P3B | 3 meters |
| R3 IB2 to Rack 1 | R3-IB2-P11A to R1-U1-P9B | 5 meters |
| R3 IB2 to Rack 2 | R3-IB2-P11B to R2-U1-P10B | 5 meters |
| R3 IB2 to Rack 4 | R3-IB2-P8B to R4-U1-P4B | 5 meters |
| R3 IB2 to Rack 5 | R3-IB2-P9A to R5-U1-P5B | 5 meters |
| R3 IB2 to Rack 6 | R3-IB2-P9B to R6-U1-P6B | 5 meters |
| R3 IB2 to Rack 7 | R3-IB2-P10A to R7-U1-P7B | 10 meters |
| R3 IB2 to Rack 8 | R3-IB2-P10B to R8-U1-P8B | 10 meters |

Table E-31 shows the cable connections for the fourth spine switch (R4-U1) when eight racks are cabled together.

Table E-31 Leaf Switch Connections for the Fourth Rack in an Eight-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|---------------------------|--------------|
| R4 IB3 within Rack 4 | R4-IB3-P8A to R4-U1-P3A | 3 meters |
| R4 IB3 to Rack 1 | R4-IB3-P10B to R1-U1-P8A | 10 meters |
| R4 IB3 to Rack 2 | R4-IB3-P11A to R2-U1-P9A | 5 meters |
| R4 IB3 to Rack 3 | R4-IB3-P11B to R3-U1-P10A | 5 meters |
| R4 IB3 to Rack 5 | R4-IB3-P8B to R5-U1-P4A | 5 meters |
| R4 IB3 to Rack 6 | R4-IB3-P9A to R6-U1-P5A | 5 meters |
| R4 IB3 to Rack 7 | R4-IB3-P9B to R7-U1-P6A | 10 meters |
| R4 IB3 to Rack 8 | R4-IB3-P10A to R8-U1-P7A | 10 meters |
| R4 IB2 within Rack 4 | R4-IB2-P8A to R4-U1-P3B | 3 meters |
| R4 IB2 to Rack 1 | R4-IB2-P10B to R1-U1-P8B | 10 meters |
| R4 IB2 to Rack 2 | R4-IB2-P11A to R2-U1-P9B | 5 meters |
| R4 IB2 to Rack 3 | R4-IB2-P11B to R3-U1-P10B | 5 meters |
| R4 IB2 to Rack 5 | R4-IB2-P8B to R5-U1-P4B | 5 meters |
| R4 IB2 to Rack 6 | R4-IB2-P9A to R6-U1-P5B | 5 meters |
| 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 E-32 shows the cable connections for the fifth spine switch (R5-U1) when eight racks are cabled together.

Table E-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-IB3-P10B to R2-U1-P8B | 10 meters |
| R5 IB2 to Rack 3 | R5-IB2-P11A to R3-U1-P9B | 5 meters |
| R5 IB2 to Rack 4 | R5-IB2-P11B to R4-U1-P10B | 5 meters |
| R5 IB2 to Rack 6 | R5-IB2-P8B to R6-U1-P4B | 5 meters |
| R5 IB2 to Rack 7 | R5-IB2-P9A to R7-U1-P5B | 5 meters |
| R5 IB2 to Rack 8 | R5-IB2-P9B to R8-U1-P6B | 10 meters |

Table E-33 shows the cable connections for the sixth spine switch (R6-U1) when eight racks are cabled together.

Table E-33 Leaf Switch Connections for the Sixth Rack in an Eight-Rack System

| Leaf Switch | Connection | Cable Length |
|----------------------|---------------------------|--------------|
| R6 IB3 within Rack 6 | R6-IB3-P8A to R6-U1-P3A | 3 meters |
| R6 IB3 to Rack 1 | R6-IB3-P9B to R1-U1-P6A | 10 meters |
| R6 IB3 to Rack 2 | R6-IB3-P10A to R2-U1-P7A | 10 meters |
| R6 IB3 to Rack 3 | R6-IB3-P10B to R3-U1-P8A | 5 meters |
| R6 IB3 to Rack 4 | R6-IB3-P11A to R4-U1-P9A | 5 meters |
| R6 IB3 to Rack 5 | R6-IB3-P11B to R5-U1-P10A | 5 meters |
| R6 IB3 to Rack 7 | R6-IB3-P8B to R7-U1-P4A | 5 meters |
| R6 IB3 to Rack 8 | R6-IB3-P9A to R8-U1-P5A | 5 meters |
| R6 IB2 within Rack 6 | R6-IB2-P8A to R6-U1-P3B | 3 meters |
| R6 IB2 to Rack 1 | R6-IB2-P9B to R1-U1-P6B | 10 meters |
| R6 IB2 to Rack 2 | R6-IB3-P10A to R2-U1-P7B | 10 meters |
| R6 IB2 to Rack 3 | R6-IB2-P10B to R3-U1-P8B | 5 meters |
| R6 IB2 to Rack 4 | R6-IB2-P11A to R4-U1-P9B | 5 meters |
| R6 IB2 to Rack 5 | R6-IB2-P11B to R5-U1-P10B | 5 meters |

Table E-33 (Cont.) Leaf Switch Connections for the Sixth Rack in an Eight-Rack System

| Leaf Switch | Connection | Cable Length |
|------------------|-------------------------|--------------|
| R6 IB2 to Rack 7 | R6-IB3-P8B to R7-U1-P4B | 5 meters |
| R6 IB2 to Rack 8 | R6-IB2-P9A to R8-U1-P5B | 5 meters |

Table E-34 shows the cable connections for the seventh spine switch (R7-U1) when eight racks are cabled together.

Table E-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-IB3-P9B to R2-U1-P6B | 10 meters |
| R7 IB2 to Rack 3 | R7-IB2-P10A to R3-U1-P7B | 10 meters |
| R7 IB2 to Rack 4 | R7-IB2-P10B to R4-U1-P8B | 10 meters |
| R7 IB2 to Rack 5 | R7-IB2-P11A to R5-U1-P9B | 5 meters |
| R7 IB2 to Rack 6 | R7-IB3-P11B to R6-U1-P10B | 5 meters |
| R7 IB2 to Rack 8 | R7-IB3-P8B to R8-U1-P4B | 5 meters |

Table E-35 shows the cable connections for the eighth spine switch (R8-U1) when eight racks are cabled together.

Table E-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 | 5 meters |
| R8 IB3 to Rack 6 | R8-IB3-P11A to R6-U1-P9A | 5 meters |
| R8 IB3 to Rack 7 | R8-IB3-P11B to R7-U1-P10A | 5 meters |
| R8 IB2 within Rack 8 | R8-IB2-P8A to R8-U1-P3B | 3 meters |

Table E-35 (Cont.) Leaf Switch Connections for the Eighth Rack in an Eight-Rack

| Leaf Switch | Connection | Cable Length |
|--------------------|--------------------------|---------------------|
| R8 IB2 to Rack 1 | R8-IB2-P8B to R1-U1-P4B | 10 meters |
| R8 IB2 to Rack 2 | R8-IB3-P9A to R2-U1-P5B | 10 meters |
| R8 IB2 to Rack 3 | R8-IB2-P9B to R3-U1-P6B | 10 meters |
| R8 IB2 to Rack 4 | R8-IB2-P10A to R4-U1-P7B | 10 meters |
| R8 IB2 to Rack 5 | R8-IB2-P10B to R5-U1-P8B | 5 meters |
| R8 IB2 to Rack 6 | R8-IB3-P11A to R6-U1-P9B | 5 meters |
| R8 IB2 to Rack 7 | R8-IB3-P1B to R7-U1-P10B | 5 meters |

A

- AC breaker panels, 2-9
- AC power, location, 2-9
- access requirements, 2-4
- access route requirements, 2-3, 2-4
- accessing servers from KVM, 7-3
- acoustic levels, 2-1, 2-2
- add subcommand, 9-18
- adding racks, 9-3
- administrative network, 3-2
 - cabling requirements, 2-12
 - configuration, 7-20
 - connections, 2-12
 - ports, 3-1
- administrative network diagram, 13-2
- administrative utilities, 1-3
- Africa
 - PDU specifications, 2-7, 2-8
- air conditioning, 2-10, 2-14, 6-2
- airflow, 2-2, 2-9
- airflow requirements, 2-10
- alerts, A-1
 - generating, 13-1
 - sent by Cloudera Manager, 4-9
- Allen wrench keys, 6-8
- altitude requirements, 2-9
- ambient relative humidity range, 2-10
- ambient temperature range, 2-10, 10-1
- Americas
 - PDU specifications, 2-5, 2-6
- Apache OpenOffice Spreadsheet, 4-2
- Asia Pacific
 - PDU specifications, 2-7, 2-8
- ASR
 - about, 5-1
 - activating assets, 5-4
 - configuration, 5-2
 - troubleshooting, 5-6
 - See also* Auto Service Request
- ASR Manager
 - port number, 4-8
- ASR validation
 - generating test events, 5-5
- Auto Service Request
 - alerts, A-1

- installation on Oracle Big Data Appliance, 9-7
- installing, 4-8, 9-17
- See also* ASR

Auto Service Request ports, 3-7

B

- base image reinstallation, 9-9
- base image utilities, 9-19
- batteries
 - life expectancy, B-2
 - replacement policy, B-2
- BDA server offline repair category, B-2
- BDA_REBOOT files, 6-14
- bdacheckcluster utility, 11-2
- bdacheckhw command, 7-19, 10-28
- bdacheckhw utility, 11-3
- bdacheckib utility, 11-4
- bda-checkip.out file, 3-9
- bdachecknet command, 11-6
- bdachecksw command, 7-20, 10-29
- bdachecksw utility, 11-7
- BdaDeploy.json file, 4-3, 7-20
- bdadiag utility, 11-8
- BdaExpansion.json file, 4-3
- bdagetclientconfig utility, 11-11
- bdaid utility, 11-12
- bdaimagevalidate utility, 11-13
- bda-install-template.pdf file, 4-4
- BDAMammoth directory, 9-13, 9-18
- BDA-O repair category, B-2
- bdaredoclientnet utility, 11-14
- bdaserials utility, 11-16
- bdaswapoff utility, 11-17
- bdaswapon utility, 11-18
- bdauupdatefw utility, 11-19
- Big Data Appliance
 - See* Oracle Big Data Appliance
- Big Data Connectors
 - See* Oracle Big Data Connectors
- black cables, C-1, D-1
- blue cables, C-1, D-1
- bolts, 6-7
- bondeth0 interface
 - host names, 4-5
 - network connections, 3-1

- bondib0 interface
 - description, 3-2
 - host names, 4-5
 - maintenance, 10-25
- bondib0 network
 - verifying operations, 10-27
- breaker panels, 2-9

C

- cable tie,
 - removing, 6-11
- cable-routing floor cutout, 6-8
- cables
 - color coding, C-1, D-1
 - spares, 1-3
 - splitter, 3-5
- cabling
 - client network, 3-4
 - eight racks, E-19
 - existing network requirements, 2-12
 - five racks, E-9
 - four racks, E-7
 - multirack, 8-2, E-1
 - seven racks, E-15
 - six racks, E-12
 - spare parts, B-15
 - three racks, E-5
 - two racks, E-4
- cage nuts, 6-3
- casters, 6-5
- cautions, 10-3
- CDH
 - cluster configuration, 9-4
 - health checks, 11-2
 - installing, 9-17
- CDH cluster logs, 11-8
- CDH ports, 3-7
- CDH services, starting, 9-17
- ceiling height requirements, 2-3
- checking firmware, 7-20
- checking network configuration, 11-6
- checking the server software, 11-7
- circuit breakers, 2-9
 - general requirements, 2-9
- circuit breakers for PDUs, 6-13
- Cisco switch
 - replacement parts, B-14
 - See* Ethernet switch
- CleanupInstall step, 9-17
- CleanupSSHroot step, 9-18
- CLI, Oracle ILOM, 13-2
- client access network
 - cabling requirements, 2-12
 - description, 3-2
 - requirements, 2-12
- client configuration distribution, 11-11
- client network
 - bandwidth, 3-5
 - cabling, 3-4

- clock
 - configuring Ethernet switch, 7-8
 - configuring InfiniBand switches, 7-13
- Cloudera Manager, 4-8, 4-9
- Cloudera's Distribution including Apache Hadoop
 - See* CDH
- cluster names, 4-5
- collecting diagnostics, 11-8
- color coding, cables, C-1, D-1
- commands
 - remote execution across servers, 1-4, 12-1
- component failure
 - environmental causes, 2-9
- components, location in the rack, C-2, D-2
- configuration
 - disk controller, 10-5
 - physical drives, 10-6
 - virtual drive, 10-6
- configuration files
 - generating, 4-2
 - location, 4-3
 - ntp.conf file, 10-31
- configuration utility, 4-1
- configuration worksheets, 2-13, 4-3
- configure terminal command, 7-5
- configuring a new disk, 10-11
- configuring ASR, 5-2
- configuring Ethernet switch, 7-4, 7-5
- configuring KVM, 7-1, 7-3
- configuring servers, 7-18
- connecting KVM to the network, 7-2
- connectors
 - See* Oracle Big Data Connectors
- cooling requirements, 2-1, 2-2
- copying files, 12-2
- CopyLicenseFiles step, 9-16
- CopySoftwareSource step, 9-16
- corrosion problems, 2-10
- CreateLogicalVolumes step, 9-16
- CreateUsers step, 9-16

D

- data center switch configuration, 2-13
- dba group, 4-9, 9-16
- dcli utility
 - description, 1-4
 - examples, 12-5
 - option descriptions, 12-4
 - return values, 12-5
 - syntax, 12-3
 - target servers, 12-2
 - using, 12-1
- device mappings
 - restoring original, 10-13
- diagnostics collection, 11-8
- dimensions
 - rack, 2-3
 - shipping, 2-3
- disablesm command, 7-15, 8-2, 8-3, 8-4, 10-27

- disk controller configuration, 10-5
- disk controllers, 10-5
- disk drives
 - spares, 1-3, B-6
- disks
 - repairing, 10-5
 - replacing, 10-6
 - software configuration, 10-11
- DNS network registration, 2-13
- DNS server requirement, 3-2, 4-6
- door dimensions, 2-4
- downgrading the factory image, 9-4

E

- earth ground cable, 6-9
- electric receptacles, 2-5
- electrical access, 2-4
- electrical requirements, 2-4, 6-2, 6-13
- electrical storms, 2-4
- electronic noise, 6-2
- electrostatic discharge, 2-10
- elevator requirements, 2-4
- email accounts
 - My Oracle Support, 5-1
- email configuration, 4-9
- emergency shutdown, 10-3
- enablesm command, 7-15, 8-2, 8-3, 8-4
- env_test command, 7-14
- environmental specifications, 2-1, 2-2
- error messages, A-1
- eth0 interface, 3-2
- Ethernet bridge ports, 11-24
- Ethernet cabling, 2-12, 3-4
- Ethernet interface, 3-2
- Ethernet over InfiniBand protocol., 3-4
- Ethernet ports, 2-2, 2-3, 3-1
- Ethernet switch
 - configuration, 7-4
 - configuring host name, 7-6
 - IP routing, 7-5
 - port speed, 7-5
 - replacement parts, B-14
 - replacement switch, 1-4
 - setting IP address, 4-7
 - setting remote password, 7-6
 - spare parts, 1-3
 - time settings, 7-8
 - validating and saving configuration, 7-10
- Europe
 - PDU specifications, 2-7, 2-8
- Exadata Database Machine
 - See Oracle Exadata Database Machine
- ext4 file system, 10-18
- external filer configuration, 9-16

F

- Fabric Management, 7-14
- factory image reinstallation, 9-9

- factory patches
 - installing, 9-16
- failover connections, 3-1, 3-4
- failover switch connections, 3-1, 3-2
- Fibre Channel over Ethernet, 1-6
- filler panels, 2-11
- firewall configuration, 2-13
- firmware updates, 11-19
- firmware upgrades
 - KVM switch, 7-3
 - PDUs, 7-17
 - restrictions, 1-5
- first server location, 7-21
- floor cut dimensions, 6-7
- floor load, 2-4
- floor tiles, 2-11
- flooring
 - raised, 2-4
 - raised pattern, 2-3
- forklift, 2-3
- FRUs, B-1

G

- gateway addresses
 - factory settings, 3-6
- gateway configuration
 - Ethernet switch, 7-5
- gateway switches
 - assigning to VNICs, 3-5
 - cabling, 3-4
 - configuring, 7-11
 - connecting to Oracle ILOM, 13-5
 - replacing, 10-26
 - using splitter cables, 3-5
 - See also InfiniBand switches
- getmaster command, 8-3, 8-4
- getting help, 11-12
- ground cable, 6-9
- grounding resistance, 6-1
- grounding to earth, 2-9
- groups, 9-16
- GRUB master boot records, 10-19

H

- Hadoop
 - cluster configuration, 9-4
 - installing, 9-17
- hadoop group, 9-16
- Hadoop ports, 3-7
- Hadoop services, starting, 9-17
- hardware diagnostics, 13-1
- hardware failures
 - automatic reporting, 5-1
- hardware profiles
 - acquiring from servers, 7-19
- hardware validation, 11-2, 11-3, 11-13
- HDFS diagnostics, 11-8
- HDFS disks

- formatting, 10-23
- HDFS partition, 10-18
- hdfs user, 9-16
- health checks
 - InfiniBand switches, 7-13
- high-voltage PDUs, 2-7, 2-8, 10-32, 10-33
- hosts
 - factory IP addresses, 3-6
- hot swap repair category, B-1
- HS repair category, B-1
- humidity range, 2-9, 2-10

I

- ibdiagnet command, 10-28
- iblinkinfo command, 11-21
- ibqueryerrors command, 10-28
- ibstatus command, 7-19
- ILOM
 - See* Oracle ILOM
- ILOM host names, 4-5
- ilom-admin
 - default password, 7-11
- ilom-admin user, 13-2
- ilom-operator user, 13-2
- image validation, 11-13
- image version
 - downgrading, 9-4
 - upgrading, 9-4
- imagehistory utility, 11-22
- imageinfo utility, 9-4, 11-23
- incline, maximum, 2-4
- InfiniBand configuration, 7-11
- InfiniBand network
 - cabling checks, 11-4
 - factory IP addresses, 3-6
 - host names, 4-5
 - verifying operations, 10-27
- InfiniBand network configuration, 7-20
- InfiniBand network interface card, 3-1
- InfiniBand private network, 3-2
- InfiniBand Subnet Managers, 10-29
- InfiniBand switches
 - checking status, 7-19
 - configuring, 7-11
 - health checks, 7-13
 - on Oracle Exadata Database Machine, 7-15
 - on Oracle Exalogic, 7-15
 - priority settings, 7-15
 - replacing, 10-26
 - setting IP addresses, 4-7
 - setting time zone, 7-13
 - spare parts, B-12, B-13
- infrastructure repair category, B-1
- installation errors, 9-4
- Installation Template, 4-4
- InstallBDASoftware step, 9-17
- InstallHadoop step, 9-17
- installing
 - base image, 9-9

- third-party software, 1-5
- installing patches, 9-16
- installing the software
 - See also* Mammoth Utility, 9-2
- International Electrotechnical Commission (IEC), 2-9
- IP address requirements, 2-2, 2-3
- IP addresses
 - custom settings, 4-6
 - factory settings, 3-6
- IP routing, 7-5
- ipmitool command, 7-19
- IR repair category, B-1

J

- Japan
 - PDU specifications, 2-5, 2-6
- jumper brackets, 6-3

K

- KVM switch
 - accessing servers, 7-3
 - cabling, D-2
 - configuration procedures, 7-1
 - connecting to the network, 7-2
 - firmware upgrades, 7-3
 - naming, 7-2
 - setting IP address, 4-7
 - starting, 7-2
- KVM switch, MAC address, 7-2

L

- leaf switches
 - configuring, 7-11
 - See also* InfiniBand switches
- LED color codes, B-3, B-4
- leveling bolts, 6-4
- leveling feet, 6-5, 6-6
- licenses, location, 9-16
- licensing restrictions, 1-4
- lift gate, 2-3
- lightning arresters, 2-4
- lights-out management, 13-1
- listlinkup command, 11-24
- loading dock, 2-3
- locating
 - first server, 7-21
- location in the rack, C-2, D-2
- locks, panel, 6-3
- log files
 - Mammoth Utility, 9-4
- logical volume for Oracle NoSQL Database, 9-16
- low-voltage PDUs, 2-5, 2-6, 10-32

M

- MAC address of KVM switch, 7-2
- MAC addresses, 11-16
- maintenance access requirements, 2-4

- makebdamage utility, 9-19
- Mammoth Reconfiguration Utility
 - syntax, 9-18
- Mammoth Utility
 - description, 9-2
 - execution steps, 9-14
 - location, 9-13
 - log files, 9-4
 - reasons for failure, 9-4
 - running, 9-2
 - synchronizing clocks, 7-13
 - syntax, 9-13
- mammoth-rack_name.params file, 4-4
- management network
 - cabling requirements, 2-12
 - configuration, 7-20
 - connecting the KVM, 7-2
 - connections, 2-12
 - factory IP addresses, 3-6
 - host names, 4-5
 - ports, 3-1
- management tools
 - Oracle ILOM, 13-2
- management utilities, 1-3
- mapred user, 9-16
- MegaCli64 command, 7-20, 10-5
- Microsoft Office Excel, 4-2
- Middle East
 - PDU specifications, 2-7, 2-8
- minimum ceiling height, 2-3
- monitoring PDUs, 10-32
- mount point
 - logical volume for Oracle NoSQL Database, 9-16
- mounting a device, 7-21
- mounting brackets, 6-4, 6-6
 - installation instructions, 6-7
- mounting tool, 6-3
- multiple-rack clusters
 - client network connections, 3-6
- multirack cabling
 - procedures, 8-1
 - three to eight, 8-3
 - two racks, 8-2
- multirack software configuration, 9-3
- My Oracle Support, 4-2
 - email account, 5-1
- MySQL Database installation, 9-16

N

- naming conventions, 4-5
- NET0, 3-2
- network
 - connecting the KVM, 7-2
 - network configuration, 7-20
 - network configuration worksheet, 4-5
 - network connections, 2-12, 7-24
 - reconfiguring 10GbE, 11-14
 - network diagram, 3-3, 13-2
 - network errors

- checking for, 3-9
- network interfaces
 - InfiniBand card, 3-1
- network operations
 - verifying InfiniBand, 10-27
- network registration, 2-13
- network requirements, 2-2, 2-3
- Network Time Protocol, 10-31
- network topology, 10-29
- network validation, 11-2, 11-6
- networks
 - administrative, 3-2
 - client access, 2-12, 3-2
 - descriptions, 3-2
 - factory IP addresses, 3-6
 - InfiniBand, 3-2
- networksetup scripts, 7-22, 7-25
- nodes
 - location in the rack, C-2, D-2
- noise, electronic, 6-2
- North America
 - PDU specifications, 2-5, 2-6
- NoSQL database
 - See Oracle NoSQL Database
- NTP server requirement, 3-2, 4-6
- NTP servers, 10-31
 - configuring, 7-9
 - setting on InfiniBand switches, 7-13
- ntp.conf file, 10-31
- nuts, 6-3

O

- ODI
 - See Oracle Data Integrator
- oinstall group, 4-9, 9-16
- operating cautions and warnings, 10-3
- operating system
 - configuring a new disk, 10-15
- operating system disks, 10-14
- operating system history, 11-22
- operating system image, 11-23
- Oracle Big Data Appliance
 - general description, 1-1
 - rack components, 1-1
 - server components, 1-2
- Oracle Big Data Appliance Configuration Worksheets, 2-13
- Oracle Big Data Connectors
 - installation, 9-4, 9-17
 - installing server components, 4-8
 - software installation, 9-4
- Oracle Data Integrator, 4-9
 - installing the agent, 4-8
- Oracle Enterprise Manager Cloud Control
 - installing support on Oracle Big Data Appliance, 9-8
- Oracle Exadata Database Machine, 4-9
 - InfiniBand firmware version, 7-15
 - network identifiers, 4-8

- Oracle Exalogic
 - InfiniBand firmware version, 7-15
- Oracle ILOM, 13-1
 - backing up configuration, 10-25
 - connecting to CLI, 13-3
 - connecting to web interface, 13-3
 - description, 13-1
 - Ethernet connections, 13-3
 - factory IP addresses, 3-6
 - opening serial connections, 13-4
 - restoring configuration, 10-26
 - user interfaces, 13-2
- Oracle ILOM port, 3-1
- Oracle Inventory Group, 4-9
- Oracle Linux, 1-1
- Oracle NoSQL Database, 9-16
 - creating logical volumes, 9-16
 - disk configuration, 10-20
 - installation, 4-8, 9-17
 - logical volume, 9-16
- Oracle NoSQL Database disks, 10-20
- Oracle Open Office Calc, 4-2
- Oracle Support, 11-12
- oracle user name, 4-9
- orange cables, C-1, D-1
- OS Watcher logs, 11-8

P

- pallet jack, 2-3
- panel locks, 6-3
- parallel execution of commands, 12-1
- partitioning a new disk, 10-15
- passwordless SSH, 11-25, 11-27, 12-1
- passwords (default)
 - Ethernet switch, 7-5
 - ilom-admin, 7-11
 - PDU metering unit, 7-17
 - servers, 7-18
- patches
 - installing, 9-16
- PatchFactoryImage step, 9-16
- path to data center, 2-3
- PDAs, location in the rack, C-2, D-2
- PDU configuration, 7-17
- PDU metering unit
 - default password, 7-17
- PDU requirements, 2-5
- PDU specifications
 - Americas, Japan, Taiwan, 2-5, 2-6
 - Europe, Middle East, Africa, Asia Pacific, 2-7, 2-8
- PDU thresholds, 10-32
- PDUs
 - circuit breakers, 6-13
 - firmware upgrade, 7-17
 - network connections, 7-17
 - powering on, 6-13
 - replacement parts, B-14
 - setting IP addresses, 4-7
 - single-phase, 6-2, 6-13

- perforated tiles, 2-11
- photocopiers, 6-2
- physical drive configuration, 10-6
- port assignments, 3-7
- port speed
 - Ethernet switch, 7-5
- ports
 - ASR Manager, 4-8
 - connecting to 40 GbE, 3-5
- power connectors, 2-5
- power cord length requirements, 2-9
- power cords, 2-9
 - connecting to rack, 6-10
 - removing cable ties, 6-11
 - routing, 6-11
- power LEDs
 - locations, 6-13
- power off switch, 10-3
- power outlet grounding, 2-9
- power receptacle location, 2-9
- power requirements, 2-1, 2-2, 2-4
- power supplies, uninterruptible, 2-9
- powering off
 - emergency, 10-3
 - sequence, 10-2
- powering on, 6-13, 10-2
- preinstallation tasks, 2-12
- preinstall-checkip script
 - description, 4-4
 - running, 3-9
- PreInstallChecks step, 9-15
- private network
 - verifying operations, 10-27
- problem reporting, 5-1
- Puppet utility, 9-15

Q

- QSFP ports, splitting, 3-5

R

- rack components, 1-1
- rack configuration file, 9-4
- rack dimensions, 2-3
- rack layout illustration, C-2, D-2
- rack names, 4-5
- rack weight, 2-1, 2-2
- racks
 - adding another, 8-1
 - connecting power cords, 6-10
 - inspecting, 6-10
 - installing
 - Oracle Big Data Appliance software, 9-2
 - moving instructions, 6-5
 - reinstalling base image, 9-9
- RAID arrays
 - repairing logical, 10-16
- RAID configuration, 7-20
- red cables, C-1, D-1

- reimagecluster utility, 9-19
- reimagerack utility, 9-20
- reimaging the servers, 9-9
- reinstalling the software, 9-2
- relative humidity, 2-10, 2-11
- remove subcommand, 9-18
- remove-root-ssh command, 11-25
- remove-root-ssh script, 12-2
- repair categories, B-1
- replacement parts, B-1
- replacing spine switch, 10-27
- restart validation, 6-14
- restrictions on use, 1-4
- root user
 - passwordless SSH, 12-1
- route requirements, 2-3
- running-config output, 7-10

S

- safety guidelines
 - installation, 6-1
- scp file copy utility, 12-2
- screwdrivers, 6-4
- screws, 6-3
- sensor monitoring, 13-1
- serial numbers, 11-16
- server components, 1-2
- server configuration, 7-18
- server diagnostics, 11-8
- server disks
 - replacing, 10-6
 - software configuration, 10-11
- server software validation, 11-7
- servers
 - connecting to Oracle ILOM, 13-4
 - default passwords, 7-18
 - factory IP addresses, 3-6
 - LED color codes, B-3, B-4
 - location in the rack, C-2, D-2
 - monitoring, 13-1
 - setting IP addresses, 4-6
 - spare parts, B-6
 - upgrading base image, 9-9
- service locations, restrictions on moving, 1-6
- services, starting, 9-17
- setsmpriority command, 7-15, 8-2, 8-3, 8-4
- SetupASR step, 9-17
- SetupInstall step, 9-14
- SetupMountPoints step, 9-16
- SetupMySQL step, 9-16
- SetupPuppet step, 9-15
- setup-root-ssh command, 11-27
- setup-root-ssh script, 10-31, 12-1
- shipping crate dimensions, 6-2
- shipping dimensions, 2-3
- shipping kit, 6-3
- shipping weight, 2-3
- show command, 7-10
- showunhealthy command, 7-14
- showusb command, 7-21
- showvlan command, 11-29
- showvnic command, 11-30
- shutdown command, 10-2
- single-phase PDUs, 2-5, 2-7, 6-2, 6-13, 10-32
- site readiness tasks, 2-12
- SMTP alerts, 5-1
- SMTP configuration, 4-9
- SNMP alerts, 5-1
- software configuration worksheet, 4-8
- software image
 - downgrading, 9-4
 - upgrading, 9-4
- software installation
 - multiple racks, 9-3
 - upgrading, 9-5
 - using Mammoth, 9-2
- software installation errors, 9-4
- software installation,
 - third-party, 1-5
- software port assignments, 3-7
- software profiles
 - acquiring from servers, 7-20
- software source code, third party, 9-16
- software validation, 11-2, 11-13
- sound levels, 2-1, 2-2
- South America
 - PDU specifications, 2-5, 2-6
- spare parts, 6-4
- spares kit, 1-3
- spine switch
 - configuring, 7-11
- spine switches
 - cabling, E-2
 - replacing, 10-26, 10-27
- splitter cables, 3-5
- spreadsheet
 - configuration utility, 4-1
- spreadsheet programs, 4-2
- spring nuts, 6-3
- SSH keys
 - checking rack distribution, 7-18
- SSH, removing passwordless, 11-25
- SSH, setting up passwordless, 11-27
- stabilizing feet, 6-5
- stabilizing the rack, 6-6
- StartHadoopServices step, 9-17
- starting KVM, 7-2
- state file, 9-4
- static discharges, 2-10
- static electricity, 6-2
- Subnet Manager Master, 8-4
- Subnet Managers, 10-29, 10-30
- subnet masks
 - factory settings, 3-6
- Sun Network 10 GbE Switch 72p, 1-5, 3-5
- Sun Network QDR InfiniBand Gateway
 - Switch, 10-30, 11-14, 11-24
- support accounts, 5-2
- support services, 11-12

- swap partition
 - restoring, 10-19
- swapping, 10-9, 11-17, 11-18
- switches, C-2, D-2
 - monitoring, 13-1
 - naming conventions, 4-6
 - replacement parts, B-14
 - replacing, 10-26
 - setting host names, 7-12
 - setting IP addresses, 4-7
 - spare parts, B-12, B-13
 - third-party, 1-6

T

- Taiwan
 - PDU specifications, 2-5, 2-6
- telnet access
 - configuring Ethernet switch, 7-6
- temperature
 - acceptable range, 2-10
 - operating range, 2-9
- third-party licenses, location, 9-16
- third-party software, 1-5
- third-party software source code, 9-16
- third-party switches
 - restrictions on use, 1-6
- three-phase PDUs, 2-6, 2-8, 10-32, 10-33
- thresholds for PDUs, 10-32
- time zone
 - configuring Ethernet switch, 7-8
 - configuring InfiniBand switches, 7-13
 - setting summer time, 7-8
- top clearance (rack), 2-3

U

- unpacking instructions, 6-4
- upgrade history, 11-22
- upgrading base image, 9-9
- upgrading the factory image, 9-4
- upgrading the software, 9-5
- UPS, 2-9
- USB devices
 - restrictions on use, 1-6
- USB flash drive, 7-21
- users, 9-16

V

- validating a cluster, 11-2
- validating hardware and software, 11-13
- validating network configuration, 11-6
- validating the server software, 11-7
- validation checks after reboot, 6-14
- ventilation requirements, 2-10
- verifying
 - database server configuration, 10-5
- version command, 7-14
- virtual drive configuration, 10-6
- VLAN configuration, 11-29

- VLANs
 - configuration example, 7-5
 - default configuration, 3-1
- VNIC configuration, 11-30
- VNICs
 - assigning to gateway switches, 3-5
 - default configuration, 3-1
 - re-creating, 11-14

W

- warnings, 10-3
- warranty restrictions, 1-4
- washers, 6-7
- weight
 - rack, 2-1, 2-2
 - shipping, 2-3
- wrenches, 6-3, 6-8
- wrist strap, 6-4